

Wiley Finance Series

Advanced Financial Risk Management

SECOND EDITION

*Tools and Techniques for
Integrated Credit Risk and
Interest Rate Risk Management*

DONALD R. VAN DEVENTER
KENJI IMAI • MARK MESLER

Advanced Financial Risk Management

Founded in 1807, John Wiley & Sons is the oldest independent publishing company in the United States. With offices in North America, Europe, Australia, and Asia, Wiley is globally committed to developing and marketing print and electronic products and services for our customers' professional and personal knowledge and understanding.

The Wiley Finance series contains books written specifically for finance and investment professionals as well as sophisticated individual investors and their financial advisors. Book topics range from portfolio management to e-commerce, risk management, financial engineering, valuation, and financial instrument analysis, as well as much more.

For a list of available titles, visit our Web site at www.WileyFinance.com.

Advanced Financial Risk Management

Second Edition

*Tools and Techniques for Integrated
Credit Risk and Interest Rate
Risk Management*

DONALD R. VAN DEVENTER
KENJI IMAI
MARK MESLER

WILEY

Cover Design: Leiva-Sposato.
Cover Image: © hudiemm/iStockphoto.

Copyright © 2013 by John Wiley & Sons Singapore Pte. Ltd.

Published by John Wiley & Sons Singapore Pte. Ltd.
1 Fusionopolis Walk, #07-01, Solaris South Tower, Singapore 138628

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as expressly permitted by law, without either the prior written permission of the Publisher, or authorization through payment of the appropriate photocopy fee to the Copyright Clearance Center. Requests for permission should be addressed to the Publisher, John Wiley & Sons Singapore Pte. Ltd., 1 Fusionopolis Walk, #07-01, Solaris South Tower, Singapore 138628, tel: 65-6643-8000, fax: 65-6643-8008, e-mail: enquiry@wiley.com.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor the author shall be liable for any damages arising herefrom.

Other Wiley Editorial Offices

John Wiley & Sons, 111 River Street, Hoboken, NJ 07030, USA
John Wiley & Sons, The Atrium, Southern Gate, Chichester, West Sussex, P019 8SQ, United Kingdom
John Wiley & Sons (Canada) Ltd., 5353 Dundas Street West, Suite 400, Toronto, Ontario, M9B 6H8, Canada
John Wiley & Sons Australia Ltd., 42 McDougall Street, Milton, Queensland 4064, Australia
Wiley-VCH, Boschstrasse 12, D-69469 Weinheim, Germany

Library of Congress Cataloging-in-Publication Data

ISBN 978-1-118-27854-3 (Hardcover)
ISBN 978-1-118-27857-4 (ePDF)
ISBN 978-1-118-27856-7 (Mobi)
ISBN 978-1-118-27855-0 (ePub)

Typeset in 10/12pt Sabon-Roman by MPS Limited, Chennai, India.

Printed in Singapore by Ho Printing.
10 9 8 7 6 5 4 3 2 1

For Yasuko, Tomoki, and Anna.
K. I.

To all of our colleagues at Kamakura.
M. M.

For Ayako, Ella, Ai, and Yuu.
D. v. D.

Contents

Introduction: Wall Street Lessons from Bubbles	xxiii
Key Fallacies in Risk Management	xxiii
Selected Events in the Credit Crisis	xxviii

PART ONE

Risk Management: Definitions and Objectives

CHAPTER 1

A Risk Management Synthesis: Market Risk, Credit Risk, Liquidity Risk, and Asset and Liability Management	3
Risk Management: Definitions and Objectives	6
Advances in Integrated Risk Management and Institutional Barriers to Progress	8
Measuring the Trade-Offs between Risk and Return	11
When Bad Things Happen to Good People	11
U.S. Savings and Loan Crisis	12
Long-Term Capital Management	13
The 2006–2011 Credit Crisis	13
A Thousand Cuts	13

CHAPTER 2

Risk, Return, Performance Measurement, and Capital Regulation	15
Practical Quantification of Risk	15
Perils and Pitfalls in the Measurement of Risk: The Impact of Selection Bias	16
Biases in Return vs. a Relative Benchmark	17
Historical Value at Risk: Selection Bias Again	18
Monte Carlo–Based Value at Risk	19
Expected Losses on Tranches of Collateralized Debt Obligations	19
Measuring Return: Market vs. Accounting Returns	20
Introduction to Transfer Pricing: Extracting Interest Rate Risk in a Financial Accounting Context	20
Bank of America, 1973–1979	21
First Interstate, 1982–1987	24
Performance Measurement and Capital Regulation	26
Perspectives on Measuring Risk: One Source of Risk or Many Sources of Risk?	26
Interest Rate Risk Management Evolution	27
Equity Risk Management Evolution	28

Option Risk Management Evolution	28
Credit Risk Management Evolution	28
Managing Risk and Strategy, Business by Business	29
Risk and Strategy Management in a Complex Financial Institution	29
What Causes Financial Institutions to Fail?	31
The Role of Capital in Risk Management and Business Strategy	32
Capital-Based Risk Management in Banking Today: Pros and Cons	35
History of Capital-Based Regulations in Commercial Banking	37

PART TWO

Risk Management Techniques for Interest Rate Analytics

CHAPTER 3

Interest Rate Risk Introduction and Overview 45

Background Information on Movements in the U.S. Treasury Yield Curve	46
A Step-by-Step Approach to Analyzing Interest Rate Risk	55
The Interest Rate Risk Safety Zone	58

CHAPTER 4

Fixed Income Mathematics: The Basic Tools 59

Modern Implications of Present Value	59
Price, Accrued Interest, and Value	60
Calculation of Accrued Interest	60
Present Value	61
The Basic Present Value Calculation	61
Example	62
Calculating the Value of a Fixed Coupon Bond with Principal Paid at Maturity	62
Calculating the Coupon of a Fixed Coupon Bond with Principal Paid at Maturity When the Value Is Known	62
Example	63
The Value of an Amortizing Loan	63
Calculating the Payment Amount of an Amortizing Bond When the Value Is Known	63
Risk Management Implications	64
Calculating the Value of a Floating-Rate Bond or Loan with Principal Paid at Maturity	64
Example	65
Risk Management Implications	65
Compound Interest Conventions and Formulas	66
Future Value of an Invested Amount Earning at a Simple Interest Rate of y Compounded m Times per Year for n Periods	66
Future Value of an Invested Amount Earning at a Simple Interest Rate of y Compounded Continuously for n Years	66
Example	67
Present Value of a Future Amount If Funds Are Invested at a Simple Interest Rate of y Compounded m Times per Year for n Periods	67

Present Value of a Future Amount If Funds Are Invested at a Simple Interest Rate of y Compounded Continuously for n Years	67
Compounding Formulas and Present Value Factors $P(t)$	67
Yields and Yield-to-Maturity Calculations	68
The Formula for Yield to Maturity	68
Yield to Maturity for Long or Short First Coupon Payment Periods	69
Calculating Forward Interest Rates and Bond Prices	69
Implied Forward Interest Rates on Zero-Coupon Bonds	69
Example	70
Implied Forward Zero-Coupon Bond Prices	70
Present Value of Forward Fixed Coupon Bond	70
Implied Forward Price on a Fixed Coupon Bond	71
Implied Forward Coupon on a Fixed Coupon Bond	71
Other Forward Calculations	71
Summary	71

CHAPTER 5

Yield Curve Smoothing 73

Example A: Stepwise Constant Yields and Forwards vs. Nelson-Siegel	77
Deriving the Form of the Yield Curve Implied by Example A	79
Fitting the Nelson-Siegel Approach to Sample Data	81
Example D: Quadratic Yield Splines and Related Forward Rates	85
Deriving the Form of the Yield Curve Implied by Example D	86
Example F: Cubic Yield Splines and Related Forwards	94
Deriving the Form of the Yield Curve Implied by	
Example F Assumptions	95
Example H: Maximum Smoothness Forward	
Rates and Related Yields	101
Deriving the Parameters of the Quartic Forward Rate Curves	
Implied by Example H Assumptions	104
Comparing Yield Curve and Forward Rate Smoothing Techniques	111
Ranking 23 Smoothing Techniques by Smoothness of the	
Forward Rate Curve	112
Ranking 23 Smoothing Techniques by Length of the	
Forward Curve	112
Trading Off Smoothness vs. the Length of the Forward Rate Curve	112
The Shimko Test for Measuring Accuracy of Smoothing Techniques	116
Smoothing Yield Curves Using Coupon-Bearing Bond Prices as Inputs	116
Appendix: Proof of the Maximum Smoothness Forward Rate Theorem	117

CHAPTER 6

Introduction to Heath, Jarrow, and Morton Interest Rate Modeling 123

Objectives of the Example and Key Input Data	124
Key Implications and Notation of the HJM Approach	129
Pseudo-Probabilities	131
The Formula for Zero-Coupon Bond Price Shifts	132
Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 2$	132

Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 4$	134
Valuation in the HJM Framework	137
Valuation of a Zero-Coupon Bond Maturing at Time $T = 4$	139
Valuation of a Coupon-Bearing Bond Paying Annual Interest	140
Valuation of a Digital Option on the One-Year U.S. Treasury Rate	140
Conclusion	140

CHAPTER 7

HJM Interest Rate Modeling with Rate and Maturity-Dependent Volatility	142
Objectives of the Example and Key Input Data	142
Key Implications and Notation of the HJM Approach	146
Pseudo-Probabilities	147
The Formula for Zero-Coupon Bond Price Shifts	147
Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 2$	149
Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 4$	150
Valuation in the HJM Framework	153
Valuation of a Zero-Coupon Bond Maturing at Time $T = 4$	155
Valuation of a Coupon-Bearing Bond Paying Annual Interest	156
Valuation of a Digital Option on the One-Year U.S. Treasury Rate	158
Conclusion	158

CHAPTER 8

HJM Interest Rate Modeling with Two Risk Factors	161
Probability of Yield Curve Twists in the U.S. Treasury Market	161
Objectives of the Example and Key Input Data	162
Introducing a Second Risk Factor Driving Interest Rates	163
Key Implications and Notation of the HJM Approach	167
Pseudo-Probabilities	171
The Formula for Zero-Coupon Bond Price Shifts with	
Two Risk Factors	171
Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 2$	173
Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 3$	174
Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 4$	175
Valuation in the HJM Framework	178
Valuation of a Zero-Coupon Bond Maturing at Time $T = 4$	183
Valuation of a Coupon-Bearing Bond Paying Annual Interest	184
Valuation of a Digital Option on the One-Year	
U.S. Treasury Rate	185
Replication of HJM Example 3 in Common Spreadsheet Software	186
Conclusion	189

CHAPTER 9**HJM Interest Rate Modeling with Three Risk Factors 190**

Probability of Yield Curve Twists in the U.S. Treasury Market	190
Objectives of the Example and Key Input Data	191
Risk Factor 1: Annual Changes in the One-Year U.S.	
Treasury Spot Rate	192
Alternative Specifications of the Interest Rate Volatility Surface	200
Key Implications and Notation of the HJM Approach	201
Pseudo-Probabilities	205
The Formula for Zero-Coupon Bond Price Shifts with	
Three Risk Factors	205
Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 2$	206
Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 3$	208
Building the Bushy Tree for Zero-Coupon Bonds	
Maturing at Time $T = 4$	208
Valuation in the HJM Framework	217
Valuation of a Zero-Coupon Bond Maturing at Time $T = 4$	219
Valuation of a Coupon-Bearing Bond Paying Annual Interest	225
Valuation of a Digital Option on the One-Year	
U.S. Treasury Rate	227
Conclusion	229

CHAPTER 10**Valuation, Liquidity, and Net Income 230**

How Many Risk Factors Are Necessary to Accurately	
Model Movements in the Risk-Free Yield Curve?	230
Revisiting the Phrase “No Arbitrage”	231
Valuation, Liquidity Risk, and Net Income	234
Risk-Neutral and Empirical Probabilities of Interest	
Rate Movements	235
Monte Carlo Simulation Using HJM Modeling	236
Common Pitfalls in Interest Rate Risk Management	238
Pitfalls in the Use of One-Factor Term Structure Models	238
Common Pitfalls in Asset and Liability Management	243
Summarizing the Problems with Interpolated Monte Carlo	
Simulation for Risk Analysis	246

CHAPTER 11**Interest Rate Mismatching and Hedging 250**

Political Factions in Interest Rate Risk Management	251
Pension Fund Considerations	251
Life Insurance Companies and Property and Casualty Insurance	
Companies	252
Commercial Banks	253

Making a Decision on Interest Rate Risk and Return:	
The Safety Zone	254
Obvious Interest Rate Risk Decisions	255
Assessing the Risk and Return Trade-Offs from a	
Change in Interest Rate Risk	255

CHAPTER 12

Legacy Approaches to Interest Rate Risk Management 257

Gap Analysis and Simulation Models	257
Measuring Interest Rate Risk: A Review	258
Legacy Rate Risk Tools: Interest Rate Sensitivity Gap Analysis	258
The Safety Zone	259
What's Wrong with Gap Analysis?	263
Legacy Rate Risk Tools: Multiperiod Simulation	264
Key Assumptions in Simulation	264
Data Aggregation in Simulation Modeling	266
Constraining the Model	266
Modeling the Maturity Structure of a Class of Assets	267
Periodicity of the Analysis	267
Exceptions to the Exact Day Count Trend	267
Legacy Rate Risk Tools: Duration and Convexity	267
Macauley's Duration: The Original Formula	268
Using Duration for Hedging	270
Comparing a Duration Hedge with Hedging in the HJM Framework	271
Duration: The Traditional Market Convention	273
The Formula for Yield to Maturity	273
Yield to Maturity for Long or Short First Coupon Payment Periods	274
Applying the Yield-to-Maturity Formula to Duration	275
Modified Duration	276
The Perfect Duration Hedge: The Difference between the	
Original Macaulay and Conventional Durations	278
Convexity and Its Uses	278
Convexity: A General Definition	279
Convexity for the Present Value Formula	280
Hedging Implications of the Convexity Concept	280
Conclusion	281

CHAPTER 13

Special Cases of Heath, Jarrow, and Morton Interest Rate Modeling 283

What Is an Academic Term Structure Model and Why	
Was It Developed?	284
The Vocabulary of Term Structure Models	284
Ito's Lemma	286
Ito's Lemma for More Than One Random Variable	287
Using Ito's Lemma to Build a Term Structure Model	287
Duration as a Term Structure Model	288
Conclusions about the Use of Duration's Parallel Shift Assumptions	290
The Vasicek and Extended Vasicek Models	292

The Merton Term Structure Model: Parallel Yield Curve Shifts	293
The Extended Merton Model	298
The Vasicek Model	300
The Extended Vasicek—Hull and White Model	303
Alternative Term Structure Models	303
Alternative One-Factor Interest Rate Models	304
Two-Factor Interest Rate Models	306
Chen's Three-Factor Term Structure Model	307
Reprising the HJM Approach	308
Appendix A: Deriving Zero-Coupon Bond Prices in the Extended Merton/Ho and Lee Model	308
Appendix B: Deriving Zero-Coupon Bond Prices in the Vasicek Model	310
Appendix C: Valuing Zero-Coupon Bonds in the Extended Vasicek Model	313

CHAPTER 14

Estimating the Parameters of Interest Rate Models	316
Revisiting the Meaning of No Arbitrage	316
A Framework for Fitting Term Structure Models	316
Fitting Zero-Coupon Bond Prices and Volatility Parameters Jointly	317
Steps in Fitting the Interest Rate Volatility Assumptions	318
Example 1: Fitting Interest Rate Volatility When Six Callable Bonds Are Observable	318
Example 2: The Consequences of Fewer Inputs	329
Example 3: The Case of One Input	329
Interest Rate Parameter Fitting in Practical Application	330

PART THREE

Risk Management Techniques for Credit Risk Analytics

CHAPTER 15

An Introduction to Credit Risk: Using Market Signals in Loan Pricing and Performance Measurement	335
Market Prices for Credit Risk	335
Critical Sources of Market Data on Credit Risk	336
Bond Prices	336
Credit Default Swap Prices	337
First to Default Swaps	337
Collateralized Debt Obligations	338
Interest Rate Swap Prices	338
Equity Prices	338
Increased Accuracy in Pricing	339
Increased Clarity in Corporate Strategy	339
Increased Sophistication in Risk Management	340
Increased Precision in Measuring the Safety and Soundness of Financial Institutions	340
Credit Default Swaps: The Dangers of Market Manipulation	341

Daily Nondealer Trading Volume for 1,090 Reference Names	347
Credit Default Swap Trading Volume in Municipals and Sub-Sovereigns	352
Credit Default Swap Trading Volume in Sovereign Credits	353
Implications of CDS Trading Volume Data	357

CHAPTER 16

Reduced Form Credit Models and Credit Model Testing	359
The Jarrow-Turnbull Model	359
The Jarrow-Turnbull Framework	360
The Jarrow Model	361
Zero-Coupon Bond Prices in the Jarrow Model	363
The Jarrow Model and the Issue of Liquidity in the Bond Market	364
The Jarrow-Merton Put Option as a Risk Index and a Practical Hedge	364
Fitting the Jarrow Model to Bond Prices, Credit Derivatives Prices, and Historical Default Databases	365
Fitting the Jarrow Model to Debt Prices	365
Fitting to Current Price Data and Historical Price Data	366
Fitting the Jarrow Model to Credit Derivatives Prices	366
Fitting the Jarrow Model to a Historical Database of Defaults	366
Fitting the Jarrow Model to Retail, Small Business, and Governmental Counterparties	370
Correlations in Default Probabilities	372
The Jarrow and Jarrow-Turnbull Models: A Summary	373
Tests of Credit Models Using Historical Data	374
An Introduction to Credit Model Testing	375
Misunderstandings about Credit Model Testing	376
The Two Components of Credit Model Performance	378
Measuring Ordinal Ranking of Companies by Credit Risk	379
The Predictive ROC Accuracy Ratio: Techniques and Results	380
The Predictive Capability of the Jarrow-Chava Reduced Form Model Default Probabilities	380
Measuring the Predictive ROC Accuracy Ratio	381
Reduced Form Model vs. Merton Model Performance	381
Consistency of Estimated and Actual Defaults	383
Recent Results from North America	383
The Falkenstein and Boral Test	383
Performance of Credit Models vs. Naïve Models of Risk	386
ROC Accuracy Ratios for Merton Model Theoretical Version vs. Selected Naïve Models	387
Tests of Credit Models Using Market Data	388
Testing Credit Models: The Analogy with Interest Rates	388
Market Data Test 1: Accuracy in Fitting Observable Yield Curves and Credit Spreads	388
Market Data Test 2: Tests of Hedging Performance	389
Market Data Test 3: Consistency of Model Implications with Model Performance	390

Market Data Test 4: Comparing Performance with Credit Spreads and Credit Default Swap Prices	391
Appendix: Converting Default Intensities to Discrete Default Probabilities	391
Converting Monthly Default Probabilities to Annual Default Probabilities	392
Converting Annual Default Probabilities to Monthly Default Probabilities	392
Converting Continuous Instantaneous Probabilities of Default to an Annual Default Probability or Monthly Default Probability	392
Converting Continuous Default Probability to an Annual Default Probability	393
Converting Continuous Default Probability to a Monthly Default Probability	393
Converting an Annual Default Probability to a Continuous Default Intensity	393
Converting a Monthly Default Probability to a Continuous Default Intensity	394
CHAPTER 17	
Credit Spread Fitting and Modeling	396
Introduction to Credit Spread Smoothing	396
The Market Convention for Credit Spreads	397
A Better Convention for Credit Model—Independent Credit Spreads	398
Deriving the Full Credit Spread of a Risky Issuer	399
Credit Spread Smoothing Using Yield Curve—Smoothing Techniques	404
Setting the Scene: Smoothing Results for the Risk-Free Curve	404
A Naïve Approach: Smoothing ABC Yields by Ignoring the Risk-Free Curve	406
Fitting Credit Spreads with Cubic Splines	409
Maximum Smoothness Forward Credit Spreads	410
Comparing Results	411
Data Problems with Risky Issuers	413
The Case of LIBOR	413
Determinants of Credit Spread Levels	415
The Credit Risk Premium: The Supply and Demand for Credit	416
Conclusion	420
CHAPTER 18	
Legacy Approaches to Credit Risk	421
The Rise and Fall of Legacy Ratings	421
Ratings: What They Do and Don't Do	422
Through the Cycle vs. Point in Time, a Distinction without a Difference	423
Stress Testing, Legacy Ratings, and Transition Matrices	425
Transition Matrices: Analyzing the Random Changes in Ratings from One Level to Another	426

Moral Hazard in “Self-Assessment” of Ratings Accuracy by Legacy Rating Agencies	426
Comparing the Accuracy of Ratings and Reduced Form Default Probabilities	429
Problems with Legacy Ratings in the 2006 to 2011 Credit Crisis	431
The Jarrow-Merton Put Option and Legacy Ratings	437
The Merton Model of Risky Debt	438
The Intuition of the Merton Model	439
The Basic Merton Model	441
Valuing Multipayment Bonds with the Merton Model of Risky Debt	444
Estimating the Probability of Default in the Merton Model	445
Implying the Value of Company Assets and Their Return Volatility σ	446
Mapping the Theoretical Merton Default Probabilities to Actual Defaults	447
The Merton Model When Interest Rates Are Random	447
The Merton Model with Early Default	447
Loss Given Default in the Merton Model	448
Copulas and Correlation between the Events of Default of Two Companies	448
Back to the Merton Case	448
Problems with the Merton Model: Summing Up	449
Appendix	450
Assumptions	450
Using Ito’s Lemma to Expand Changes in the Value of Company Equity	450

CHAPTER 19

Valuing Credit Risky Bonds	453
The Present Value Formula	453
Valuing Bonds with No Credit Risk	454
Simulating the Future Values of Bonds with No Credit Risk	454
Current and Future Values of Fixed Income Instruments: HJM Background and a Straight Bond Example	455
Valuation of a Straight Bond with a Bullet Principal Payment at Maturity	461
Valuing an Amortizing Loan	461
Valuing Risk-Free, Floating-Rate Loans	465
Valuing Bonds with Credit Risk	465
Simulating the Future Values of Bonds with Credit Risk	471
Valuing the Jarrow-Merton Put Option	472

CHAPTER 20

Credit Derivatives and Collateralized Debt Obligations	473
Credit Default Swaps: Theory	474
Credit Default Swaps: Practice	477
Collateralized Debt Obligations: Theory	480
Collateralized Debt Obligations: A Worked Example of Reduced Form Simulation	483

Collateralized Debt Obligations: Practice	486
The Copula Method of CDO Valuation: A Postmortem	487
Valuing the Jarrow–Merton Put Option	490

PART FOUR

Risk Management Applications: Instrument by Instrument

CHAPTER 21

European Options on Bonds	495
Example: European Call Option on Coupon-Bearing Bond	501
Example: Coupon-Bearing Bond with Embedded European Call Option	503
European Options on Defaultable Bonds	509
HJM Special Case: European Options in the One-Factor Vasicek Model	509
Options on Coupon-Bearing Bonds	511
The Jarrow-Merton Put Option	512

CHAPTER 22

Forward and Futures Contracts	513
Forward Contracts on Zero-Coupon Bonds	514
Forward Rate Agreements	520
Eurodollar Futures-Type Forward Contracts	524
Futures on Zero-Coupon Bonds: The Sydney Futures Exchange Bank Bill Contract	527
Futures on Coupon-Bearing Bonds: Dealing with the Cheapest to Deliver Option	528
Eurodollar and Euroyen Futures Contracts	529
Defaultable Forward and Futures Contracts	530

CHAPTER 23

European Options on Forward and Futures Contracts	531
Valuing Options on Forwards and Futures: Notations and Useful Formulas	531
European Options on Forward Contracts on Zero-Coupon Bonds	532
European Options on Forward Rate Agreements	538
European Options on a Eurodollar Futures-Type Forward Contract	540
European Options on Futures on Coupon-Bearing Bonds	546
European Options on Money Market Futures Contracts	546
Defaultable Options on Forward and Futures Contracts	546

CHAPTER 24

Caps and Floors	548
Caps as European Options on Forward Rate Agreements	550
Forming Other Cap-Related Securities	550
Valuing a Cap	550
Valuing a Floor	554
Valuing a Floating Rate Loan with a Cap	557

Value of a Loan with a Cap and a Floor	563
Variations on Caps and Floors	565
Measuring the Credit Risk of Counterparties on Caps and Floors	565

CHAPTER 25

Interest Rate Swaps and Swaptions	567
Interest Rate Swap Basics	567
Valuing the Interest Rate Swaps	568
The Observable Fixed Rate in the Swap Market	574
An Introduction to Swaptions	574
Valuation of European Swaptions	578
Valuation of American Swaptions	579
Defaultable Interest Rate Swaps and Swaptions	579

CHAPTER 26

Exotic Swap and Options Structures	580
Arrears Swaps	580
Digital Option	586
Digital Range Notes	588
Range Floater	588
Other Derivative Securities	593
Credit Risk and Exotic Derivatives Structures	594

CHAPTER 27

American Fixed Income Options	596
An Overview of Numerical Techniques for Fixed Income Option Valuation	597
An Example of Valuation of a Callable Bond with a Three-Factor HJM Bushy Tree	598
What Is the Par Coupon on a Callable Bond?	613
An Example of Valuation of a Rationally Prepaid Amortizing Loan	613
Monte Carlo Simulation	615
Conclusions	618
Finite Difference Methods	618
Binomial Lattices	619
Trinomial Lattices	619
HJM Valuation of American Fixed Income Options When Default Risk Is Present	620

CHAPTER 28

Irrational Exercise of Fixed Income Options	622
Analysis of Irrationality: Criteria for a Powerful Explanation	623
The Transactions Cost Approach	624
Irrational Exercise of European Options	625
The Irrational Exercise of American Options	626

A Worked Example Using an Amortizing Loan with Rational and Irrational Prepayment Behavior	626
Implied Irrationality and Hedging	636
Credit Risk and Irrational Prepayment Behavior	637

CHAPTER 29

Mortgage-Backed Securities and Asset-Backed Securities 639

Transactions Costs, Prepayments, Default, and Multinomial Logit	640
Legacy Prepayment Analysis of Mortgage-Backed Securities	643
Legacy Approaches: Prepayment Speeds and the Valuation of Mortgages	643
Constant Prepayment Speeds Are Simply a Principal Amortization Assumption	644
Legacy Approaches: Option-Adjusted Spread	645
Implications for OAV Spread, CMOs, and ARMs	647
Logistic Regression, Credit Risk, and Prepayment	648
Mortgage-Servicing Rights: The Ultimate Structured Product	648
An Introduction to the Valuation of Mortgage-Servicing Rights	649
Comparing Best Practice and Common Practice in Valuing and Hedging Mortgage-Servicing Rights	650
Valuation Yield Curve for Cash Flows	650
Simulation of Random Movements in Yields	651
The Role of Home Prices in Defaults and Prepayments	652
Other Sources of Cash Flow Related to Mortgage-Servicing Rights	653
Incorrect Hedging of Mortgage-Servicing Rights	653
Conclusion	654

CHAPTER 30

Nonmaturity Deposits 656

The Value of the Deposit Franchise	657
Total Cash Flow of Nonmaturity Deposits	658
Specifying the Rate and Balance Movement Formulas	659
The Impact of Bank Credit Risk on Deposit Rates and Balances	669
Case Study: German Three-Month Notice Savings Deposits	672
The Regulators' View	673
Conclusion	674

CHAPTER 31

Foreign Exchange Markets 675

Setting the Stage: Assumptions for the Domestic and Foreign Economies	675
Foreign Exchange Forwards	676
Numerical Methods for Valuation of Foreign Currency Derivatives	677
Legacy Approaches to Foreign Exchange Options Valuation	678
Implications of a Term Structure Model-Based FX Options Formula	680
The Impact of Credit Risk on Foreign Exchange Risk Formulas	681

CHAPTER 32**Impact of Collateral on Valuation Models: The Example of Home Prices in the Credit Crisis 682**

The Impact of Changing Home Prices on Collateral Values in the Credit Crisis	682
Modeling Variations in Collateral Values	683
The Impact of Collateral Values on a Rationally Prepaid Mortgage	684
Conclusions about the Impact of Collateral Values	693

CHAPTER 33**Pricing and Valuing Revolving Credit and Other Facilities 694**

Analyzing Revolving Credit and Other Facilities	695
Fluctuating Credit Risk and Revolving Credit Drawdowns	696
Incorporating Links between Credit Quality and Line Usage	697
Is a Line of Credit a Put Option on the Debt of the Issuer?	697

CHAPTER 34**Modeling Common Stock and Convertible Bonds on a Default-Adjusted Basis 700**

Modeling Equities: The Traditional Fund Management Approach	701
Modeling Equities: The Derivatives Approach	702
Modeling Equities: A Credit Risk–Adjusted Approach	703
Options on the Common Stock of a Company That Can Go Bankrupt	704
Convertible Bonds of a Company That Can Go Bankrupt	706

CHAPTER 35**Valuing Insurance Policies and Pension Obligations 708**

Life Insurance: Mortality Rates vs. Default Probabilities	708
Cyclicalities in Default Probabilities and Mortality Rates	711
Valuing Life Insurance Policies	711
Pension Obligations	712
Property and Casualty Insurance	713
The Jarrow-Merton Put Option	714

PART FIVE**Portfolio Strategy and Risk Management****CHAPTER 36****Value-at-Risk and Risk Management Objectives Revisited at the Portfolio and Company Level 719**

The Jarrow-Merton Put Option as a Measure of Total Risk: An Example	719
A Four-Question Pass–Fail Test for Financial Institutions’ CEOs and Boards of Directors	723
Why Do These Four Questions Matter?	724
An Alphabet of 26 Extra-Credit Questions	724
Is Your Value-at-Risk from Value-at-Risk?	726

VaR vs. the Put Option for Capital Allocation	728
Why Are the VaR and Put Approaches So Different:	
Self-Insurance vs. Third-Party Insurance	729
Calculating the Jarrow-Merton Put Option Value and	
Answering the Key 4 + 26 Questions	731
Valuing and Simulating the Jarrow-Merton Put Option	732
What's the Hedge?	733
Liquidity, Performance, Capital Allocation, and	
Own Default Risk	734

CHAPTER 37

Liquidity Analysis and Management: Examples from the Credit Crisis	735
Liquidity Risk Case Studies from the Credit Crisis	735
Case Studies in Liquidity Risk	736
Largest Funding Shortfalls	736
American International Group (AIG)	737
Consolidated JPMorgan Chase, Bear Stearns, and	
Washington Mutual	744
State Street	746
Morgan Stanley	749
Dexia Credit Local New York Branch	751
Implications of the Credit Crisis History for Liquidity	
Risk Management and Analysis	758
Types of Liquidity Events	758
Liquidity Risk and Credit Risk Linkages	759
Measuring Liquidity Risk as a Line of Credit in the	
Jarrow-Merton Put Option Sense	760
Integrating Managerial Behavior and Market Funds Supply	
in Liquidity Risk Measurement	761
Determining the Optimal Liquidity Strategy	763
Summing Up	763

CHAPTER 38

Performance Measurement: Plus Alpha vs. Transfer Pricing	765
Transaction-Level Performance Measurement vs. Portfolio-	
Level Performance Measurement	766
Plus Alpha Benchmark Performance vs. Transfer Pricing	767
Why Default Risk Is Critical in Performance Measurement	
of Equity Portfolios	768
"Plus Alpha" Performance Measurement in Insurance and Banking	769
Decomposing the Reasons for Plus or Minus Alpha in a Fixed	
Income Portfolio	770
A Worked Example of Modern Fixed Income Performance Attribution	772
The Jarrow-Merton Put Option and Capital	780
Using the Jarrow-Merton Put Option for Capital Allocation	780
Introduction	780
Using the Jarrow-Merton Put Option Concept for Capital Allocation	780

Extending the Jarrow-Merton Capital Allocation to a Multiperiod Framework	782
Summing Up	782
CHAPTER 39	
Managing Institutional Default Risk and Safety and Soundness	783
Step 1: Admitting the Possibility of Failure	783
Managing the Probability of Failure	785
Are Ratings a Useful Guide?	785
Are CDS Spreads a Useful Guide?	786
Using Quantitative Default Probabilities	787
Controlling the Probability of Failure through the Credit Cycle	789
Hedging Total Risk to Maximize Shareholder Value	790
Implications for Basel II, Basel III, and Solvency II	791
Simulating Your Own Probability of Default	792
CHAPTER 40	
Information Technology Considerations	793
Common Practice in Risk Management Systems: Dealing with Legacy Systems	793
Upgrading the Risk Infrastructure: The Request for Proposal Process	795
Paid Pilots as Final Proof of Concept	796
Keys to Success in Software Installation	797
Vendor Size: Larger Vendor or Small Vendor?	798
Being a Best Practice User	799
CHAPTER 41	
Shareholder Value Creation and Destruction	800
Do No Harm	800
Measure the Need to Change	801
Rating Your Primary Risk System	803
Master the Politics and Exposition of Risk Management:	
Shareholder Value Creation	803
Daily Management Reporting of Total Risk	805
Moving from Common Practice to Best Practice	806
The Senior Management Perspective	807
The Middle Management Perspective	807
The Working-Level Perspective	807
Getting Help to Create Shareholder Value	808
Postscript	808
Bibliography	809
Index	819

Introduction: Wall Street Lessons from Bubbles

The credit crisis that began to unfold in the United States and Europe in 2006 contains a treasure trove of lessons for risk managers that we have tried to incorporate into this book. Since we have each worked in Japan, we felt strongly that the collapse of the Japanese bubble, which peaked in late 1989, contained equally useful lessons for risk managers. As you'll note in the "key fallacies in risk management" discussed below, many ignored the lessons of the Japanese bubble because of the common fallacy that "if it hasn't happened to me yet, it won't happen to me, even if it's happened to someone else."

Now that the United States and much of Europe are experiencing the collapse of a bubble much like that which burst in Japan, the lessons from each of these bubbles seem much more relevant to risk managers around the world.

We have worked hard in the second edition of this book to severely de-emphasize the discussion of financial models that are obviously inaccurate, misleading, or clearly inferior to another modeling approach. We make this judgment on the basis of cold hard facts (via model testing) or because of assumptions that are known to be false. The list of models that failed under the duress of the credit crisis is a long one, and we make no apologies for reflecting those failures in this book. We've also worked hard to explain which models performed well during the credit crisis. Again, we base that judgment on model testing and the logical consistency and accuracy of the assumptions behind those models.

We believe in a "multiple models approach" to risk management. That doesn't mean, however, that all models used are equally accurate. Nothing could be further from the truth. The use of a multiple models approach, however, makes it clear when a better model has been brought to the risk management discipline and when it's time for an older model to be removed from the toolkit. One of our British friends provided the elegant observation that "I don't think it's gentlemanly to compare the accuracy of two models." The authors, by contrast, believe that such comparisons are a mandatory part of corporate governance and best practice risk management.

With that brief introduction, we turn to a short summary of common fallacies in risk management that have been exposed as fallacies very starkly in the wake of the recent credit crisis.

KEY FALLACIES IN RISK MANAGEMENT

Summarizing the dangerous elements of conventional wisdom in risk management isn't easy. We've restricted ourselves to the seven most dangerous ways of thinking about risk. Each of them has much in common with this famous quote of John Maynard Keynes from

“The Consequences to the Banks of the Collapse of Money Values” in 1931: “A sound banker, alas, is not one who foresees danger and avoids it, but one who, when he is ruined, is ruined in a conventional way along with his fellows, so that no one can really blame him.” We summarize them here and discuss each briefly in turn:

- If it hasn’t happened to me yet, it won’t happen to me, even if it’s happened to someone else.
- Silo risk management allows my firm to choose the “best of breed” risk model for our silo.
- I don’t care what’s wrong with the model. Everyone else is using it.
- I don’t care what’s wrong with the assumptions. Everyone else is using them.
- Mathematical models are superior to computer simulations.
- Big North American and European banks are more sophisticated than other banks around the world and we want to manage risk like they do.
- Goldman says they do it this way and that must be right.

We discuss each of these fallacies in turn.

If It Hasn’t Happened to Me Yet, It Won’t Happen to Me, Even If It’s Happened to Someone Else.

The perceived risk in any given situation is often a function of the age, the location, and the personal experience of the person making the risk assessment. There are lots of examples. Some motorcycle riders don’t wear helmets. Some smokers think cancer happens to other people. In risk management, perhaps the biggest mistake in 2005 and 2006 was the assumption that “home prices don’t go down” in the United States, in spite of considerable evidence in a major feature on the home price bubble from *The Economist* (June 16, 2005). Exhibit I.1 compares, from the perspective of January 2006, the relative home price appreciation during the Japanese bubble (since 1982) and Case-Shiller home prices indices for Los Angeles and a 10-city composite (since 1996). U.S. market participants, in spite of evidence like this, continued to assume U.S. home prices would rise.

Of course, this assumption was dramatically wrong. It’s best illustrated by Richard C. Koo (“The World in Balance Sheet Recession: Causes, Cure, and Politics,” Nomura Research Institute, 2011) in Exhibit I.2, showing how similar the collapse in U.S. and Japan home prices has been.

Obviously, just because something hasn’t happened to you yet doesn’t mean that it won’t. Why that isn’t obvious to more risk managers is a mystery to us.

Silo Risk Management Allows My Firm to Choose the “Best of Breed” Risk Model for Our Silo.

Alan Greenspan, former Chairman of the Board of Governors of the Federal Reserve, made an important confession in this quote from the *Guardian*, five weeks after the collapse of Lehman Brothers, on October 23, 2008: “I made a mistake in presuming that the self-interests of organizations, specifically banks and others, were such that they were best capable of protecting their own shareholders and their equity in the firms.” In the same manner, management of large financial institutions often ignores




EXHIBIT I.1 Feeling Homesick?




EXHIBIT I.2 U.S. and Japan—Collapse in Home Prices

Sources: Bloomberg, Real Estate Economic Institute, Japan, S&P, S&P/Case Shiller[®] Home Price Indices, as of October 5, 2011.

the fact that their staff acts in their own best interests, with the best interests of the organization being a secondary consideration at best. How else to explain why institutions often use separate and inconsistent risk systems for market risk, credit risk, liquidity risk, interest rate risk, capital allocation, and so on? Using today's financial and computer technology, there is no systems or financial theory reason for such systems to be separate and distinct.

Disparate risk management units almost never cooperate in consolidating risk systems, even when there is an obvious increase in accuracy in risk assessment, unless it is forced from the top down. With the recent decline in U.S. home prices driving all of these silo risks, using a suite of separate risk systems (most of which could not analyze home prices as a risk factor) does nothing for the shareholders, but it builds barriers that preserve the independence of silo risk management organizations.

I Don't Care What's Wrong with the Model. Everyone Else Is Using It.

One of us gave a presentation to the European meeting of the International Association of Credit Portfolio Managers in Zurich in May 2007, just as the credit crisis was beginning to become obvious to more people. The thrust of the presentation was that the commonly used copula method for valuation of collateralized debt obligations gave dramatically more optimistic valuations than an alternative approach (the "reduced form approach"), in which a series of key macroeconomic (macro) factors are common drivers of default probabilities for each reference name. After the presentation, two bankers illustrated the accuracy of Keynes's previous quote by saying, "That was a good presentation, but there were 100 people in this room and 99 of them think you're wrong." Alas, the presentation was right on the mark, but the power of conventional wisdom blinded many to their model risk.

I Don't Care What's Wrong with the Assumptions. Everyone Else Is Using Them.

In this book, we devote a lot of time to pointing out the false assumptions of various models as an important part of assessing model risk. Surprisingly, many risk managers are unconcerned about false assumptions, especially if many others are using the same false assumptions. Two of the most spectacular falsehoods have to be mentioned. The first is the use of the Black model to value interest rate caps and floors, even though the Black model assumes interest rates are constant! The second is the assumption of the most common form of copula model for CDO valuation, that every pair of reference names has the identical correlation in default risk. The lack of anxiety about the false assumptions in these two cases in particular has resulted in a lot of job losses at major financial institutions. Beware.

Mathematical Models Are Superior to Computer Simulations.

Consistent with Alan Greenspan's earlier quote, it would not be accurate to assume that the academic builders of financial models have the shareholders' interests at heart. They don't. They have their own interests at heart, and primary among those interests is to be published in prestigious academic journals as often as possible. Even the most casual reader of such journals will observe that the overwhelming majority of articles about valuation models presents a formula for an answer, not a

simulation. In order to get a mathematical formula for an answer, even the best academic financial theorists are forced to make simplifying assumptions that are not accurate. Among the many such assumptions are assumptions that a variable is normally distributed when it's not, assumptions that returns from period to period are independent when they're not, and the assumption that the number of risk factors that determine risk are small in number (1, 2, or 3) when in fact that's not the case.

A common theme in this book is that such simplifications have caused many institutions to fail and many risk managers to lose their jobs. A realistic simulation that has no "closed form" mathematical solution is usually the only accurate way to describe risk accurately. Sadly, such a simulation will normally not be published in a prestigious academic finance journal.

Big North American and European Banks Are More Sophisticated Than Other Banks around the World and We Want to Manage Risk Like They Do.

In spite of the second \$1 trillion bail-out of U.S. financial institutions in the last 25 years, many persist in the belief that large North American and European financial institutions are the most skilled risk managers in the world. This hallucination persists in spite of a mass of public evidence to the contrary. Two quotes with respect to Citigroup in the aftermath of the credit crisis illustrate the point quite well. The *New York Times* (November 22, 2008) suggests the risk management expertise of top management at Citigroup:

"Chuck Prince going down to the corporate investment bank in late 2002 was the start of that process," a former Citigroup executive said of the bank's big C.D.O. push. "Chuck was totally new to the job. He didn't know a C.D.O. from a grocery list, so he looked for someone for advice and support. That person was Rubin. And Rubin had always been an advocate of being more aggressive in the capital markets arena. He would say, 'You have to take more risk if you want to earn more.'"

In *Fortune* magazine (November 28, 2007), Robert Rubin makes the point even more strongly: "I tried to help people as they thought their way through this. Myself, at that point, I had no familiarity at all with CDOs."

Besides this evidence, there are the public records of the bailouts of Citigroup, Bank of America, Merrill Lynch, Bear Stearns, Royal Bank of Scotland, HBOS, and a host of European institutions. Why does this belief—a hallucination, perhaps—in the risk management expertise of North American and European institutions go on? We think that management teams in many countries are much more sophisticated about risk management on a day-to-day basis than U.S. banks were during the credit crisis. Working with risk managers in 32 countries, we still see this differential in expertise on a daily basis.

Goldman Says They Do It This Way and That Must Be Right.

We continue to be astonished that many financial market participants believe the modeling approach described to them by the big Wall Street firm's local sales rep. The

use of financial model discussions in the mid-1980s by Salomon Brothers in Tokyo was a legendary example of how to do what Wall Street does, fleecing clients by building up a “relationship” about financial models in which naïve buy-side clients reveal their proprietary modeling approach in trade for hearing about Salomon’s “proprietary” yen fixed-income models. In reality, the modeling conversation was strictly to exploit the clients’ naiveté by model arbitrage. Total profits in Tokyo for Salomon during this time averaged more than \$500 million per year for more than a decade. Model arbitrage was at the heart of the CDO losses on Wall Street, with an interesting twist. Many traders arbitrated their own firms and then moved on when the game ended. It is important to remember that Wall Street and Wall Street traders have the same “relationship” with clients that you might have with a hamburger at McDonald’s.

SELECTED EVENTS IN THE CREDIT CRISIS

As the credit crisis recedes into history, leaving only U.S. and European government deficits behind, it is important to record the credit crisis history before it’s lost. This section summarizes recent events to confirm the accuracy of the assertions above. In the chapters that follow, we describe an approach to risk management designed to protect our readers from the kinds of events that unfolded in the credit crisis.

The credit crisis chronology below was assembled from many sources:

- A contemporaneous set of press clippings and news articles maintained by Kamakura Corporation from the very early dates of the credit crisis
- A credit risk chronology maintained by the Federal Reserve Bank of St. Louis
- A credit risk chronology maintained by the University of Iowa. We would like to thank @bionicturtle via twitter for bringing this chronology to our attention
- The “Levin report,” released by Senator Carl Levin on April 13, 2011

The full citation for the Levin report is as follows:

Wall Street and the Financial Crisis: Anatomy of a Financial Collapse, Majority and Minority Staff Report, Permanent Subcommittee on Investigations, United States Senate, April 13, 2011.

This list is an abridged version of a longer chronology on www.kamakuraco.com listed in the bibliography.

September 2, 2004	Washington Mutual chief risk officer Jim Vanasek sends internal Washington Mutual memo stating “At this point in the mortgage cycle with prices having increased far beyond the rate of increase in personal incomes, there clearly comes a time when prices must slow down or perhaps even decline.” (Levin report, p. 66)
May 31, 2005	“Fed Debates Pricking the U.S. Housing Bubble,” <i>New York Times</i> . (Levin report, p. 272)
June 3, 2005	“Yale Professor Predicts Housing Bubble Will Burst,” <i>National Public Radio</i> . (Levin report, p. 272)
June 16, 2005	“The global housing boom: In come the waves. The worldwide rise in house prices is the biggest bubble in

	history. Prepare for the economic pain when it pops.” <i>The Economist</i> .
September 30, 2005	Case-Shiller home price index for Boston, MA peaks. (Standard & Poor’s)
April 30, 2006	In an April memo discussing Countrywide’s issuance of subprime 80/20 loans, which are loans that have no down payment and are comprised of a first loan for 80 percent of the home’s value and a second loan for the remaining 20 percent of value, resulting in a loan to value ratio of 100 percent, Countrywide CEO Angelo Mozilo wrote “In all my years in the business I have never seen a more toxic pr[o]duct.” (Levin report, p. 232)
May 2, 2006	Ameriquist Mortgage closes retail branch network and lays off 3,600 employees. (<i>Orange County Register</i>)
May 5, 2006	Merit Financial Inc. files for bankruptcy. (SeattlePI)
July 31, 2006	Case-Shiller home price index for 20-city Composite Index peaks. (Standard & Poor’s)
August 31, 2006	Case-Shiller home price index for Las Vegas peaks. (Standard & Poor’s)
September 30, 2006	Case-Shiller home price index for Los Angeles peaks. (Standard & Poor’s)
October 31, 2006	S&P director in an internal e-mail “note also the ‘mailing in the keys and walking away’ epidemic has begun.” (Levin report, p. 268)
December 13, 2006	Dan Sparks informed the Firm-Wide Risk Committee of Goldman Sachs that two more subprime originators had failed in the last week and that there was concern about early payment defaults, saying, “Street aggressively putting back early payment defaults to originators thereby affecting the originators business. Rumors around more failures in the market.” (Levin report, p. 478)
January 31, 2007	By January 2007, nearly 10 percent of all subprime loans were delinquent, a 68 percent increase from January 2006. (Levin report, p. 268)
February 2, 2007	Dan Sparks reported to senior Goldman Sachs executives, “The team is working on putting loans in the deals back to the originators (New Century, Wamu, and Fremont, all real counterparties) as there seem to be issues potentially including some fraud at origination . . .” (Levin report, p. 484)
February 22, 2007	HSBC fires head of its U.S. mortgage lending business as losses reach \$10.5 billion
April 2, 2007	New Century, subprime lender, declares bankruptcy. (Bloomberg; Levin report, p. 47)
June 17, 2007	Two Bear Stearns subprime hedge funds collapse. (Levin report, p. 47)

July 6, 2007	UBS fires CEO and the heir apparent for Chairman of the Board Peter Wuffi. (<i>Financial Times</i>)
July 9, 2007	Credit Suisse releases a report that shows overall market CDO losses could total up to \$52 billion. (Bloomberg)
August 2, 2007	Bailout of IKB Deutsche Industriebank AG due to losses of up to €1 billion on mortgage-related CDOs. (Bloomberg)
August 3, 2007	AXA rescues money market funds after subprime losses. (Reuters)
August 9, 2007	BNP freezes \$2.2 billion of funds over subprime. (Reuters, Bloomberg)
August 14, 2007	17 Canadian structured investment vehicles fail when commercial paper is denied by Canadian banks. (Bloomberg)
August 16, 2007	Countrywide taps \$11.5 billion commercial paper backup line. (Bloomberg)
September 14, 2007	Bank of England rescues Northern Rock over UK mortgage losses. (Reuters)
October 1, 2007	UBS announces a \$3.7 billion write-down and, after the announcement, the chief executive of its investment banking division, Huw Jenkins, was replaced. (<i>Financial Times</i>)
October 5, 2007	Merrill Lynch writes down \$5.5 billion in losses on subprime investments. (Reuters)
October 16, 2007	Citigroup announces \$3 billion in write-offs on subprime mortgages. (<i>Financial Times</i>)
October 18, 2007	Bank of America writes off \$4 billion in losses (Bloomberg)
October 24, 2007	Merrill Lynch writes down \$7.9 billion on subprime mortgages and related securities. (Bloomberg; <i>Financial Times</i>)
October 30, 2007	Merrill Lynch CEO O'Neal fired. (Reuters)
November 5, 2007	Citigroup CEO Prince resigns after announcement that Citigroup may have to write down as much as \$11 billion in bad debt from subprime loans. (Bloomberg)
November 7, 2007	Morgan Stanley reports \$3.7 billion in subprime losses. (Bloomberg)
November 13, 2007	Bank of America says it will have to write off \$3 billion in bad debt. (BBC News)
November 15, 2007	Barclays confirms a \$1.6 billion write-down in the month of October on their subprime holdings. The bank also released that more than £5 billion in exposure to subprime loan packages could lead to more write-downs in the future. (Bloomberg)
November 21, 2007	Freddie Mac announces \$2 billion in losses from mortgage defaults (<i>Financial Times</i>)
December 6, 2007	The Royal Bank of Scotland announces that it expects to write down £1.25 billion because of exposure to the U.S. subprime market. (BBC News)

December 10, 2007	UBS announces another \$10 billion in subprime write-downs, bringing the total to date to \$13.7 billion. UBS also announced a capital injection of \$11.5 billion from the government of Singapore and an unnamed Middle East investor. (MarketWatch.com)
December 12, 2007	Federal Reserve establishes Term Auction Facility to provide bank funding secured by collateral. (Levin report, p. 47)
December 19, 2007	Morgan Stanley announces \$9.4 billion in write-downs from subprime losses and a capital injection of \$5 billion from a Chinese sovereign wealth fund. (<i>Financial Times</i>)
January 15, 2008	Citigroup reports a \$9.83 loss in the fourth quarter after taking \$18.1 billion in write-downs on subprime mortgage-related exposure. The firm also announced it would raise \$12.5 billion in new capital, including \$6.88 billion from the Government of Singapore Investment Corporation. (<i>Financial Times</i>)
January 17, 2008	Merrill Lynch announces net loss of \$7.8 billion for 2007 due to \$14.1 billion in write-downs on investments related to subprime mortgages. (BBC News)
January 22, 2008	Ambac reports a record loss of \$3.26 billion after write-downs of \$5.21 billion on its guarantees of subprime mortgage-related bonds. (<i>Financial Times</i>)
February 14, 2008	UBS confirmed a 2007 loss of \$4 billion on \$18.4 billion in write-downs related to the subprime crisis.
February 17, 2008	Britain announces the nationalization of Northern Rock, with loans to Northern Rock reaching 25 billion pounds sterling. (<i>Financial Times</i>)
February 28, 2008	AIG announces a \$5.2 billion loss for the fourth quarter of 2007, its second consecutive quarterly loss. AIG announced write-downs of \$11.12 billion pretax on its credit default swap portfolio. (<i>Financial Times</i>)
March 3, 2008	HSBC, the largest UK bank, reports a loss of \$17.2 billion in write-downs of its U.S. mortgage portfolio. (BBC News)
March 14, 2008	Federal Reserve and JPMorgan Chase agree to provide emergency funding for Bear Stearns. Under the agreement, JPMorgan would borrow from the Federal Reserve discount window and funnel the borrowings to Bear Stearns. (<i>Forbes</i> ; DataCenterKnowledge.com)
March 16, 2008	JPMorgan Chase agrees to pay \$2 per share for Bear Stearns on Sunday, March 16, a 93 percent discount to the closing price on Friday March 14. JPMorgan agreed to guarantee the trading liabilities of Bear Stearns, effective immediately. (<i>New York Times</i>)
April 1, 2008	UBS, whose share price fell 83 percent in the last year, reports it will write down \$19 billion in the first quarter on its U.S. holdings. (<i>Financial Times</i>)

April 1, 2008	UBS CEO Ospel resigns after announcement that UBS total losses are almost \$38 billion. (Bloomberg)
April 18, 2008	Citigroup reports \$5.11 billion in first quarter losses and \$12 billion in write-downs on subprime mortgage loans and other risky assets. The bank plans to cut 9,000 jobs in addition to the 4,200 jobs cut in January. (BBC News)
May 9, 2008	AIG reports \$7.81 billion in first quarter losses and \$9.11 billion of write-downs on the revaluation of their credit default swap portfolio. AIG Holding Company was also downgraded to AA– by two major rating agencies. (SeekingAlpha.com).
June 2, 2008	Wachovia CEO Thompson is ousted following large losses that resulted from the acquisition of a big mortgage lender at the peak of the housing market. (Reuters)
June 25, 2008	Shareholders of Countrywide, a troubled mortgage lender, approve the acquisition of the company by Bank of America. (<i>Financial Times</i>)
July 11, 2008	IndyMac Bank, a \$30 billion subprime mortgage lender fails and is seized by FDIC after depositors withdraw \$1.55 billion. (Levin report, pp. 47 and 234)
July 17, 2008	Merrill Lynch writes down \$9.4 billion on mortgage related assets. (<i>Financial Times</i>)
July 19, 2008	Citigroup lost \$5.2 billion and had \$7.2 billion of write downs in the second quarter. (<i>Financial Times</i>)
July 22, 2008	Marking the largest loss in the history of the fourth largest U.S. bank, Wachovia loses \$8.9 billion in the second quarter. (<i>Financial Times</i>)
September 7, 2008	U.S. takes control of Fannie Mae & Freddie Mac. The two companies currently have \$5,400 billion in outstanding liabilities and guarantee three-quarters of new U.S. mortgages. The government agreed to inject up to \$100 billion in each of them and will buy mortgage-backed bonds. (Levin report, p. 47; <i>Financial Times</i>)
September 11, 2008	Lehman Brothers reports its worst ever third quarter as it lost \$3.9 billion total and \$7.8 billion in credit linked write-downs. The company plans to shrink its size. (<i>Financial Times</i>)
September 13, 2008	The U.S. Treasury and Federal Reserve refuse to provide public funds to help with a rescue takeover for Lehman Brothers as they did for Bear Sterns. Bank of America backs out of negotiations with Lehman Brothers because of the lack of government funds. (<i>Financial Times</i>)
September 15, 2008	Lehman Brothers bankruptcy. (Levin report, p. 47)
September 15, 2008	Merrill Lynch announces sale to Bank of America. (Levin report, p. 47)

September 16, 2008	Federal Reserve offers \$85 billion credit line to AIG; Reserve Primary Money Fund NAV falls below \$1. (Levin report, p. 47)
September 18, 2008	Lloyds rescues HBOS, the largest mortgage lender in the UK. (Reuters)
September 21, 2008	Goldman Sachs and Morgan Stanley convert to bank holding companies. (Levin report, p. 47)
September 25, 2008	Washington Mutual fails, subsidiary bank is seized by the FDIC and sold to JPMorgan Chase for \$1.9 billion. JPMorgan Chase immediately wrote off \$31 billion in losses on the Washington Mutual assets. (<i>The Guardian</i> ; Levin report, p. 47)
October 3, 2008	Congress and President Bush establish Troubled Asset Relief Program (TARP), which is created by the Emergency Economic Stabilization Act (EESA). The revised bailout plan allows the Treasury to restore stability to the financial system by buying \$700 billion in toxic debt from banks. (Levin report, p. 47; SIGTARP report, p. 2; CNN)
October 4, 2008	Wells Fargo offers \$15.1 billion to buy Wachovia, outbidding Citigroup's \$2.2 billion bid. (<i>Financial Times</i>)
October 6, 2008	Germany announces €50 billion bail-out of Hypo Real Estate AG. (<i>USAToday</i>)
October 6, 2008	Germany announces unlimited guarantee of €568 billion in private bank deposits. (<i>USAToday</i>)
October 8, 2008	The UK government implements £400 billion rescue plan that includes government investing in banking industry, guaranteeing up to £250 billion of bank debt, and adding £100 billion to the Bank of England's short-term loan scheme. (<i>Financial Times</i>)
October 13, 2008	The UK government injects £37 billion in the nation's three largest banks, kicking off the nationalization process. (<i>Financial Times</i>).
October 14, 2008	The U.S. government announces capital injection of \$250 billion, of which \$125 billion will go to nine large banks as part of the Capital Purchase Program (CPP) in exchange for more government control on items such as executive compensation. (<i>Financial Times</i> ; SIGTARP report, p. 1)
October 16, 2008	Switzerland government gives UBS a \$59.2 billion bailout. (Bloomberg)
October 28, 2008	U.S. uses TARP to buy \$125 billion in preferred stock at nine banks. The nine banks held over \$11 trillion in banking assets or roughly 75 percent of all assets owned by U.S. banks. (Levin report, p. 47; SIGTARP report, p. 1)

November 23, 2008	U.S. gives government bailout to Citigroup, agreeing to cover losses on roughly \$306 billion of Citigroup's risky assets. (Reuters)
December 18, 2008	President George W. Bush reveals plan to rescue General Motors and Chrysler by lending them a total of \$17.4 billion. (<i>Financial Times</i>)
January 15, 2009	The U.S. government gives Bank of America an additional \$20 billion as part of TARP's Targeted Investment Program (TIP), which allows the Treasury to make additional targeted investments than what was given under TARP's Capital Purchase Program. Furthermore, the government agrees to guarantee nearly \$118 billion of potential losses on troubled assets. (SIGTARP report, p. 1; <i>Financial Times</i>)
March 20, 2009	German parliament passes Hypo Real Estate Nationalization bill. (Deutsche Welle)
September 12, 2010	German government adds another €50 billion in aid to Hypo Real Estate AG bringing total support to €142 billion. (Business Standard)

PART

One

Risk Management: Definitions and Objectives

A Risk Management Synthesis

Market Risk, Credit Risk, Liquidity Risk, and Asset and Liability Management

The field of risk management has undergone an enormous change in the past 40 years and the pace of change is accelerating, thanks in part to the lessons learned during the credit crisis that began in late 2006.

It hasn't always been this way in the risk management field, as Frederick Macaulay must have realized nearly 40 years after introducing the concept of duration in 1938. The oldest of the three authors entered the banking industry in the aftermath of what seemed at the time to be a major interest rate crisis taking place in the United States in the years 1974 and 1975. Financial institutions were stunned at the heights to which interest rates could rise, and they began looking for ways to manage the risk. Savings and loan associations, whose primary asset class was the 30-year fixed rate mortgage, hurriedly began to offer floating-rate mortgages for the first time. In the midst of this panic, where did risk managers turn? To the concept of mark to market and hedging using Macaulay duration? (We discuss these in Chapters 3 to 13.) Unfortunately, for many of the institutions involved, the answer was no.

During this era in the United States, a mark-to-market approach came naturally to members of the management team who rose through the ranks on the trading floor. In this era, however, and even today, chief executive officers who passed through the trading floor on their way to the top were rare. Instead, most of the senior management team grew up looking at traditional financial statements and thinking of risk on a net income basis rather than a mark-to-market basis. This is partly why Citicorp CEO Charles Prince was described in the Introduction as an executive who "didn't know the difference between a CDO and a grocery list."

As a result, the first tool to which management turned was simulation of net income, normally over a time horizon of one to two years. Given the Wall Street analyst pressures that persist even to this day, it is not a surprise that net income simulation was the tool of choice. What is surprising, however, is that it was often the only choice, and the results of that decision were fatal to a large number of U.S. financial institutions when interest rates rose to 21 percent in the late 1970s and early 1980s. One trillion dollars later, U.S. financial institutions regulators had bailed out hundreds of failed financial institutions that disappeared because of the

unhedged interest rate risk and the credit risk that was driven by the high interest rate environment.

From the perspective of 2012, risk management has taken two steps forward and one step backward. The Federal Reserve's Comprehensive Capital Analysis and Review 2012 (CCAR, 2012) appropriately focuses on a lengthy list of macroeconomic factors that contributed heavily to the credit crisis of 2006–2011. Those macro factors represent the two steps forward. The step backward was that the reporting of the CCAR 2012 results is still heavily oriented toward financial accounting and net income measures instead of mark-to-market risk measures.

From a perspective in the 1980s, the failure of net income simulation-focused risk management must have been only mildly satisfying to the fans of Frederick Macaulay and his analytical, mark-to-market approach to risk management. Even in the mid-1970s, when a moderate crisis had just occurred and hints of the crisis to come were getting stronger, resistance to mark-to-market concepts in risk management were strong and the costs to its advocates were high in career terms. The eldest of the three authors, encouraged by the U.S. Comptroller of the Currency, pushed hard for the adoption of mark-to-market-based risk management at the sixth-largest bank in the United States from 1977 to 1982. The chief financial officer, one of the brightest people in banking at the time, listened carefully but did not buy the concept, his PhD from Stanford University notwithstanding. He was so programmed by accounting concepts by that stage of his career that the mark-to-market approach to risk management was a foreign language that he would never feel comfortable speaking. The advocate of the mark-to-market approach to risk management had to change firms.

This type of financial accounting-focused executive is still in the majority today. The departed Charles Prince at Citigroup is a prime example, but the former CEOs at failed firms like Lehman Brothers, Bear Stearns, Northern Rock PLC, Royal Bank of Scotland PLC, Merrill Lynch, Bank of America, Wachovia, and Washington Mutual all suffered from the same fate. The risk that CEOs face from ignoring mark-to-market risk measurement is much larger and subtler now than it was during the savings and loan crisis. Trading activities are larger, financial instruments are more complex, and the compensation system at large financial institutions has caused the interests of traders and the interests of management and the shareholders to diverge wildly.

On the one hand, CEOs understand how critical risk management can be. Lehman CEO Dick Fuld was quoted by McDonald and Robinson (2009) in *A Colossal Failure of Common Sense*, as saying, "The key to risk management is never putting yourself in a position where you cannot live to fight another day" (317). On the other hand, the view held by the staff of management's expertise showed a clear understanding of what management did and did not understand. McDonald and Robinson note, "The World War I British Army was once described by a German general as 'lions led by donkeys.' Mike Gelband's [head of fixed income] opinion of the chain of command in Lehman's little army could scarcely have been more succinctly phrased" (234).

Every once in a while, the donkeys at the top realized that there are both lions and donkeys below. *Business Insider* quoted John Thain after he was brought in to rescue Merrill Lynch in 2008:

*“The bankers and traders dealing in CDOs didn’t understand what they were doing,” John Thain said in a recent speech. “To model correctly one tranche of one CDO took about three hours on one of the fastest computers in the United States. There is no chance that pretty much anybody understood what they were doing with these securities. Creating things that you don’t understand is really not a good idea no matter who owns it,” the former Merrill Lynch chief executive said in a speech this month, according to Financial News.*¹

The failures of many financial institutions during the credit crisis could be directly traced to a fundamental lack of understanding of risk in a large financial institution. When important macroeconomic factors change (e.g., interest rates, home prices, oil prices, foreign exchange rates, the business cycle), all of the risks that the financial institution faces are affected, even if the financial institution calls these risks by different names and attempts to manage the risks with separate groups of people and separate risk systems that go by the labels “credit risk,” “market risk,” “asset and liability management,” “liquidity risk,” and so on.

Comments from executives who were brought in to rescue failing firms show a clear recognition of this problem. Here are a few recent samples:

- John Thain, CEO brought in to rescue Merrill Lynch, said, “There were at least two major problems. One was that credit risk management was separate from market risk management, and that doesn’t make sense, because they are both permutations of the other. We are combining market and credit risk. Merrill had a risk committee. It just didn’t function. So now when we have a weekly meeting, the head of fixed income and equities show up. I show up, and the risk heads show up. It functions and functions across the businesses.”²
- Bloomberg.com, April 24, 2008, reporting on the Merrill Lynch shareholders meeting, wrote, “Ann Reese, chairwoman of Merrill’s audit committee, said the board had had ‘numerous discussions’ with management about its investments in the months before the credit crisis. The board initially didn’t realize that prices of CDOs were linked to the U.S. housing market. . . . ‘The CDO position did not come to the board’s attention until late in the process,’ said Reese, a former chief financial officer of ITT Corp. who now is co-executive director of the nonprofit Center for Adoption Policy. ‘For reasons that we have subsequently explored, there was not a sense that these triple-A securities should be included in the overall exposure to residential real estate.’”³
- The UBS AG “Shareholders’ Report on UBS’s Write-downs” stated, “Whilst there were a number of credit spread RFL [risk factor limit] limits in place, there was no RFL that specifically addressed certain factors relevant to Subprime exposure, such as delinquency rates or residential real estate price developments.”⁴
- Vikram Pandit, CEO at Citigroup after the departure of Charles Prince, said the following on the PBS’s *Charlie Rose Show*: “What went wrong is we had tremendous concentration in the sense we put a lot of our money to work against U.S. real estate. . . . We got here by lending money, and putting money to work in the U.S. real estate market, in a size that was probably larger than what we ought to have done on a diversification basis.”⁵

Besides management's failure to understand how macroeconomic factors drove risk, there were other serious problems. One of the biggest was the evolution of compensation packages for traders after the merging of commercial banking and investment banking in the 1990s. Mervyn King, Governor of the Bank of England, was quoted by the New York Times on May 6, 2008, as saying "Banks have come to realize in the recent crisis that they are paying the price for having designed compensation packages which provide incentives that are not, in the long run, in the interests of the banks themselves, and I would like to think that would change."⁶

The result of the change in compensation systems was simple to describe. Traders increasingly had an incentive to "stuff" their own institution with overpriced assets, like the super-senior tranches of collateralized debt obligations, that traders would otherwise have to sell at a loss. For example, the Financial Crisis Inquiry Commission explained, "Prince and Robert Rubin, chairman of the Executive Committee of the board, told the FCIC that before September 2007, they had not known that Citigroup's investment banking division had sold some CDOs with liquidity puts [the right to sell the CDO tranches back to Citigroup] and retained the super-senior tranches of others."⁷

Clearly, mark-to-market risk management is necessary but not sufficient to correctly assess the risk of the institution as a whole. Complete information on the positions of the institution, something Rubin and Prince did not have, is also necessary but not sufficient. The heart of this book is to outline what tools and techniques are both necessary and sufficient to have a complete and accurate understanding of the risk of the firm.

RISK MANAGEMENT: DEFINITIONS AND OBJECTIVES

In the past decade, the definition of what risk management is all about has changed dramatically for reasons we outline in the rest of this chapter. At this point, the best practice definition of risk management can be summarized as follows:

Risk management is the discipline that clearly shows management the risks and returns of every major strategic decision at both the institutional level and the transaction level. Moreover, the risk management discipline shows how to change strategy in order to bring the risk return trade-off into line with the best long- and short-term interests of the institution.

This definition knows no boundaries in terms of the nature of the institution. It applies to

- Pension funds
- Insurance companies
- Industrial corporations
- Commercial banks
- Cooperative financial institutions like savings banks
- Securities firms
- National government treasuries
- Foundations and charities

Similarly, the definition knows no political or organizational boundaries like those that have traditionally constrained the discipline of risk management from achieving this integrated approach to risk management. This definition of risk management includes within it many overlapping and inseparable subdisciplines. This is but a partial list:

- Credit risk
- Market risk
- Asset and liability management
- Liquidity risk
- Capital allocation
- Regulatory capital calculations
- Operational risk
- Performance measurement
- Transfer pricing

The primary focus of this book is to show how to execute the practice of risk management in a way that is fully integrated and makes no distinction between such subdisciplines. There should be no distinctions between them because they are simply different views of the same risk. They share the same mathematical roots, the same data needs, the same management reporting needs, and increasingly the same information technology infrastructure. (A detailed discussion is provided in Chapter 40.)

Even more important, as the incidents at Citigroup illustrate, this definition of risk management applies to all layers of management and those in positions of responsibility for the institution, including the Board of Directors, the company's auditors, and the institution's regulators, if any. All layers of management share a common obligation and, yes, burden with respect to the practice of risk management. The Basel Committee on Banking Supervision has recognized this (thousands of pages of pronouncements can be found at BIS.org), as have banking regulators all over the world, in their increasingly strict separation of duties requirements because of incidents such as the following:

Internal auditors at one of the biggest firms on Wall Street, ABC & Co., approached a trader in exotic options and asked how they could come up with an independent estimate of a certain type of volatility. The trader replied "Call Joe at XYZ Brothers at 555-1234." The auditors retired to their offices to dial Joe but before they did the trader had called Joe himself, an old drinking buddy, and told him what answer to give to the auditors' question.⁸

Conflicts of interest become starkly apparent in the execution of best practice risk management. The willingness of traders at Citigroup to provide liquidity puts on the CDO tranches they sold reflected their own best interests, not the best interests of the shareholders of Citigroup. All too often over the past 40 years, the only people expert enough to assess the risk of certain complex transactions have had a vested interest in misstating the risk, and often the valuation, of the transactions.

This incident confirms the risk that comes from an expertise differential between trading and risk control. Fortunately, the developments in financial technology over

the past 40 years have led to a sharp narrowing of the gap in spite of the incident just mentioned. More important, the incident at Citigroup confirms the need for this progress to continue. We can trace the rapid advance in risk management technology to a number of factors.

ADVANCES IN INTEGRATED RISK MANAGEMENT AND INSTITUTIONAL BARRIERS TO PROGRESS

For most of the past 40 years, there have been sharply defined “silos” that compartmentalized risk management in narrowly defined areas. Market risk was focused on a narrowly defined set of instruments accounted for on a market value basis, traded on a trading floor, with prices that were generally observable in the market. Credit risk was itself split into two subdisciplines that were often strangely disconnected:

- The decision about whether or not to originate a particular loan or private placement
- The ongoing assessment of portfolio quality on instruments that have already been put on the balance sheet

The first discipline was largely based on internal and external ratings of riskiness, which we discuss in Chapter 18, and traditional financial analysis. The second discipline was conceptually more difficult, but lower-quality analytics normally prevailed—it consisted mainly of measuring the slowly changing migration of the portfolio from one distribution of ratings in the portfolio to another. Neither discipline led to useful insights on pricing, valuation, or hedging of risk.

The asset and liability management risk silo in major banks and insurance companies was normally confined to very short-term forecasts for financial accounting–based net income. Best practice consisted of using a small number of scenarios for testing net income sensitivity to specific changes in the shape and level of yield curves. In commercial banks, at least both sides of the balance sheet were included in this exercise. In life insurance and pension funds, the analysis of income and cash flow were generally restricted to the investment portfolio—the asset side of the institution only. The actuaries ruled the liability side and were not part of the income simulation process until the last moment, when the asset side analysis and liability side analysis were pasted together. A classic sign of this bifurcation of risk is the concept of the replicating portfolio, the asset strategy that best matches the total cash flows on the liability side in n scenarios. This concept is used in insurance and nowhere else because of the huge political Great Wall of China between the actuaries and everyone else in the insurance business. If that wall came down, there are many more accurate and efficient ways to set risk strategy than the replicating portfolio concept.

Performance measurement was another silo with substantial differences among types of institutions, even though their balance sheets were nearly identical in composition. To a pension plan, performance measurement applied only to the asset side of the organization and it referred solely to whether the investment portfolio, sector

by sector, outperformed a specific benchmark index. To a commercial bank, performance measurement meant “funds transfer pricing,” the assignment of financial accounting profits to each asset and liability on the balance sheet of the bank and the subsequent compilation of financial accounting profits for each business unit. This process was further complicated by capital allocation that assigned capital to business units even if they had no liabilities and could have used the same performance measurement approach as a pension plan. At an investment bank or on the trading floor of a commercial bank, performance measurement was simply a raw calculation of profit or loss with little or no adjustment for risk or performance versus a naïve benchmark.

In Chapters 36 to 41, we explore these institutional differences in great detail and reconcile the differences in implications and management actions. These differences across institutions and industry segments are large, even though the financial institutions and major corporate treasury operations we are discussing are very similar to each other.

A number of key breakthroughs in financial theory, financial instruments, and computer science have proven that the silo approach to risk management is no longer tenable. Furthermore, these developments have shown that the differences in risk management approaches by institutional type are unnecessary and often counterproductive.

The key events in this realization are easy to summarize:

Black-Scholes Options Model

The introduction of the options pricing model by Fisher Black and Myron Scholes in 1973 showed how to value a financial security that until that point had been considered impossible to value if there was not an observable market price. Just as important, Black and Scholes, and associates such as Robert Merton, showed explicitly how to incorporate riskiness into security valuation. This innovation provided the very foundation for modern risk management. Recently, some pundits have blamed Black and Scholes for the credit crisis that began in 2006. (For an example of this genre, see Ian Stewart’s “The Mathematical Equation that Caused the Banks to Crash,” *The Observer*, February 11, 2012.) Given that Fischer Black, one of the giants of finance, passed away in August 1995, Stewart’s thesis is the financial equivalent of blaming Henry Ford, who died in 1947, for twenty-first-century automobile crashes.

Term Structure Models and the Heath, Jarrow, and Morton (HJM) Approach

In the middle of the 1970s, Robert Merton (1973) and Oldrich Vasicek (1977) introduced models of the yield curve that were consistent with the no-arbitrage foundations of Black and Scholes and based on the assumption that the short-term rate of interest varies randomly and drives the rest of the yield curve. Following an explosion of related term structure model developments in the late 1980s and early 1990s, some of it by Fischer Black, Heath, Jarrow, and Morton (HJM) published an imposing synthesis in three papers from 1990 to 1992. The HJM approach was a general solution that described what kind of random yield curve movements were consistent with no arbitrage for any

given specification for the volatility of interest rates. (The HJM approach and its special cases are a special emphasis of Chapters 6 to 14.)

Interest Rate Risk Crisis and the Credit Crisis

The interest rate spikes of the late 1970s and early 1980s created a strong need to put the new risk management technologies to work as soon as possible. The collapse of the Japanese bubble and the U.S. and European credit crisis of 2006–2011 made it quite clear that legacy approaches to risk management and the silo organization of the risk management effort did not prevent the failures of financial firms.

Advent of the Personal Computer

The advent of the personal computer and popular spreadsheet software democratized risk management analysis, bringing to the working level the potential to go far beyond the insights that could be gained by large, old-fashioned mainframe systems then dominating the financial services industry. High-quality risk management analysis could be explored and proven much more quickly and economically. Once these “proof of concept” explorations were accepted by management, risk managers had the enterprise-wide political support necessary to commit major resources to a larger, state-of-the-art system.

Expansion of the Interest Rate Derivatives Market

The interest rate derivatives market had a chicken-and-egg relationship with the exploding market in interest rate swaps, caps, and floors in the middle and late 1980s. New financial analytics for valuing these instruments made popular acceptance of these risk management tools very rapid. A recent lawsuit by Charles Schwab (filed August 23, 2011) accuses a large number of prominent financial institutions of manipulating the London interbank offered rate (LIBOR) that has been the basis for the floating rate side of interest rate swaps since the 1980s. While the outcome of the legal process is not known as of this writing, improvements in the integrity of the interest rate derivatives market are already under way.

Computer Storage Capacity and Multithreading

Rapid declines in the cost of computer storage capability meant that transaction-level risk management could become a reality in two senses. First, multiple instances of the same data could be used by many analysts, who no longer were blocked from access to the (formerly) single mainframe source of the data. Second, processing speed and the level of precision took a great leap forward with the development of multichip processors and the ability to multithread, spreading calculations and input–output over multiple computer chips working simultaneously.

Monte Carlo Simulation and 64-Bit Operating Systems

Another important development in the mid-1980s was the rapid adoption of Monte Carlo simulation as a risk management tool, further accelerating the

analysis of risk and return by major financial institutions. Faster computer chips were an essential component in making this step forward possible. The biggest boost in the popularity and power of Monte Carlo simulation came after 2000, however, with the development of the 64-bit operating system. This allowed a vastly greater amount of computer memory to be used in the course of processing, eliminating the need to store intermediate calculations on disk for later use. The impact of this development cannot be overstated.

Development of Quantitative Credit Models and the Introduction of Credit Derivatives

The next major development that accelerated the progress of risk management was the introduction of quantitative models of credit risk and the proliferation of credit derivatives such as credit default swaps, first-to-default swaps, and collateralized debt obligations. Researchers such as Robert Jarrow (1999, 2001), Jarrow and Turnbull (1995), and Duffie and Singleton (1999) constructed elegant models of credit risk on a random interest rate framework, completing the analytical integration of credit risk and interest rate risk. Tyler Shumway (2001, 2008 with S. Bharath); Jarrow and Chava (2004); and Campbell, Hilscher, and Szilagyi (2008, 2011) used logistic regression on historical databases to implement the mathematical credit models.

MEASURING THE TRADE-OFFS BETWEEN RISK AND RETURN

The implications of these developments for integrated risk management and a more sophisticated trade-off between risk and return were huge. For the first time, risk could be measured instead of just debated. In addition, alternative strategies could be assessed and both risk and return could be estimated using transaction-level detail for the entire institution and any subportfolio within the institution. Just as important, these developments dashed any rationale for the silo management of risk except that of sheer corporate politics. Market risk, credit risk, liquidity risk, and asset and liability management all use the same mathematics. They use the same data. They are caused by the same macroeconomic factors. And they impact their institutions in the same way. Any argument that they are separate and distinct from an analytical point of view is a hard one to justify from the perspective of the twenty-first century.

We hope that this book convinces any doubters by the time the last chapter is read.

WHEN BAD THINGS HAPPEN TO GOOD PEOPLE

In Chapter 2, we discuss the nature of risk and return measurement in great detail and the organizational structure consistent with best practice in risk management. Before we begin with a general discussion of risks that turned out badly, we provide some practical insights into the kinds of analytics needed to assess the likelihood, magnitude, and timing of specific risks.

These three factors are closely linked. The likelihood of a risk occurring depends on the time interval over which we measure the probability of occurrence. A given risk—the March 11, 2011, earthquake in Japan—can have different magnitudes; and the probability of an event of size X is not the same as the probability of an event of size Y . Finally, timing is everything. An earthquake tomorrow has a greater present value cost than an earthquake in the year 2525.

All of these points are obvious, and yet much of risk management has been slow to progress from simple one-period models that assume many different risks that are homogeneous in nature. For example, a common calculation in the collateralized debt obligation market is to assume there is only one period (of perhaps five years in length), that all reference names have constant default probabilities, and that all pairs of reference names have the same pairwise correlation in their default probability. This simplistic analysis seriously understates the risk of buying a CDO tranche, and it is one of the primary reasons for the massive losses in the collateralized debt obligation market. We discuss this in detail in Chapter 20.

We can understand why the likelihood, magnitude, and timing of risks have to be analyzed in a more sophisticated manner by talking about how a specific type of institution can fail and what actions would have had to be taken to save the institution.

U.S. SAVINGS AND LOAN CRISIS

The U.S. savings and loan crisis of the 1980s was predominantly due to interest rate risk, with a little interest rate–induced credit risk thrown in for spice. The savings and loans owned huge portfolios of home mortgage loans with fixed rates for 30 years. Most of the coupons on these loans were in the 8 to 10 percent range, although older loans bore much lower coupons. Passbook savings accounts were federally insured and yielded less than half of the mortgage coupons at the beginning of the crisis. Interest rates on the accounts were capped by U.S. federal regulation Q. As short-term rates began to rise, the savings and loans were reasonably complacent. They funded a fairly modest portion of their balance sheets with certificates of deposit.

As rates began to rise, two things became increasingly clear to the public, if not to management of the savings and loan associations. The first obvious point was that the value of the mortgages was declining rapidly. The second obvious point was that the savings and loans had very limited ability to hedge the interest rate risk exposure they had so that things were sure to get worse as interest rates rose. And finally, the public began to realize that the value of the assets of the savings and loans was less than the value of the savings and loans' obligations to depositors.

What happened next? Even with U.S. government deposit insurance in place, depositors voted with their feet. Interest rates on certificates of deposit of the savings and loan associations, even with U.S. deposit insurance, rose to a very substantial premium over U.S. Treasury bills that, in theory, had the same risk. Regular savings depositors withdrew their deposits in droves and savings and loans were unable to fund their cash needs in the certificate of deposit market. The greater fool theory ("even though my net worth is negative on a mark-to-market basis, someone will be dumb enough to buy my certificates of deposit") failed as a funding strategy.

LONG-TERM CAPITAL MANAGEMENT

The collapse of Long-Term Capital Management, the prominent hedge fund, in the late 1990s was a similar story. The fund took very large positions arbitraging credit spreads using a very high degree of leverage. Credit spreads moved against LTCM in a significant way. As a private institution, it was very difficult for market participants to know what capital remained and so they did the prudent thing—they declined to continue providing funding against LTCM positions, forcing LTCM to dump their very large positions at very disadvantageous terms. The LTCM partners were too smart to rely on the greater fool theory and sought help to liquidate the firm in a smooth way, but they weren't smart enough to have correctly assessed the likelihood, magnitude, and timing of the risk inherent in their portfolio.

THE 2006–2011 CREDIT CRISIS

In the introduction, we listed a long chronology of the massive losses that rippled through major financial institutions because of a series of miscues. A fall in home prices caused a sharp increase in mortgage defaults, which in turn led to massive losses in collateralized debt obligations with mortgage-backed securities as underlying reference collateral. The economics of these losses was combined with fraud in the mortgage origination and securitization process and a lack of transparency that hid the true state of affairs from both management and shareholders.

Once investors realized that CDO tranches were worth much less than the values that major firms claimed in their financial statements, short-term funding for major financial institutions with the biggest holdings disappeared quickly. A partial list of corporate victims is given in the introduction. We discuss the funding shortfalls, institution by institution, in Chapter 37.

The result of the credit crisis was the second \$1 trillion bailout of U.S. financial institutions in the past 25 years, with losses of similar magnitude throughout Europe.

A THOUSAND CUTS

These three case studies show how the magnitude of risk, likelihood of risk, and timing of risk are intimately linked. They also show that the greater fool theory rarely works for long once the dice have rolled and come up “snake eyes.”

That being said, risk “gone bad” can still destroy value without leading to the demise of the institution. The experience of both Bank of America and Citigroup during the credit crisis shows how a risk going wrong can destroy value and yet leave the institution able to go forward, wounded but still alive, thanks to generous government bailout funding.

That is why our definition of risk management earlier in this chapter is so important—institutions will sometimes lose their bets, but if management has been careful about analyzing the trade-offs between risk and return, the magnitude of the losses will be small enough to allow the institution to recover and move on. If management has not been careful, “death from a thousand cuts” can still happen.

With this in mind, we now turn to the nature of the risk-return trade-off.

NOTES

1. Yael Bizouati, "John Thain, Wild And Unleashed," *Business Insider*, October 7, 2009, http://articles.businessinsider.com/2009-10-07/wall_street/30069922_1_john-thain-merrill-stock-merrill-lynch.
2. Susanne Craig and Randall Smith, "Merrill's Risk Manager, New Chief John Thain on What Led to the Losses and Why He's Hiring Goldman Sachs Executives," *Wall Street Journal*, January 18, 2008, C1.
3. Bradley Keoun, "Merrill's Thain Calls Criticism '20/20 Hindsight' (Update4)," Bloomberg, April 24, 2008, www.bloomberg.com/apps/news?pid=newsarchive&sid=aSyltYwgVTdU&refer=home.
4. UBS AG, "Shareholders' Report on UBS's Write-downs," April 18, 2008, 19.
5. Jonathan Stempel, "Citigroup CEO: Bank 'Went Wrong' in Real Estate," Reuters, November 25, 2008, www.reuters.com/article/2008/11/26/citigroup-pandit-idUSN2531169720081126.
6. James Saft, "The Illogic of Financial-Services Compensation," *New York Times*, May 6, 2008, www.nytimes.com/2008/05/06/business/worldbusiness/06iht-rtrcol07.1.12603554.html.
7. Financial Crisis Inquiry Commission, *Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States* (Washington, DC: Government Printing Office, February 25, 2011), 260.
8. We wish to thank our colleague Tatsuo Kishi for this story about his former employer, a well-known investment banking firm.

Risk, Return, Performance Measurement, and Capital Regulation

In Chapter 1, we defined risk management as a discipline designed to help management understand the relative risks and returns from different strategies, both at the portfolio level (i.e., the perspective of the CEO or chief executive officer) and at the transaction level (i.e., the perspective of the trader, portfolio strategist, or lending officer). In this chapter, we focus on the nature of risk and return in a practical sense, not a philosophical sense. The practical definition of risk and return reflects the differences in perception of different parts of the financial world. The definitions also introduce a number of biases into perceived risk and return that can be very dangerous. We point out these dangers in this chapter and then spend the rest of the book talking about how to “avoid these pitfalls” in the conventional wisdom.

There is an ongoing debate in the financial services industry as to how risk and return should be measured. Fund managers and the asset side of pension funds are almost always on a mark-to-market basis. Insurance companies and commercial banks use a mix of mark-to-market risk/return management and a variation on financial accounting called *transfer pricing*, which separates the interest rate risk portion of net income from the non-rate-related components. We introduce the common practice of transfer pricing and its history in the last half of this chapter. We discuss the best practice of transfer pricing and performance measurement in the remainder of this book (especially in Chapter 38).

PRACTICAL QUANTIFICATION OF RISK

In December 2002, at a major risk management conference in Geneva, Nobel Prize winner Robert Merton told a story about the equity salesman from a major securities firm who was trying to convince a portfolio manager to switch the emphasis of his fund from fixed income securities to equities. Merton quoted the salesman as saying, “Over the 40-year time horizon that your fund has, you have a 99 percent probability that an all equity portfolio will return more than a portfolio of bonds. There’s no way that you’ll have less money in 40 years than the current par value of the bonds you own now.” Merton said the portfolio manager thought for a minute and then said to the salesman, “Well, if you’re right, I’d like to buy a put option from your firm to put the equity portfolio you want me to create back to your firm in 40 years at my

current portfolio value. Since your argument is so persuasive, I am sure that put option will be very cheap.” Naturally, he never heard from the salesman again.

The salesman Merton spoke of was using a tried-and-true technique, talking about expected returns and ignoring risk. This was the sales approach used by Wall Street to foist low-quality, overrated collateralized debt obligations (CDOs) on hapless investors in the lead-up to the 2006–2011 credit crisis. In the Merton story, however, the portfolio manager replied in a way that reflects the highest standard of risk management theory and practice: he asked the price of a contract that insures perfectly against the risk that concerns him.¹ This approach is being pursued very aggressively by practitioners and academics who feel, like we do, that many traditional measures of risk are much too simple. Jarrow and Purnanandam (2005), in their paper “Generalized Coherent Risk Measures: The Firm’s Perspective,” argue the virtues of put premiums like the one requested by the portfolio manager as a particularly high quality measure of risk.²

Other measures of risk have understated true risk in a more subtle way. Some of the understatement is an accidental reflection of how much more we understand now about risk than we understood when the risk measure was introduced. Sometimes, risk measures that understate risk have a longer life than they deserve when some parties have a vested interest in perpetuating the understatement.

We turn to various risk measures to point out these potential pitfalls in the next section.

PERILS AND PITFALLS IN THE MEASUREMENT OF RISK: THE IMPACT OF SELECTION BIAS

A major bank with a significant portfolio of loans in the Asia Pacific region once retained one of the authors as a consultant in assessing the risk of the portfolio. As part of the assignment, the Chief Financial Officer (CFO) of the bank requested a historical value-at-risk analysis of the bank’s 100 largest counterparties in Asia. We, of course, did the work he requested but provided the analysis in the context of a letter that explained the nature of the results. The work was done, we explained, using industry standards and best practices of value-at-risk analysis. However, the ninety-ninth percentile worst case, which we described in the results section, was in fact the “best case worst case” that the bank would experience.

“How does that help me with bank regulators?” the CFO asked, not fully appreciating that we were trying to save his job. We explained the bias in value at risk patiently to him as follows:

1. You have selected 100 of your current counterparties in Asia.
2. By definition, they are your counterparties now because they have not defaulted.
3. You asked us to do a historical analysis of value at risk using the bonds of these counterparties over the past 12 months.
4. We have done that and produced a ninety-ninth percentile “worst case” for you as you requested, *but* . . .
5. It is the best case–worst case scenario because we are doing a historical analysis on 100 names over a period in which none of them defaulted. The very nature of the risk measure and the nature of the counterparties guarantee that this

seriously underestimates the credit risk in this portfolio. If the CFO wanted a more accurate assessment of the risk, the test should have either been based on 100 counterparties randomly chosen as of a year before (and indeed some of them subsequently defaulted) or we should have done a forward-looking simulation that allows for default on a name-by-name basis.

He conceded the point and hedged himself with the board, avoiding the career-ending scenario of a risk that could have turned out to be much worse than the advertised worst case.

We now turn to some of the most common examples of this selection bias in measuring risk.

BIASES IN RETURN VS. A RELATIVE BENCHMARK

In the fund management and pension fund business, most portfolio managers are judged by measures of their plus alpha or risk-adjusted return versus a predefined benchmark. An example might be a fund manager who manages large-cap portfolios of common stock and is given the Standard & Poor's 500 stock index as his benchmark index. This fund manager is told that he must keep a tracking error of less than x percent versus the index and that he is expected to earn a return in excess of the return on the S&P 500. This is a very common kind of risk-and-return measurement system on the "buy side" of financial markets.

What is wrong with this as a risk-and-return measurement system? Everything that was wrong in the story in the previous section.³ We can see this by taking apart this risk management system piece by piece.

First, how is tracking error measured? There are many variations, but in most cases a regression is run on two data series: the stock price returns for ABC Company and the returns on the S&P 500 index. By definition, tracking error will only be known for ABC Company if that stock price series exists (i.e., if the company has not been delisted from the exchange due to bankruptcy). Our risk measure for ABC Company is a no-bankruptcy-risk risk measure. Similarly, what about the bankruptcy risk of the S&P 500? Generally speaking, as a company's credit quality (and/or market capitalization) declines, they are dropped from the S&P 500. This happens after the close of business, and on the subsequent opening, the price of the stock drops considerably in a way that is unhedgeable unless one predicts the change in S&P 500 composition. This drop does not affect the S&P index, but the holder of the common stock of the company dropped from the index will suffer a large loss. The composition of the S&P 500 has included some firms that went bankrupt while they were in the index. We suspect, however, that the index is ultimately restated to offset the impact of these failures.

As you can see, the most common performance measure used in the equity markets ignores credit risk almost completely, even though on average between 1.00 percent and 1.5 percent of listed U.S. companies have defaulted annually over the 1963–2012 time period. We are substantially understating the risk component of this traditional performance measure.

How does the risk of bankruptcy affect our fund manager if the risk occurs? It depends on whether:

1. ABC company is in the manager's portfolio or not.
2. ABC company is in the S&P 500 or not.
3. ABC goes bankrupt while it is in the index or is dropped from the S&P 500 for credit quality reasons before going bankrupt.

If ABC Company is dropped from the S&P 500 index, the manager will have a decline of as much as 20 percent or more when the “drop” is announced, and yet there will be no corresponding change in the index. The manager will have a big negative alpha because he (or she) couldn't get out of ABC Company until after the announcement that ABC Company is not in the index any longer. If ABC, in fact, goes bankrupt while still in the index, the drop in stock price could be as much as 100 percent, but at least the index will also be affected. We strongly recommend that managers with an equity index as a benchmark independently confirm how a bankruptcy affects the index.

If the manager owns a stock that goes bankrupt and it is not in the S&P 500 index, he has a negative alpha versus the index that tracking error completely ignores. Even if the company is merely dropped from the index, the manager will suffer a negative 20 percent or so stock price decline that is not offset by the movement of the index and that is not reflected in measured “tracking error.” This same kind of problem affects fixed income performance measurement versus a benchmark such as the former Lehman Brothers Government Bond Index (now renamed after Barclays Capital), but in more subtle ways.

We next turn to another common measure of risk with similar pitfalls.

HISTORICAL VALUE AT RISK: SELECTION BIAS AGAIN

There are as many approaches to value at risk as there are risk management analysts, but in general value-at-risk (VaR) calculations fall into three basic categories:

1. A historical VaR that measures dollar price changes for specific securities over a period in history
2. A variance/covariance approach that measures the variances and covariances of past returns and implies (from this implicit assumption of normally distributed returns) the n th percentile worst case
3. A forward-looking simulation of returns over the desired time horizon

The comments of this section apply only to methods 1 and 2.

What are the concerns about these calculations, which have been much used and much discussed over the past decade? Credit risk is completely ignored. We have the selection bias discussed earlier in this chapter because, by definition, none of the current securities holdings would include the securities of a bankrupt firm.

Still, there are two other serious problems with these two approaches to VaR that are just as serious:

1. They ignore all cash flows between time zero and the VaR horizon (10 days, 1 month, 3 months, 1 year) since these approaches are effectively a single-period

analysis that implicitly assumes you own the same things at the end of the period that you owned at time zero

2. They cannot answer the question posed by the portfolio manager in Robert Merton's anecdote: what is the hedge and how much does it cost?

Both of these drawbacks seriously affect VaR as a risk measure. In the end, on a fully credit adjusted basis, we need to know what the hedge is and how much it costs—just as Merton and Jarrow advise.

MONTE CARLO–BASED VALUE AT RISK

Many thoughtful risk managers have sought to bring credit risk into traditional historical or variance–covariance value at risk by either simulating default/no default during the VaR time period or by using a transition matrix approach.

As an “interim fix,” a transition matrix approach is a step forward. But again, some key problems remain that understate the risk remain:

- Common macroeconomic factors that drive correlated default are not specified. This will generally result in a significant underestimate of the “fat tails” from credit losses. Van Deventer and Imai (2003) discuss the impact of macroeconomic factors in great detail and we turn to them later in this book. The obvious lesson of the collapse of the Japanese bubble and the 2006–2011 credit crisis is that macroeconomic factors affect everything. Prior to 2006, a VaR calculation that did not simulate home price movements accurately would have been grossly misleading.
- The default probabilities are implicitly held constant over the time period. In reality, these default probabilities risk and fall through the credit cycle, and holding them constant again understates the level of credit risk implicit in the portfolio.
- The timing of defaults during the period is ignored, when in fact the exact timing is very important (see the next section for the reasons).
- Interim cash flows are again ignored (unless it is a multiperiod simulation).
- What's the hedge? Again, we can't answer the portfolio manager's question in the Robert Merton anecdote.

EXPECTED LOSSES ON TRANCHES OF COLLATERALIZED DEBT OBLIGATIONS

In the past decade, the market for credit default swaps and collateralized debt obligations has exploded and then collapsed completely. Ratings have been assigned and collateralized debt obligation (CDO) tranches have been marketed on the basis of the “expected loss.” Thus, with amazing speed, these CDO tranches were marked down to “junk” or default.

The focus up through 2007 was purely on expected losses and ignored the fat tails that are typical of the worst part of the credit cycle. The introduction summarizes

some of the tuition payments made by large international banks and securities firms in this regard.

As always, to understand the risk completely, we need the full probability distribution of the losses and, most important, we need to be able to stress test it in order to calculate the hedge.

We turn now to the other side of the risk-return horizon.

MEASURING RETURN: MARKET VS. ACCOUNTING RETURNS

The nature of risk measurement is complex enough, as we have seen in previous sections, but there is more complexity involved: the definition of return. There are a few key determinants of what approach is used to return measurement:

- The nature of the institution involved
- The nature of the regulatory agencies involved
- The availability of observable market prices
- The time horizon of greatest concern

For managers of common stock portfolios, the decisions are pretty straightforward. The market prices of the common stock are observable, regulation is market-value based, and there is an observable daily benchmark. Moreover, industry expert bodies, such as the CFA Institute (formerly the Association for Investment Management and Research), have instituted a step-by-step approach to daily performance benchmarking.

In other industries, things are more complex. How do you deal with life insurance policy liabilities and an asset mix of equities, bonds, real estate, and commodities? What would you do if your liabilities were pension fund obligations? What if they were passbook savings?

The result is often a system of splitting risk and return using a technique called transfer pricing, which in this context has a much different meaning than its meaning to tax experts. We turn now to an introduction of how commercial banks have established this practice and what the current state of the art is. (We discuss the mechanics of transfer pricing in Chapter 38 in light of the analytical tools we discuss in the first 37 chapters of this book.)

INTRODUCTION TO TRANSFER PRICING: EXTRACTING INTEREST RATE RISK IN A FINANCIAL ACCOUNTING CONTEXT

The subject of transfer pricing represents one of the greatest differences in asset and liability management practice between large banks internationally and between banks and the fund managers mentioned in the previous section. In the United States and Australia, the so-called “matched maturity transfer pricing system” is accepted as an absolute necessity for the successful management of a large financial institution. In Japan, by way of contrast, the first bank to use multiple rate transfer pricing adopted a very simple system more than 20 years after Bank of America initiated a much more sophisticated system in 1973. As of this writing, transfer pricing has

become an accepted banking discipline in a large number of countries, but it is still rare in the insurance industry.

The rationale for transfer pricing as an alternative to the investment management business' reliance on a mark-to-market approach is a simple one. With employee counts of tens of thousands of people and hundreds of business units that invest in assets with no observable price, it is literally impossible to manage a commercial bank on a mark-to-market basis at the working level. Transfer pricing is the trick or device by which the rules of the game are changed so that (1) actions are consistent with a mark-to-market approach, but (2) the vehicle by which the rules of the game are conveyed is in financial accounting terms.

We start by presenting a selective history of the transfer pricing discipline in the United States. (In Chapter 38, we provide a detailed description of the current "common practice" in transfer pricing, with examples of its practical use.)

Bank of America, 1973–1979

Interest rate risk management has advanced in bursts of energy that usually follows a period of extremely high interest rates in the U.S. and other markets. Interest rate risk has, up until recently, been the main focus of transfer pricing. In the early 1970s, the entire home mortgage market in the United States was made up of fixed rate loans. The floating-rate mortgage had yet to be introduced. Interest rates were beginning to be deregulated, and the first futures contracts (on U.S. Treasury bills) had not yet been launched. In this environment, a rise in market interest rates created a double crisis at banks with a large amount of retail business: consumer deposits flowed out of banks (which had legal caps on deposit rates at that time) into unregulated instruments such as Treasury bills or newly introduced money market funds. Simultaneously, the banks suffered from negative spreads on their large portfolios of fixed rate loans.

A spike in interest rates beginning in 1969 triggered a serious reexamination of financial management practice at the largest bank in the United States (at the time), the Bank of America. A team of executives that included C. Baumhefner, Leland Prussia, and William "Mack" Terry recognized that changing interest rates were making it impossible for existing management "accounting" systems to correctly allocate profitability among the bank's business units and to set responsibility for managing interest rate risk. Until that time, the bank had been using a single internal transfer pricing rate that was applied to the difference between assets and liabilities in each business unit to allocate some form of interest expense or interest income in a way that the units' assets and liabilities were equal. This rate, known as the pool rate system for transfer pricing, was a short-term interest rate at Bank of America, calculated as a weighted average of its issuance costs on certificates of deposit.

The liabilities of such a single-rate pool rate system are outlined in later chapters of this book. Given the sharp changes in rates that were beginning to look normal in the United States, senior management at the Bank of America came to the conclusion that the existing system at the bank was dangerous from an interest rate risk point of view and made proper strategic decision making impossible. Strategic decision making was handicapped because the single rate transfer pricing system made it impossible to know the true risk-adjusted profitability of each business unit and line of business. The numbers reported by the existing system had mixed results due to

business judgment (“skill”) and results due to interest rate mismatches (“luck,” which could be either good or bad).

Although the bank had a very sizable accounting function, the responsibility for designing the new transfer pricing system was given to the brilliant young head of the Financial Analysis and Planning Department of the bank, MIT graduate Wm. Mack Terry. This division of labor at the bank was the beginning of a trend in U.S. banks that resulted in the total separation of the preparation of official financial reports, a routine and repetitive task that required a high degree of precision, and a low level of imagination, from management information (not managerial “accounting”) on how and why the bank was making its money. Mack Terry, who reported directly to CFO Lee Prussia, took on the new task with such vigor that most knowledgeable bankers of that era acknowledge Mack Terry as the father of the new discipline of “matched maturity transfer pricing.”

Mack Terry and the bank were faced by constraints that continue to plague large banks today:

- The bank was unable to get timely information from loan and deposit mainframe computer application systems regarding the origination, maturity, payment schedules, and maturities of existing assets and liabilities.
- The bank had neither the time nor the financial resources to develop a new transfer pricing system on a mainframe computer.

Personal computers, needless to say, were not yet available as management tools.

The danger of future rate changes and incorrect strategic decisions was so great that bank management decided to make a number of crude but relatively accurate assumptions in order to get a “quick and dirty” answer to the bank’s financial information needs. The fundamental principle developed by the Financial Analysis and Planning Department was the “matched maturity” principle. A three-year fixed rate loan to finance the purchase of an automobile would be charged a three-year fixed rate. A 30-year fixed rate mortgage loan would be charged the cost of fixed rate 30-year money. At other U.S. banks, it was more common to use the bank’s average cost of funds or its marginal cost of new three-month certificates of deposit. The average cost of funds method had obvious flaws, particularly in a rising rate environment, and it ultimately contributed to the effective bankruptcy of Franklin National Bank and many U.S. savings and loan associations. (We discuss how these transfer pricing rates are calculated in later chapters.)

The Financial Analysis and Planning (FAP) Department team agreed that the bank’s “marginal cost of funds” represented the correct yield curve to use for determining these matched maturity transfer pricing rates. The bankers recognized that the bank had the capability to raise small amounts of money at lower rates but that the true “marginal” cost of funds to the bank was the rate that would be paid on a large amount, say \$100 million, in the open market. This choice was adopted without much controversy, at least compared to the controversy of the “matched maturity” concept itself. In more recent years, many banks have used the London Interbank Offered Rate (LIBOR) and the interest rate swap curve as a proxy for their marginal costs of funds curve. It is clear from the perspective of today, however, that the LIBOR-swap curve is a poor substitute for an accurate marginal cost of funds curve.

One of the first tasks of the FAP team was to estimate how the internal profitability of each unit would change if this new system were adopted. It was quickly determined that there would be massive “reallocations” of profit and that many line managers would be very upset with the reallocation of profit away from their units. In spite of the controversy that management knew would occur, management decided that better strategic decision making was much more important than avoiding the short-run displeasure of half of the management team (the other half were the ones who received increased profit allocations).

Once a firm decision was made to go ahead with the new system, implementation decisions had to be dealt with in great detail. The system’s queue was so long at the bank that it was impossible to keep track of transfer pricing rates on a loan-by-loan or deposit-by-deposit basis. The first unit of measurement, then, was to be the portfolio consisting of otherwise identical loans that differed only in maturity and rate, not in credit risk or other terms. In addition, the bank (to its great embarrassment) did not have good enough data on either its own loans or its historical cost of funds to reconstruct what historical transfer pricing rates would have been for the older loans and deposits making up most of the bank’s balance sheet. Estimates would have to do.

These estimates were at the heart of the largest of many internal political controversies about the matched maturity transfer pricing system. Much of the fixed rate real estate portfolio was already “underwater”—losing money—by the time the transfer pricing system was undergoing revision. If the current marginal cost of funds were applied at the start of the transfer pricing system to the mortgage portfolio, the profitability of the portfolio would have been essentially zero or negative. On the other hand, Mack Terry and his team recognized that in reality this interest rate risk had gone unhedged and that the bank had lost its interest rate bet. Someone would have to “book” the loss on the older mortgage loans.

Ultimately, the bank’s asset and liability management committee was assigned a funding book that contained all of the interest rate mismatches at the bank. The funding book was the unit that bought and sold all funds transfer priced to business units. At the initiation of the matched maturity system, this funding book was charged the difference between the historical marginal cost of funds that would have been necessary to fund the mortgage portfolio on a matched maturity basis and the current marginal cost of funds. This “dead weight loss” of past management actions was appropriately assigned to senior management itself.

Because of the lack of data, the bank’s early implementation of the matched maturity system was based on the use of moving average matched maturity cost of funds figures for each portfolio. While this approximation was a crude one, it allowed a speedy implementation that ultimately made the bank’s later troubles less severe than they otherwise would have been.

Finally, the system was rolled out for implementation with a major educational campaign aimed at convincing lending officers and branch managers of the now well-accepted discipline of *spread pricing*, pricing all new business at a spread above the marginal cost of matched maturity money.⁴ Controversy was expected, and expectations were met. A large number of line managers either failed to understand the system or didn’t like it because their reported profits declined. Soon after the announcement of the system, Mack Terry began receiving anonymous “hate mail” in the interoffice mail system from disgruntled line managers.

Nonetheless, senior management fully backed the discipline of the new system, and for that the Bank of America receives full credit as the originator of the matched maturity transfer pricing concept. In an ironic footnote, the bank suffered heavily from an interest rate mismatch as interest rates skyrocketed in the 1979–1980 time period. The transfer pricing system made it clear that senior management was to blame since the asset and liability management committee had consciously decided not to hedge most of the interest rate risk embedded in the bank's portfolio.

First Interstate, 1982–1987

In the years to follow, a number of banks adopted the matched maturity transfer pricing system. The oral history of U.S. banking often ranks Continental Illinois (later acquired by the Bank of America) as the second bank to move to a matched maturity transfer pricing system, not long after the Bank of America implemented the idea. At most banks, however, the idea was still quite new and the early 1980s were largely consumed with recovering from the interest rate-related and credit risk-related problems of the 1979–1981 period. By the mid-1980s, however, banks had recovered enough from these crises to turn back to the task of improving management practices in order to face the more competitive environment that full deregulation of interest rates had created. First Interstate Bancorp, which at the time was the seventh-largest bank holding company in the United States, approached the transfer pricing problem in a manner similar to that of many large banking companies in the mid-1980s.

First Interstate's organization was much more complex than Bank of America's since the company operated 13 banks in nine western states, all of which had separate treasury functions, separate management, separate government regulation, and separate systems.⁵ In addition, the bank holding company's legal entity, First Interstate Bancorp, had a number of nonbank subsidiaries that required funding at the holding company (parent) level. Most U.S. bank holding companies consisted of a small parent company whose dominant subsidiary was a lead bank that typically made up 90 percent of the total assets of the company in consolidation. The lead banks were almost always considered a stronger credit risk than the parent companies because of the existence of Federal deposit insurance at the bank level but not the parent level and because of the richness of funding sources available to banks compared to bank holding companies.

In the First Interstate case, things were more complex. The lead bank, First Interstate Bank of California, represented only 40 percent of the assets of the holding company and, therefore, its credit was generally felt by market participants to be weaker than that of the parent. Moreover, the First Interstate Banks did not compare funding needs, and as a result, it was often said that a New York bank could buy overnight funding from one First Interstate Bank and sell it to another First Interstate Bank at a good profit. The transfer pricing system at First Interstate had to cure this problem as well as address the correct allocation of profits and interest rate risk as in the Bank of America case.

Management took a twofold approach to the problem. At the holding company level, the corporate treasury unit began "making markets" to all bank and nonbank units within the company. Because First Interstate was much more decentralized than the Bank of America, funds transfers between the holding company and subsidiaries were voluntary transfers of funds, not mandatory. In addition, since each unit was a

separate legal entity, a transfer pricing transaction was accompanied by the actual movement of cash from one bank account to another.

The holding company transfer pricing system began in early 1984 under the auspices of the holding company's funding department. The department agreed to buy or sell funds at its marginal cost at rates that varied by maturity from 30 days to 10 years. No offers to buy or sell funds were to be refused under this system. The transfer pricing system, since it was voluntary, was received without controversy and actually generated a high degree of enthusiasm among line units. For the first time, the units had a firm "cost of funds" quotation that was guaranteed to be available and could be used to price new business in line units on a no-interest-rate risk basis. Demand from line units was very strong, and the parent company became an active issuer of bonds and commercial paper to support the strong funding demand from both bank and nonbank subsidiaries.

Among bank subsidiaries, the "on-demand" transfer pricing system had the effect of equalizing the cost of funds across subsidiary banks. High cost of funds banks immediately found it cheaper to borrow at the lower rates offered by the parent company. Regulatory restrictions kept the parent company from borrowing from subsidiary banks, but there was an equivalent transaction that achieved the same objective. After the transfer pricing system had been in operation for some months, the parent company had acquired a substantial portfolio of certificates of deposit of subsidiary banks. These certificates of deposit could be sold to other banks within the system. By selling individual bank certificates of deposit to other First Interstate banks, the holding company reduced the size of this portfolio and effectively borrowed at its marginal cost of funds, the yield it attached to the certificates that it sold.

The transfer pricing system at the holding company did have implementation problems of a sort. Generally, rates were set at the beginning of a business day and held constant for the entire day. The parent company soon noticed that borrowings from affiliates, and one subsidiary in particular, would increase when open market rates rose late in the day, allowing subsidiaries to arbitrage the parent company treasury staff by borrowing at the lower rate transfer price set earlier in the day. This got to be a big enough problem that rates for all borrowings above \$10 million were priced in real time. Most large banks use a similar real-time quotation system now for pricing large corporate borrowings. What will surprise many bankers, however, is that in-house transactions ultimately also have to be priced in real time in many cases.

At the same time that the parent company was implementing this system, the parent company's asset and liability management department and the financial staff of First Interstate Bank of California (FICAL) began to design a Bank of America-style matched maturity transfer pricing system for FICAL. Personal computer technology at the time did not permit PCs to be used as the platform, so the company undertook a very ambitious mainframe development effort. After a complex design phase and a development effort that cost close to \$10 million and two years of effort, the system was successfully put into action with considerably less controversy than in the Bank of America case. Such a system today would cost much, much less from a third-party software vendor.

The biggest practical problem to arise in the FICAL transfer pricing system was a subtle one with significant political and economic implications. During the design phase of the system, the question was raised about how to handle the prepayment of

fixed rate loans in the transfer pricing system. The financial management team at First Interstate dreaded the thought of explaining option-adjusted transfer pricing rates to line managers and decided to do the following: internal transfer pricing would be done on a nonprepayable basis. If the underlying asset were prepaid, then the Treasury unit at the center of the transfer pricing system would simply charge a mark-to-market prepayment penalty to the line unit, allowing it to extinguish its borrowings at the same time that the asset was prepaid. This simple system is standard practice in many countries, including the retail mortgage market in Australia.

This decision led to unforeseen consequences. During the latter half of the 1980s and continuing into the 1990s, interest rates declined in the United States and were accompanied by massive prepayments of fixed rate debt of all kinds. As a result, the FICAL transfer pricing system's reported profits for line units were soon dominated by huge mark-to-market penalties that line managers didn't have control over and generally didn't understand.

Bank management quickly moved to a transfer pricing system that allowed for "costless prepayment" by incorporating the cost of a prepayment option in transfer prices. This trend is firmly established as standard practice in most banks today. With this political and organizational background on transfer pricing in mind, we now put it in the context of performance measurement and capital regulation. (We return to the mechanics of transfer pricing in Chapter 38.)

PERFORMANCE MEASUREMENT AND CAPITAL REGULATION

The primary purpose in this chapter and subsequent chapters is to present common practice for performance measurement in various wings of the financial services business and to contrast the approaches used by different institutions. When we see institution A and institution B managing similar risks, but using different approaches to the measurement and management of risk, we will carefully note the differences and seek an explanation. Often, institutional barriers to change delay the synthesis and common practices that one would expect from two different groups of smart people managing the same kind of risk.

This difference is particularly stark in the role of capital in risk management and performance measurement. Commercial banks are extensively focused on capital-based risk measures, and very few financial services businesses are. Why? We begin to answer this question with some additional perspectives on the measurement and management of risk. We then talk about managing risk and strategy in financial institutions, business line by business line, and how capital comes into play in this analysis. We then turn to the history of capital-based risk regulations in the commercial banking business and discuss its pros and cons.

PERSPECTIVES ON MEASURING RISK: ONE SOURCE OF RISK OR MANY SOURCES OF RISK?

Risk management has been an endless quest for a single number that best quantifies the risk of a financial institution or an individual business unit. In Chapter 1, we discussed how Merton and Jarrow suggest that the value of a put option that fully

insures the risk is the current best practice in this regard. This risk measure also provides a concrete answer to the question that best reveals the strengths or weaknesses of a risk management approach: “What is the hedge?”

Over the past 50 years, the evolution of risk management technology toward the Merton and Jarrow solution has gone through two basic steps:

1. Assume there is a single source of risk and a risk measurement statistic consistent with that source of risk
2. Recognize that in reality there are multiple sources of risk and revise the measure so that we have an integrated measure of all risks and sources of risk

INTEREST RATE RISK MANAGEMENT EVOLUTION

In Chapters 3 to 14, we review traditional and modern tools of interest rate risk management. The first sophisticated interest rate risk management tool was the Macaulay (1938) duration concept. In its traditional implementation, managers who own fixed income assets shift a yield curve up or down by a fixed percentage (often 1 percent) and measure the percentage change in the value of the portfolio. This measure shifts yield curves at all maturities by the same amount, essentially assuming that there is only one type of interest rate risk, parallel movements in the yield curve. As we show in Chapter 3, that is not the case in reality and, therefore, the duration concept measures some but not all interest rate risks.

The next step forward in sophistication was to recognize that yield curves in fact move in nonparallel ways. The term structure models that we discuss in Chapters 6 to 14 allow analysts to move one or more key term structure model parameters (typically including the short-term rate of interest) and see the response of the full yield curve. This provides a more sophisticated measure of duration that allows for nonparallel shifts of the yield curve.

Fixed income managers in both commercial banks and other financial services companies recognize, however, that both of these approaches are abstractions from reality. As we show in Chapter 3, there are multiple factors driving any yield curve, whether it is a credit-risk-free government yield curve or a yield curve that embeds the default risk and potential losses of a defaultable counterparty.

In recognition of these n factors driving the yield curve, fixed income managers have overlaid practical supplements on the theoretical models discussed above:

1. Interest rate sensitivity gaps (see legacy approaches to interest rate risk in Chapter 12) that show the mismatches between the maturities of assets and liabilities (if any) period by period. These gaps provide visibility to managers that allow them to implicitly reduce the risk that interest rates in gap period K move in a way different from what the assumed duration or term structure model specifies.
2. Multiperiod simulation of interest rate cash flows using both single scenario interest rate shifts and true Monte Carlo simulation of interest rates. Managers can measure the volatility of cash flows and financial accounting income period by period, again to supplement the interest rate models that may understate the true complexity of interest rate risk movements

What is the equivalent of the Merton and Jarrow “put option” in the interest rate risk context? It is the value of a put option to sell the entire portfolio of the financial institution’s assets and liabilities at a fixed price at a specific point in time. From a pension fund perspective, a more complex string of options that guarantee the pension fund’s ability to provide cash flow of $X(t)$ in each of n periods to meet obligations to pensioners would be necessary. (We return to this discussion in more detail in Chapters 36 to 41.)

EQUITY RISK MANAGEMENT EVOLUTION

Risk management in the equity markets began in a similar fashion with the powerful and simple insights of the capital asset pricing model. This model initially suggested that all common stocks had their returns driven by a single common factor, the return on the market as a whole, and an idiosyncratic component of risk unique to each stock. By proper diversification, the idiosyncratic component of risk could be diversified away. The single risk factor, the exposure to changes in return on the market, was measured by the beta of both individual common stocks and on the portfolio held by the equity manager.

As in the interest rate case, risk managers quickly seized on the insights of the model and then generalized it to recognize there are multiple drivers of risk and the multiperiod nature of these risk drivers. Earlier in this chapter, we introduced the need to add credit risk to the multiperiod equity risk models that reflect current “common practice.”

OPTION RISK MANAGEMENT EVOLUTION

The Black-Scholes option model brought great hope to options risk managers in much the same way that the capital asset pricing model did. Options portfolio managers (and managers of broader portfolios of instruments with options characteristics) quickly adopted the delta of an options position, the equivalent position in the underlying stock, as an excellent single measure of the risk of an options position.

With experience, however, analysts realized that there was a *volatility smile* that is, a graph showing that the Black-Scholes options model implies a different level of volatility at different strike prices with the same maturity. Analysts have tended to “bend” the Black-Scholes options modeling by using different volatilities at different strike prices rather than looking for a more elegant explanation for the deviation of options prices from the levels predicted by the Black-Scholes model. In effect, analysts are using the Black-Scholes model as an n -factor model, not a single-factor model with the stock price as the risk driver with a constant volatility. See Jarrow (2011) on the dangers of such an ad hoc approach.

CREDIT RISK MANAGEMENT EVOLUTION

The integration of credit risk, market risk, asset and liability management, liquidity risk management, and performance measurement is one of the central themes of this

book. Analysts have moved from traditional credit analysis on a company-by-company basis to more modern technology over the past 25 years. The more traditional credit analysis was a single-company analysis that involved financial ratios and occasionally simulation to estimate a default probability or rating that summarized the risk level of each company, one by one.

Now the kind of quantitative credit models that we discuss in Chapters 15 to 18 allow us to recognize the M common macroeconomic risk factors that cause correlated defaults among a portfolio of companies. These insights allow financial institutions to move to a macro-hedging program for a large portfolio of credits for the first time.

It also provides a framework for the quantification of risk as Merton and Jarrow suggest, with a put option on the credit portfolio.

We now turn to the management process in commercial banks, insurance companies, and pension funds, and analyze how and why capital came to play a role in commercial banking risk management and strategy but not in other institutions.

MANAGING RISK AND STRATEGY, BUSINESS BY BUSINESS

One of the major differences between financial institutions is the number of people who are involved in asset generation. In a life insurance company, a vast majority of the people are involved in the generation, servicing, and pricing of life insurance policies. On the asset side, investment activities are much more “wholesale” and the number of people involved per one million dollars of assets is much less. In an asset management company, the total staff count is also generally less than 1,000 people. The same is true for pension funds. Bank of America, by contrast, had a staff of 288,000 at the end of 2011. For very large commercial banks, the number of people involved is very large because almost all banks these days are heavily involved in the retail generation of assets and liabilities.

There is another important difference between commercial banks and the rest of the financial community. For asset managers, pension funds, and insurance companies, most of the people managing the assets of the company work in one central location. In a commercial bank, there can be thousands of branches involved in asset generation across a broad swath of geography.

For that reason, in what follows we will largely discuss management of risk and business strategy in a banking context because most of what happens in other financial institutions is a special case of the banking business, including even the actuarial nature of liabilities of all of these types of financial institutions.

RISK AND STRATEGY MANAGEMENT IN A COMPLEX FINANCIAL INSTITUTION

From this chapter’s introduction, we know that in the banking business various parts of the institution are managed by one or more bases from both a risk and a return point of view:

- A mark-to-market orientation for the institution as a whole, the trading floor, and the transfer pricing book

- A financial accounting point of view for both the institution as a whole and almost all nontrading business units

Moreover, as we say in the introduction to this chapter, multiple risk measures are applied depending on the business unit involved:

- *Complex interest rate risk measures* are applied to the institution as a whole, the trading floor, and the transfer pricing book.
- *Less sophisticated interest rate risk measures* are usually sufficient for most business units, because the transfer pricing process in most institutions removes the interest rate risk from line units. Depending on the financial institution, the basis risk (a lack of synchronicity between the asset floating-rate pricing index and market rates) usually stays with the business unit if the head of that business has the responsibility for selecting that pricing index for the bank.
- *Yes/no credit risk measures* are applied at the “point of sale” in asset-generating locations, often controlled by a central location
- *Portfolio credit risk measures* are applied by the business unit responsible for maintaining or hedging credit risk quality after the generation of assets.
- *Market risk and credit risk measures* are almost applied together on the trading floor because so much of the market-risk-taking business has taken the form of credit risk arbitrage, even in the interest rate swap market. The percentage of derivatives contracts where at least one side is held by one of the top ten dealers is astronomical, especially as the bank merger trend continues in the United States.
- *Return measures are stated either on a market-based total return basis or a net income basis*, usually after interest rate risk is removed via the transfer pricing system (see Exhibit 2.1).

We defined risk management in Chapter 1 as the discipline by which management is made aware of the risk and returns of alternative strategies at both the transaction level and the portfolio level. Once we have coherent measures of risk and return, management often runs across situations like the following graph, where the risk and returns of three alternative business units A, B, and C are plotted. Unit C is




EXHIBIT 2.1 Comparing Risk-and-Return Profiles

clearly superior to Unit B because it has a higher expected return than B. Similarly, Unit A is superior to Unit B because Unit A has the same expected return but less risk than Unit B. It is much harder to compare Units A and C. How do we select business units that are “overachievers” and those that are “underachievers”? This is one of the keys to success in financial institutions management.⁶ The Sharpe ratio is commonly cited as a tool in this regard, but there is no simple answer to this question. We hope to provide extensive tools to answer this question in the remainder of the book. There are a number of hallmarks to best practice in this regard:

- A tool that cannot provide useful information about which assets are the best to add at the *transaction level* cannot be useful at the portfolio level, since the portfolio is just a sum of individual transactions
- *Buy low/sell high* contains a lot of valuable wisdom as simple as it seems. If the financial institution is offered the right to buy asset 1, with a market value of 100, at a price of 98, it should generally say yes.
- *If the bank has to choose* between asset 1 and asset 2, which has a market value of 100 and can be purchased for a price of 96, the bank should buy asset 2. If the bank is solvent and its risk is not adversely affected at the portfolio level, the bank should buy both. The market price summarizes risk and return into one number.
- Risk-return tools should provide value from *both a risk management perspective and the perspective of shareholder value creation*. These tools are not separate and distinct—they are at the heart of good management.

WHAT CAUSES FINANCIAL INSTITUTIONS TO FAIL?

In deciding at the transaction and portfolio levels which assets to select for a financial institution as a whole, the criteria for selection are not independent of which assets and liabilities are already on the books of the financial institution.⁷

The dominant reason for financial institutions to fail in the twenty-first century is credit risk, now that the interest rate risk tools and techniques discussed in the next few chapters are used with care. Even before the obvious lessons of the 2006–2011 credit crisis, this had been confirmed by many studies of bank failures in the United States, most importantly the Financial Institutions Monitoring System introduced by the Financial Institutions Monitoring System implemented by the Board of Governors of the Federal Reserve in the United States in the mid-1990s.⁸ The effort to better predict the failure of financial institutions and the subsequent losses to deposit insurance funds gained new momentum with the announcement of its new Loss Distribution Model by the Federal Deposit Insurance Corporation on December 10, 2003.⁹ The FDIC Loss Distribution Model, with our colleague Robert A. Jarrow as lead author, correctly predicted that the FDIC insurance fund was grossly underfunded a full three years before the advent of the 2006–2011 credit crisis. We discuss the loss distribution model in great detail in later chapters.

Given the importance of credit risk in the failure of financial institutions, an integrated treatment of credit risk, market risk, asset and liability management, and performance measurement is critical or we may totally miss the boat from a risk perspective. In part for this reason, capital has become a critical component of both

regulatory and management supervision of the diverse business units involved in a complex financial institution.

THE ROLE OF CAPITAL IN RISK MANAGEMENT AND BUSINESS STRATEGY

Given the diverse business units in a large, complex financial institution, lots of complexities come into play. First of all, some types of financial institutions have a very large institutional divide between the asset side of the organization and the liability side. Pension funds, life insurance companies, property and casualty insurance companies, and asset management companies are almost always clearly divided between the two halves of this balance sheet. This means that most managers either oversee assets or liabilities but almost never both. The authors would argue that this split is unfortunate from a “best practice” risk management point of view since many risk factors affect both sides of the balance sheet. Nonetheless this political reality simplifies managerial reporting.

In commercial banks, senior management is faced with three types of business units:

- Those who own only assets (specialized lending offices)
- Those who own only liabilities (the unit issuing wholesale certificates of deposit)
- Those who own both assets and liabilities (retail branches, business banking offices, trading floors doing repurchase agreements, etc.)

The largest number of business units falls into the last category, where the business units have both assets and liabilities. Management has two challenges:

1. Put the business units on a common basis
2. Adjust differences between units so that an apples-to-apples performance comparison can be done

Implicit in these two tasks are the assumptions that a common basis is possible and that backward-looking market data or backward-looking financial statements are a good indicator of future performance. The latter assumption is more tenuous than the first.¹⁰

There are two steps in accomplishing these two tasks:

1. Specify the method by which business units with different risks can be compared to each other (especially in the case of a comparison like Unit A and Unit C graphed previously).
2. Do the calculations that make this comparison possible.

The “common practice” in this regard is what we focus on in this section. In the last section of this chapter, we talk about the pitfalls of the common practice and spend the last chapters of this book refining these calculations using the tools that we develop in the rest of the book.

- Step 1: *Choose the basis for the calculations (market value basis or financial accounting basis).* As mentioned earlier in this chapter, most commercial banks use a mix of these approaches but choose a financial accounting basis for the vast majority of business units. We assume that choice is made here.
- Step 2: *Select the adjustments to nominal net interest income* for the business unit that must be made to allow for units of different risk to be compared
- Step 3: *Select a common basis for comparing business units* of different levels of net income to better understand relative risk-adjusted performance.

For a pension fund, this process is pretty easy and straightforward. The basis is the market value–based return on the asset portfolio being analyzed. There are no adjustments to this basis except to subtract the relative performance of the benchmark index by which performance is judged. A check is also done to ensure that the risk (usually tracking error) of the funds invested was within an acceptable tolerance around the risk of the benchmark.

For a bank, this process is much more convoluted. Some of the complexity is productive and some is done to ensure maximum employment in the performance measurement unit.

Here is a sample of the kinds of calculations usually made on a sample portfolio of external transactions like this:

Assets

3-year fixed rate auto loans	100

Liabilities

90-day certificates of deposit	120

All assets are funded with matched maturity borrowings from the transfer pricing center as discussed earlier in this chapter. Transfer prices are assigned using the techniques of the following chapters. All funds from the certificates of deposits gathered are sold to the transfer pricing center. Capital is then assigned depending on the nature of the assets and the nature of liabilities. Proceeds of the capital are sold back to the transfer pricing center with a maturity to match the period over which the amount of capital is held constant. The result is a financial accounting–based management information system that would look something like Exhibit 2.2.

The portfolio of three-year auto loans yields 6.76 percent. The exact-day count matched maturity transfer pricing system assigns \$100 in borrowings to fund the three-year auto loans on a matched maturity basis. These borrowings have a yield of 5.18 percent. We discuss how to calculate a transfer price for auto loans like these in Chapters 4 and 5. In a similar way, the 90-day funds raised by the business unit from issuing certificates of deposit are sold on a matched maturity basis to the transfer pricing center at 1.50 percent, a premium of 0.13 percent over the 1.37 percent cost of the certificates of the deposit to the business unit. This premium is a function of the marginal cost of issuing large-lot, 90-day certificates of deposit, not an arbitrary premium.

EXHIBIT 2.2 Risk-Adjusted Return on Capital

Assets	Balance Sheet Amount	Yield (Percent)	Annual Income
3-year fixed rate auto loans	100.0	6.76	6.76
90-day loans to the transfer pricing center	120.0	1.50	1.80
Capital-related loans to the transfer pricing center	7.7	1.50	0.12
Total Assets	227.7		8.68
Liabilities			
90-day certificates of deposit	120.0	1.37	1.64
3-year borrowings from the transfer pricing center	100.0	5.18	5.18
Total Liabilities	220.0		6.82
Capital			
Capital assigned for 90-day certificates of deposit	1.2		
Capital assigned for 3-year auto loans	6.5		
Total Capital Assigned	7.7		
Total Liabilities and Capital	227.7		1.86
Risk-adjusted return on auto loans			24.31%
Risk-adjusted return on deposits			13.00%
Total risk-adjusted return			24.05%

The next step is where there is a considerable lack of consensus among financial institutions for the reasons we outline in the next section. Most leading international banks assign “risk-adjusted capital” to each asset and (less often) each liability. The rationale for this assignment of capital varies widely from bank to bank, but they include the following common themes:

1. The total amount of capital assigned, when added across business units, should sum to the total amount of capital in the organization. The “total amount of capital in the organization” in some banks means financial accounting capital and in other banks it means the market value of the bank’s common stock.
2. The amount of capital assigned to each asset depends on the risk of the asset class. Van Deventer and Imai (2003) show how the Shimko, Tejima, and van Deventer (1993) model of risky debt can be used to do this while considering both the interest rate risk and the credit risk of the asset category.
3. The amount of capital assigned to each liability depends on its maturity and the “liquidity” protection the deposits provide. This is perhaps the most controversial step in the calculation.

4. Risk adjustments come in two places. Interest rate risk is removed from the business unit by the assignment of matched maturity transfer pricing assets and liabilities to each transaction. Credit risk is reflected in the risk-weighted capital assignment.
5. The result is a “risk-adjusted return on capital.” For the auto loan portion of the portfolio, the interest rate risk-adjusted net income is the \$6.76 in income on the auto loans less the \$5.18 cost of matched maturity funding on that portfolio, for an interest rate risk-adjusted net income of \$1.58. This is a 24.31 percent return on the \$6.50 in risk-adjusted capital assigned to the auto loan portfolio. Note that the interest rate adjustment comes in the numerator of the return calculation (since net income is net of interest rate risk) and the credit risk adjustment comes in the denominator via the risk-adjusted capital assigned.
6. In a similar way we can calculate a 13 percent risk-adjusted return on the deposit portfolio and a 24.05 percent risk-adjusted return for the business unit as a whole.

For a comprehensive survey of this risk-adjusted capital allocation system, interested readers should see Matten (1996).

Common practice in the banking industry says that, armed with these figures, we can do the following:

1. We can clearly judge which business unit creates the most shareholder value.
2. We can rank the business units from best to worst.
3. We can decide which business units should be expanded (receive more capital) and which business units should be discontinued.
4. We can carry this managerial process to the individual transaction level and set correct pricing for every asset and liability for a clear “accept/reject” criterion. More generally, this risk-adjusted capital allocation system correctly calculates the true value/market price of each asset and liability and the bank should buy assets that have an offered price (not yield) below this in the marketplace and issue liabilities that have a market price (not yield).

If this is true, we have a powerful tool on our hands. If this is more than a touch of overselling the concept, we need to be careful how we use it. We turn to those issues next.

CAPITAL-BASED RISK MANAGEMENT IN BANKING TODAY: PROS AND CONS

The previous section outlines the “common practice” for risk-adjusted asset and liability selection in the commercial banking business. We discussed the same process in the pension fund industry earlier in this chapter:

1. Select a benchmark index that best matches the nature of the assets being analyzed.
2. This benchmark should be a naïve strategy that any institution could achieve with a minimum of internal staff and analysis (such as buy every stock in the S&P 500 or buy every U.S. Treasury bond outstanding).

3. Measure the “tracking error” of the asset class being studied to insure that the asset class is within the limits for tracking error versus the benchmark.
4. A good manager who has assets whose total return in excess of the benchmark return is considered “plus alpha,” someone who has generated a positive risk adjusted return.

This approach is starkly different than the approach taken by the banking industry:

1. It is based on market returns, not financial accounting returns.
2. Risk adjustment is reflected in the returns of the benchmark and the limits on tracking error.
3. There is no “transfer pricing” of matched maturity funds.
4. There is no capital allocation.
5. There is no risk-adjusted return on capital calculation.
6. There is also no pretense that the methodology provides a tool for pricing the assets in question, only that skill of the manager selecting the assets can be measured.

While we discussed the concerns about this calculation above, it has a simplicity and elegance and lack of arbitrariness that distinguishes it as a managerial tool.

Why did the banking industry turn to such a convoluted calculation instead?

- The banking industry has many more people to manage.
- The banking industry has many classes of assets and liabilities that don’t have observable market prices.
- The banking industry feels it cannot manage its huge number of business units on a mark-to-market basis.
- The banking industry has many business units that have both assets and liabilities so the concept of “capital” is a natural one.
- In many business units, they originate more liabilities than assets, so there has to be some “adjustment” to a normal basis.
- Many of the regulatory restrictions on banks revolve around capital.

There are a few more important points about the common practice in banking risk-adjusted capital that need to be mentioned. Van Deventer and Imai (2003) in Chapters 9 and 10 go into these points at length:

- The risk-adjusted return on capital is backward looking, not forward looking. It has long been known that past returns on a security are fully reflected in its market price and therefore provide no guidance on the level of future returns.
- The risk-adjusted return on capital assigns capital in a well-intended way unsupported by a sophisticated theory of capital structure (although the approach van Deventer and Imai [2003] recommend is a step in this direction).
- No theory of asset valuation in financial economics yet incorporates the capital of the potential buyer in the valuation formula.

Are these concerns just academic concerns or do they indicate severe problems with the approach taken in the banking industry?

A few examples prove that the problems are severe:

1. The common practice in banking industry risk-adjusted return on capital indicates a negative risk-adjusted return on holding government securities because they have a negative spread versus the marginal cost of funds for the bank. This signal says no banks should own risk-free government securities, but in fact almost all banks do. Therefore, the decision rule is wrong. Banks seek to overcome this obvious “miss” in the rule with ad hoc adjustments
2. Anecdotal evidence from many banks that have carried this analysis to the transaction level has accumulated that the prices indicated are consistently “off the market,” which is no surprise, since the transactions are not benchmarked in current market prices
3. If this method were superior to the approach used by fund managers, then fund managers would use it.

Calculation of “shareholder value added,” which measures true cash flow rather than financial net income, is a partial step in the right direction. It falls short of the mark in applying a minimum hurdle rate for acceptable return, which is not consistent with actual returns in the market (in contrast to the approach that pension funds take). For more on the conventional implementation of shareholder value-added, see Chapter 2 of Uyemura and van Deventer (1993).

As a result of these concerns about risk-adjusted capital allocation, the authors recommend that the risk-adjusted capital allocation approach be restricted to the largest consolidated business units in the banking business (for example, wholesale corporate lending, investment banking, retail banking, and so on) and not be carried down to the smaller business unit level or the transaction level. They do not work on a microlevel and their errors at the macrolevel are obvious enough that senior management can bear the problems in mind when making decisions.

The methodology is not sufficiently accurate to provide a high quality guide to asset and liability selection. (We discuss advanced methodologies that provide better guidance to asset and liability management selection in Chapters 36 to 41.) We now turn to regulatory views of capital management.

HISTORY OF CAPITAL-BASED REGULATIONS IN COMMERCIAL BANKING

Capital allocation by banks and regulators has had a complex chicken-and-egg relationship that has obscured bank management’s ability to adopt the correct risk-adjusted strategies for asset selection. When regulatory constraints are imposed that are not based on market values and market returns, bank management teams are forced into regulatory “arbitrage” that stems from their conflicting dual obligation to do (1) what regulators insist they do and, at the same time, (2) seek to maximize risk-adjusted returns to shareholders over the long run.

As Uyemura and van Deventer (1993) outline, the regulatory process in the United States has historically had a tripartite structure:

1. System liquidity provided by the Board of Governors of the Federal Reserve System.
2. Interest rate–related regulations, the now defunct Regulation Q.
3. Deposit insurance managed by the Federal Deposit Insurance Corporation.

The role of the Federal Reserve in providing both institution-specific liquidity when necessary (for example in the wake of the September 11, 2001, attacks on New York City) and systemic liquidity during the 2006 to 2011 credit crisis is well known. We discuss the credit crisis liquidity support by the Federal Reserve for specific institutions in Chapter 37.

What is less well known is why the United States regulators maintained a limit on the maximum deposit interest rate that any commercial bank could pay, known as Regulation Q. For many years the maximum rate that could be paid on savings deposits with balances of \$100,000 or less was at levels of 5.00 to 5.25 percent. The theory was that this limit would prevent aggressive risk-seeking banks from bidding up deposit rates to expand their activities rapidly.

The emergence of *money market funds*, which pooled the funds of small depositors so they could earn the higher returns paid on deposits of more than \$100,000, spelled the end of Regulation Q in the United States, which required regulators to take a more sophisticated approach to risk management. Ironically, the low rates that have prevailed since the 2006–2011 credit crisis and the “breaking of the buck” (the \$1.00 net asset value of each dollar provided to a money fund) by the Reserve Fund (as listed by Bloomberg, September 16, 2008) in the wake of the bankruptcy of Lehman Brothers have impacted the growth of money market funds in a major way.

In 1980, the passage of the Financial Deregulation and Monetary Control Act ended Regulation Q and led financial regulators in the United States to turn to capital-based regulations to insure the safety and soundness of financial institutions. The regulators’ first attempt in this regard was to define the concept of *primary capital*. The regulatory definition of capital was firmly rooted in financial accounting-based capital, not market values. They included in primary capital common equity and retained earnings; perpetual preferred stock (an instrument, which until then had rarely been issued); the reserve for loan losses (even though this amount was usually set by commercial banks to be consistent with expected losses); and so-called “mandatory convertible notes,” a specific type of debt instrument that became the primary vehicle for arbitrage of the primary capital concept by commercial banks.

This regulation triggered a number of actions by bank management, which led regulators to conclude that the primary capital calculations were misguided and that the actions they engendered increased, rather than decreased, the risk in the banking system:

- Banks reduced their assets subject to the capital regulations by selling their safest, most liquid assets (government securities), increasing liquidity risk at many banks.
- Banks sold their head office buildings to book an accounting “gain” that increased capital ratios, but as a result added to the fixed costs that banks would have to incur on an ongoing basis.

- Banks began to engage in an increasing number of “off-balance sheet transactions,” obscuring the accuracy of financial statements on which regulations were based.

Bank failures dramatically increased throughout the 1980s so it is hard to argue that the concept of primary capital was effective in reducing the risk in the banking system. The primary capital concept was ultimately scrapped in favor of a system that was both more international in scope and more sophisticated.

In 1986, the Board of Governors of the Federal Reserve System and the Bank of England announced the concept of risk-based capital. The risk-based capital concept was specifically designed to address some of the major deficiencies of the primary capital concept:

1. Make a distinction among asset classes by riskiness
2. Incorporate off-balance sheet transactions to reduce regulatory capital arbitrage
3. Adopt common standards internationally for a fair international competition in financial services

The risk-based capital regulations had two basic components:

1. A revised definition of regulatory capital
2. An introduction of the “risk-weighted asset concept”

Risk-weighted assets are a variation of the bank risk-adjusted capital concept outlined in the previous section, except that the riskiness of an asset class is reflected in risk-weighted assets associated with the financial accounting assets for that class rather than the risk-adjusted capital assigned to the asset class. Low-risk assets were assigned risk-weighted assets that were a small fraction of their financial accounting amounts or market values. Higher risk assets were assigned a risk-weighted asset amount closer to or equal to their financial accounting values. This change in capital regulations was designed to encourage banks to retain liquid assets on their balance sheets by reducing the capital ratio penalty of the primary capital concept. These changes were ultimately labeled “Basel I” after the headquarter’s city of the Bank for International Settlements, which coordinated the international negotiations for the capital accords.

Bankers and regulators quickly noted that the regulations were too simple (and therefore inaccurate) in two respects. First, they almost completely ignored interest rate risk. Second, the risk-weighting schemes of Basel I were so few in number that finer gradations of credit risk were ignored, again penalizing asset classes with a risk level lower than average in their Basel I risk category.

The Board of Governors of the Federal Reserve in the United States moved to address the first point with a proposal that would link interest rate risk to capital levels. This proposal made use of Macaulay’s duration concept (1938) that we discussed in Chapter 1 and review in detail in later chapters. This proposal was never adopted as the mass of banks in the United States were incapable of understanding the basics of duration, or at least they claimed as much. Cynics responded that it was harder to get a driver’s license than to be a bank CEO, where no license is required that certifies a basic understanding of how to run a bank. The Fed, even today, is very

nervous about safety and soundness regulations that exceed a minimum level of complexity and that has affected U.S. participation in Basel II, to which we now turn.

The Basel Committee on Banking Supervision began its deliberations on how to address the shortcomings of Basel I in the late 1990s, and it has issued a steady stream of pronouncements and proposals beginning in 2001. In June 2004, the Basel II guidelines were announced that rely on a combination of capital requirements, supervisory review, and market discipline to reign in bank risk taking. Van Deventer and Imai (2003) point out the biggest weakness of the Basel II approach: Capital ratios, however derived, are weak predictors of the safety and soundness of financial institutions and they significantly underperform models like the recently announced Loss Distribution Model of the Federal Deposit Insurance Corporation.¹¹

In 2010 and 2011, Basel III provisions were announced by the Basel Committee on Banking Supervision in series of pronouncements that now total over 1,000 pages. The proposals add liquidity ratios to the Basel II approach, but with much less insight and usefulness than the best practice approaches we discuss in Chapter 37 on liquidity risk. The Basel III regulations add leverage restrictions to their inverse, the capital ratios of Basel II. In this edition of *Advanced Financial Risk Management*, we have been candid in our de-emphasis of discussions about legacy approaches to risk management that have failed the test of time. The credit crisis of 2006–2011 resulted in the failure or nationalization of the biggest banks in Europe and the United States. Neither the Basel II regulations nor the proposals for Basel III would have prevented these failures. Indeed, the implication of the Basel II and Basel III regulations, that government securities are the only truly safe asset, is now a source of amusement in the wake of ISDA credit events in the credit default swap market for Ecuador and Greece, with an artful dodge of an ISDA default for Ireland on a technicality. Government financial difficulties continue at this writing in the Eurozone.

We understand that the Basel Committee on Banking Supervision was charged with an impossible task: forging an international standard for banking risk regulation among a group of countries none of which had yet succeeded in defining such a standard in their own financial systems. Rather than review the lowest common denominator proposals that emerged, we focus in the remainder of this book on best practice and emerging best practice.¹²

We begin our journey through best practice risk management with the very foundation of financial institutions risk management, interest rate risk.

NOTES

1. Assuming, of course, that the securities firm providing the insurance is riskless.
2. Working paper, Cornell University and Kamakura Corporation, February 2001.
3. The authors would like to thank Leo de Bever, Barbara Zvan, Francois Gingras, and Barbara Chun for their insights and inspiration of this section of the book.
4. From the perspective of 2012, this emphasis on the marginal cost of funds as a basis for pricing is showing its age. We discuss problems with this concept in later chapters as part of the continual battle between those with a financial accounting basis (“How can we make money if we don’t charge at least as much of our cost of funds?”) and those with a market orientation. (“This credit is priced 10 basis points over market for AA-rated companies and the fact that we are a triple-B bank is irrelevant.”)

5. The multistate nature of commercial banking is commonplace in the United States and most other countries in 2012, but it was a significant barrier to implementation of comprehensive risk management in the mid-1980s, just as national barriers are today.
6. See Chapter 1 of Dennis Uyemura and Donald R. van Deventer, *Financial Risk Management in Banking* for more on this chart.
7. Donald R. van Deventer and Kenji Imai discuss this issue extensively in *Credit Risk Models and the Basel Accords* (2003) in Chapters 1 and 2.
8. See Chapter 9 in van Deventer and Imai, *Credit Risk Models and the Basel Accords* (2003).
9. A copy of the FDIC Loss Distribution Model is available on the research page on www.kamakuraco.com.
10. See Chapter 9 in van Deventer and Imai, *Credit Risk Models and the Basel Accords* (2003).
11. The December 10, 2003, press release of the FDIC is available on www.kamakuraco.com or the FDIC website, www.fdic.gov.
12. Readers interested in the evolution of the Basel Accords are encouraged to review the proposals in their original form on www.bis.org.

PART

Two

Risk Management Techniques for Interest Rate Analytics

Interest Rate Risk Introduction and Overview

Interest rate risk is at the heart of every silo of truly integrated risk management: credit risk, market risk, asset and liability management (ALM), liquidity risk, performance measurement, and even operational risk. Credit risk analysis that is not built on a random interest rate risk framework completely misses one of the key macroeconomic factors driving default. This is even truer when credit risk omits explicit modeling of a key macroeconomic factor such as home prices. Both interest rates and home prices, as predicted by Jarrow et al. (2003) in the FDIC Loss Distribution Model, are critical drivers of correlated bank defaults as the \$1 trillion bailout of the U.S. savings and loan industry in the 1980s and 1990s and the \$1 trillion bailout of the U.S. “too big to fail” institutions in 2008 and 2009 confirms. Market risk at most institutions includes a large proportion of fixed income instruments, so we have to deal with interest rate risk in a comprehensive way to deal with market risk. Asset and liability management at most institutions involves both sides of the balance sheet and a range of instruments with complexities that go far beyond those found in market risk, such as pension liabilities, insurance policies, nonmaturity deposits, credit card advances, and so on. Analysis of each of these instruments must be based on a random interest rate framework as well. In ALM, the complexity of the task is aggravated by the frequent need to do the analysis both on a mark-to-market basis and on a net income simulation basis looking forward. Liquidity risk analysis, which is focused on how cash flows respond to changes in macroeconomic factors and their impact on credit risk, market risk, and ALM risk, is a “derivative” risk in the sense that it stems from some other problem.

For all these reasons, we pay special attention to interest rate risk in this book. In this chapter, we focus on the big picture, in particular, on two key questions:

1. How have interest rates behaved in the past?
2. Is there an optimal interest rate risk position for a financial institution and, if so, what are the determinants of what that position should be?

We start by focusing on the first question because we want our analysis of interest rate risk to be grounded in reality. Too often, analysts of interest rate risk assume a yield curve smoothing method or a model of interest rate movements that is simply inconsistent with market prices of fixed income instruments and with

historical (and with high probability) and future movements in interest rates. We begin this chapter by summarizing key facts from the 50-year history of movements in the U.S. Treasury yield curve. We select this data set for a simple reason—the sheer size of the data and length of the time period for which data is available.

BACKGROUND INFORMATION ON MOVEMENTS IN THE U.S. TREASURY YIELD CURVE

Space does not permit a detailed, day-by-day analysis of the movements of the U.S. Treasury yield curve, but these references provide exactly such an analysis and we strongly encourage the serious reader to review them:

- Daniel T. Dickler, Robert A. Jarrow, and Donald R. van Deventer, “Inside the Kamakura Book of Yields: A Pictorial History of 50 Years of U.S. Treasury Forward Rates,” Kamakura Corporation memorandum, September 13, 2011.
- Daniel T. Dickler and Donald R. van Deventer, “Inside the Kamakura Book of Yields: An Analysis of 50 Years of Daily U.S. Treasury Forward Rates,” Kamakura Corporation blog, www.kamakuraco.com, September 14, 2011.
- Daniel T. Dickler, Robert A. Jarrow, and Donald R. van Deventer, “Inside the Kamakura Book of Yields: A Pictorial History of 50 Years of U.S. Treasury Zero Coupon Bond Yields,” Kamakura Corporation memorandum, September 26, 2011.
- Daniel T. Dickler and Donald R. van Deventer, “Inside the Kamakura Book of Yields: An Analysis of 50 Years of Daily U.S. Treasury Zero Coupon Bond Yields,” Kamakura Corporation blog, www.kamakuraco.com, September 26, 2011.
- Daniel T. Dickler, Robert A. Jarrow, and Donald R. van Deventer, “Inside the Kamakura Book of Yields: A Pictorial History of 50 Years of U.S. Treasury Par Coupon Bond Yields,” Kamakura Corporation memorandum, October 5, 2011.
- Daniel T. Dickler and Donald R. van Deventer, “Inside the Kamakura Book of Yields: An Analysis of 50 Years of Daily U.S. Par Coupon Bond Yields,” Kamakura Corporation blog, www.kamakuraco.com, October 6, 2011.

These publications are either available on the Kamakura Corporation website at www.kamakuraco.com or available by request from e-mail: info@kamakuraco.com. The Kamakura Corporation website also includes three videos that show movements in interest rates on a daily basis from January 2, 1962, to August 22, 2011. We encourage readers to also review the narrated videos shown in Exhibit 3.1.

The series of pictorial volumes and analysis by Dickler, Jarrow, and van Deventer make a number of facts clear about the historical movements of interest rates in the U.S. Treasury market, one of the largest fixed income markets in the world. A reader who is thinking, “But my market is different” is guilty of risk fallacy number 1 that we discussed in the Introduction: “If it hasn’t happened to me yet, it won’t happen to me, even if it’s happened to someone else.” This is the same mistake made by U.S. risk managers who thought that home prices wouldn’t fall like they did in Japan. It is the same mistake made by U.S. ALM experts who thought that the low

Interest Rate Risk Video: 50 Years of Par
Coupon Bond Yields, narrated by Dr.
Donald R. van Deventer, founder of
Kamakura Corporation

Interest Rate Risk Video: 50 Years of
Zero Coupon Bond Yields, narrated by Dr.
Donald R. van Deventer, founder of
Kamakura Corporation

Interest Rate Risk Video: 50 Years of
Forward Rate Movements, narrated by Dr.
Donald R. van Deventer, founder of
Kamakura Corporation

EXHIBIT 3.1 Video References

interest rates in the wake of the collapse of the Japanese bubble would never happen anywhere else. Ideally, this book would analyze interest rates in 20 countries, not one, but space does not permit that.

Here is a list of the key facts that the Dickler, Jarrow, and van Deventer studies identified about the U.S. Treasury market experience.

Fact 1: Historical data is typically available at maturities that vary over time, as seen in Exhibit 3.2 of the 12,395 data points of daily data on U.S. Treasury yields from January 2, 1962, to August 22, 2011.

Fact 2: It is very rare for the forward rate curve to be monotonically upward sloping or downward sloping. We discuss the derivation of forward rate curves in the following chapters. The chart in Exhibit 3.3 shows the distribution of the forward curves among three types of shapes: humped, upward sloping, or downward sloping.

In percentage terms, the forward rate curve has been humped for 93.32 percent of the observations since 1962 (see Exhibit 3.4).

Fact 3: It is very common for the forward rate curve to show a large number of local maximums and local minimums. Exhibit 3.5 summarizes the breakdown of the number of local minimums and maximums (“humps”) over the 12,395 business days in the data set.

In percentage terms, an interest rate modeler who wants to be able to replicate the shape of the forward rate curve correctly at least 90 percent of the time needs a model that can create at least six local maximums or minimums (see Exhibit 3.6).

Fact 4: During some extended periods of time, the forward rate curve may have no or few humps. As shown in Exhibit 3.7, 2011 was such a year.

Fact 5: It is equally true that extended periods of extreme humps can prevail, such as the forward rate curves for 2009 in the aftermath of Federal Reserve efforts to speed recovery from the 2008 collapse of firms like Bear Stearns, Lehman Brothers, FHLMC, FNMA, AIG, Wachovia, and Washington Mutual (see Exhibit 3.8).

Dickler, Jarrow, and van Deventer find that the most frequent local minimums or maximums are at months 5, 10, 22, 34, 58, 79, 113, 144, 202.

EXHIBIT 3.2 Analysis of Federal Reserve H15 Release Data Regimes

Start Date	End Date	Data Regime	Number of Observations
1/2/1962	6/30/1969	1, 3, 5, and 10 years	1,870
7/1/1969	5/28/1976	1, 3, 5, 7, and 10 years	1,722
6/1/1976	2/14/1977	1, 2, 3, 5, 7, and 10 years	178
2/15/1977	12/31/1981	1, 2, 3, 5, 7, 10, and 30 years	1,213
1/4/1982	9/30/1993	3 and 6 months with 1, 2, 3, 5, 7, 10, and 30 years	2,935
10/1/1993	7/30/2001	3 and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	1,960
7/31/2001	2/15/2002	1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	135
2/19/2002	2/8/2006	1, 3, and 6 months with 1, 2, 3, 5, 7, 10, and 20 years	994
2/9/2006	8/22/2011	Second era: 1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	1,388
		Total	12,395

Source: Kamakura Corporation.

EXHIBIT 3.3 Kamakura Corporation, U.S. Treasury Monthly Forward Rates, January 2, 1962, to August 22, 2011, Number of Forward Rate Curve Shapes by Data Regime

Start Date	End Date	Data Regime	Number of Observations	Forward Curve Shape		
				Humped	Monotonically Downward Sloping	Monotonically Upward Sloping
1/2/1962	6/30/1969	1, 3, 5, and 10 years	1,870	1,503	40	327
7/1/1969	5/28/1976	1, 3, 5, 7, and 10 years	1,722	1,707	11	4
6/1/1976	2/14/1977	1, 2, 3, 5, 7, and 10 years	178	146		32
2/15/1977	12/31/1981	1, 2, 3, 5, 7, 10, and 30 years	1,213	1,163	15	35
1/4/1982	9/30/1993	3 and 6 months with 1, 2, 3, 5, 7, 10, and 30 years	2,935	2,930		5
10/1/1993	7/30/2001	3 and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	1,960	1,960		
7/31/2001	2/15/2002	1, 3 and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	135	135		
2/19/2002	2/8/2006	1, 3, and 6 months with 1, 2, 3, 5, 7, 10, and 20 years	994	669		325
2/9/2006	8/22/2011	Second era: 1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	1,388	1,354		34
		Total	12,395	11,567	66	762

Source: Kamakura Corporation.

EXHIBIT 3.4 Kamakura Corporation, U.S. Treasury Monthly Forward Rates, January 2, 1962, to August 22, 2011, Number of Forward Rate Curve Shapes by Data Regime

Start Date	End Date	Data Regime	Number of Observations	Probability of Forward Curve Shape		
				Humped	Monotonically Downward Sloping	Monotonically Upward Sloping
1/2/1962	6/30/1969	1, 3, 5, and 10 years	1,870	80.37%	2.14%	17.49%
7/1/1969	5/28/1976	1, 3, 5, 7, and 10 years	1,722	99.13%	0.64%	0.23%
6/1/1976	2/14/1977	1, 2, 3, 5, 7, and 10 years	178	82.02%	0.00%	17.98%
2/15/1977	12/31/1981	1, 2, 3, 5, 7, 10, and 30 years	1,213	95.88%	1.24%	2.89%
1/4/1982	9/30/1993	3 and 6 months with 1, 2, 3, 5, 7, 10, and 30 years	2,935	99.83%	0.00%	0.17%
10/1/1993	7/30/2001	3 and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	1,960	100.00%	0.00%	0.00%
7/31/2001	2/15/2002	1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	135	100.00%	0.00%	0.00%
2/19/2002	2/8/2006	1, 3, and 6 months with 1, 2, 3, 5, 7, 10, and 20 years	994	67.30%	0.00%	32.70%
2/9/2006	8/22/2011	Second era: 1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	1,388	97.55%	0.00%	2.45%
		Total	12,395	93.32%	0.53%	6.15%

Source: Kamakura Corporation.

EXHIBIT 3.5 Kamakura Corporation, U.S. Treasury Monthly Forward Rates, January 2, 1962, to August 22, 2011, Number of Local Maximums and Minimums by Data Regime

Data Regime	Number of Local Maximums and Minimums										Total
	0	1	2	3	4	5	6	7	8	9	
1, 3, 5, and 10 years	367	579	924								1,870
1, 3, 5, 7, and 10 years	15	475	572	660							1,722
1, 2, 3, 5, 7, and 10 years	32	29	25	55	37						178
1, 2, 3, 5, 7, 10, and 30 years	50	77	291	202	440	153					1,213
3 and 6 months with 1, 2, 3, 5, 7, 10, and 30 years	5	25	173	343	1,042	671	427	249			2,935
3 and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years		25	10	237	96	788	217	580	7		1,960
1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years				1	48	80	5	1			135
1, 3, and 6 months with 1, 2, 3, 5, 7, 10, and 20 years	325	59	206	96	120	7	116	65			994
Second era: 1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	34	167	110	500	88	177	44	197	61	10	1,388
Total	828	1,436	2,312	2,141	1,903	1,801	805	1,091	68	10	12,395

Data Regime	Probability of Local Maximums and Minimums										Total
	0	1	2	3	4	5	6	7	8	9	
1, 3, 5, and 10 years	19.6%	31.0%	49.4%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
1, 3, 5, 7, and 10 years	0.9%	27.6%	33.2%	38.3%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
1, 2, 3, 5, 7, and 10 years	18.0%	16.3%	14.0%	30.9%	20.8%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
1, 2, 3, 5, 7, 10, and 30 years	4.1%	6.3%	24.0%	16.7%	36.3%	12.6%	0.0%	0.0%	0.0%	0.0%	100.0%
3 and 6 months with 1, 2, 3, 5, 7, 10, and 30 years	0.2%	0.9%	5.9%	11.7%	35.5%	22.9%	14.5%	8.5%	0.0%	0.0%	100.0%
3 and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	0.0%	1.3%	0.5%	12.1%	4.9%	40.2%	11.1%	29.6%	0.4%	0.0%	100.0%
1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	0.0%	0.0%	0.7%	35.6%	59.3%	3.7%	0.7%	0.0%	0.0%	0.0%	100.0%
1, 3, and 6 months with 1, 2, 3, 5, 7, 10, and 20 years	32.7%	5.9%	20.7%	9.7%	12.1%	0.7%	11.7%	6.5%	0.0%	0.0%	100.0%
Second era: 1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	2.4%	12.0%	7.9%	36.0%	6.3%	12.8%	3.2%	14.2%	4.4%	0.7%	100.0%
Total	6.7%	11.6%	18.7%	17.3%	15.4%	14.5%	6.5%	8.8%	0.5%	0.1%	100.0%

Source: Kamakura Corporation.

EXHIBIT 3.6 Cumulative Probability of Local Maximums and Minimums

Data Regime	Cumulative Probability of Local Maximums and Minimums									
	0	1	2	3	4	5	6	7	8	9
1, 3, 5, and 10 years	19.6%	50.6%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
1, 3, 5, 7, and 10 years	0.9%	28.5%	61.7%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
1, 2, 3, 5, 7, and 10 years	18.0%	34.3%	48.3%	79.2%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
1, 2, 3, 5, 7, 10, and 30 years	4.1%	10.5%	34.5%	51.1%	87.4%	100.0%	100.0%	100.0%	100.0%	100.0%
3, and 6 months with 1, 2, 3, 5, 7, 10, and 30 years	0.2%	1.0%	6.9%	18.6%	54.1%	77.0%	91.5%	100.0%	100.0%	100.0%
3 and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	0.0%	1.3%	1.8%	13.9%	18.8%	59.0%	70.1%	99.6%	100.0%	100.0%
1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	0.0%	0.0%	0.7%	36.3%	95.6%	99.3%	100.0%	100.0%	100.0%	100.0%
1, 3, and 6 months with 1, 2, 3, 5, 7, 10, and 20 years	32.7%	38.6%	59.4%	69.0%	81.1%	81.8%	93.5%	100.0%	100.0%	100.0%
Second era: 1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	2.4%	14.5%	22.4%	58.4%	64.8%	77.5%	80.7%	94.9%	99.3%	100.0%
Total	6.7%	18.3%	36.9%	54.2%	69.5%	84.1%	90.6%	99.4%	99.9%	100.0%

Source: Kamakura Corporation.




EXHIBIT 3.7 Monthly Forward Rates for U.S. Treasury Bonds, Maximum Smoothness Forward Rate Smoothing, January 1 to August 22, 2011

Source: Kamakura Corporation.




EXHIBIT 3.8 Monthly Forward Rates for U.S. Treasury Bonds, Maximum Smoothness Forward Rate Smoothing, January 1 to December 31, 2009

These optima are close to knot points (data at which the Federal Reserve provides data) at 6, 12, 24, 36, 60, 84, 120, and 240 months.

Fact 6: The volatility of the daily changes in forward rates can be calculated both using all data and using all data except the first day after the Federal Reserve changes the “data regime,” the maturities at which U.S. Treasury yields are made available. The volatilities of daily changes do not have a simple upward sloping, downward sloping, or constant pattern. In fact, an assumption that volatilities have a simple shape is dramatically wrong, as shown in Exhibit 3.9.

Fact 7: Dickler, Jarrow, and van Deventer find that the smoothing techniques employed in the figures above and explained in later chapters produce very




EXHIBIT 3.9 U.S. Treasury Monthly Forward Rates Using Maximum Smoothness Forward Rates Standard Deviation of Daily Changes, with and without First Data of Data Regime Change, January 2, 1962, to August 22, 2011

smooth forward rate curves. Adams and van Deventer (1994) adopt a common definition of smoothness used in engineering and computer graphics, the integral of the squared second derivative of the forward rate over the length of the forward rate curve. We can approximate this by taking the sum of the squared second differences of the (up to) 360 monthly forward rates on each of the 12,395 business days in the data set.

$$Z[t] = \sum_{i=2}^{360} \left[(f(i) - f(i-1)) - (f(i-1) - f(i-2)) \right]^2$$

The lowest values of $Z[t]$ are those days with the smoothest forward rate curves. The highest values of $Z[t]$ are those days on which the forward rate curve is least smooth. Not surprisingly, Dickler, Jarrow, and van Deventer find that the least smooth forward rate curve comes about on September 17, 2008, two days after the formal filing of bankruptcy by Lehman Brothers. The Federal Reserve's efforts to provide liquidity via the Treasury market and direct lending result in dramatic variation of the forward rate curve. This volatility was inevitable because of the short-term rates reported by the Federal Reserve on its H15 statistical release for September 17, 2008: 3 months = 0.03 percent, 6 months = 1.03 percent, and 1 year = 1.50 percent. Exhibit 3.10 shows how the Fed's attempts to supply liquidity caused dramatic movements in the forward rate curve. Note that, due to limitations in the graphics routine used, forward rates are plotted for every third month.




EXHIBIT 3.10 Monthly Forward Rates for U.S. Treasury Bonds, Maximum Smoothness Forward Rate Smoothing, September 1 to December 31, 2008

Source: Kamakura Corporation.

Fact 8: Nominal interest rates can indeed turn negative. This is not obvious in the Federal Reserve’s H15 data for a simple reason: the Federal Reserve has acknowledged that it will not publish negative interest rates because the U.S. Treasury “will not issue” bonds with a negative coupon. Academics have argued that nominal interest rates cannot be negative “because it is costless to store cash.” In fact, it is *not* costless to store cash and the proof of that is the August 5, 2011, announcement by the Bank of New York Mellon that it would charge a fee to clients who supplied zero interest demand deposits to the bank with balances of more than \$50 million. The *Wall Street Journal* report on the fee noted that 1-month U.S. Treasury bill rates were trading at negative yields on that day. Negative nominal rates have been reported on the website of the Hong Kong Monetary Authority and observed in Switzerland and Japan as well.

With this data as background, we now turn to the mechanics of how a financial institution can and should manage interest rate risk.

A STEP-BY-STEP APPROACH TO ANALYZING INTEREST RATE RISK

Besides reviewing historical experience, our goal in this chapter is to address the big picture of determining the optimal degree of interest rate risk for a major financial institution. We start with a very fundamental question: *What number is interest rate risk management all about?*

In Chapter 1, we noted how Merton and Jarrow have both proposed that a properly described “put option” is the single number that best describes the total risk of a financial institution in most situations. If we are thinking only about the firm’s current portfolio, this put option is a European put option exercised only on one

specific date on the value of the portfolio we hold now. If we are thinking about a multiperiod analysis of the firm's balance sheet going forward, this put option can be a American put option exercisable at many dates in the future.

We can describe almost all common risk management risk measures using the Merton and Jarrow analysis of the put option. When one thinks about it, the put option is a much more precise summary of the risk than other terms that are more commonly used to describe risk.

For example

- Instead of the *10-day VaR of a trading portfolio*, what is the value of a 10-day put on the firm's current portfolio with an exercise price equal to the portfolio's current market value? The price of the put will increase sharply with the risk of the firm's portfolio, and the put's price will reflect all possible losses and their probability, not just the ninety-ninth percentile loss as is traditional in value at risk analysis.
- Instead of *stress testing the 12-month net income* of the financial institution to see if net income will go below \$100 million for the year, what is the price of a put option in month 12 that will produce a gain in net income exactly equal to the shortfall of net income versus the \$100 million target? The more interest rate risk in the balance sheet of the financial institution, the more expensive this put will be. The put will reflect all levels of net income shortfall and their probability, not just the shortfalls detected by specific stress tests.
- Instead of the *Basel II or Basel III risk-weighted capital ratio* for the bank, what is the price of the put option that insures solvency of the bank in one year's time? This put option measures all potential losses embedded in the financial institution's balance sheet and their probability of occurrence, including both interest rate risk and credit risk, as we discuss at the end of this chapter.
- Instead of *expected losses on a collateralized debt obligation tranche's B tranche*, what is the price of a put option on the value of the tranche at par at maturity? This put option reflects all losses on the tranche, not just the average loss, along with their probability of occurrence.
- Instead of *expected losses on the Bank Insurance Fund* in the United States, the Federal Deposit Insurance Corporation has valued the put option of retail bank deposits at their par value as discussed in the FDIC's loss distribution model announced on December 10, 2003.

The market values of put options on bank total risk are visible in the marketplace. Exhibit 3.11 shows put option prices at the close of trading on April 20, 2011, for JPMorgan Chase. The closing stock price for JPMorgan Chase was \$42.72.

With this key Merton-Jarrow insight in mind, a lot of the debate about the proper focus of interest rate risk management fades away. Three different dimensions of a financial institution's risk have dominated past arguments about how much interest rate risk is the right amount:

- *Net interest income (or net income)*, a measure that by definition is multiperiod in nature and includes instruments that the financial institution owns today and those that it will own in the future. Net interest income as the focus of risk analysis is unheard of on the trading floor but is much discussed by senior

Strike	Symbol	Last	Chg	Bid	Ask	Vol	Open Int
16.00	JPM120421P00016000	0.02	0.00	N/A	0.01	32	550
17.00	JPM120421P00017000	0.11	0.00	N/A	0.01	12	12
19.00	JPM120421P00019000	0.04	0.00	N/A	0.02	16	150
20.00	JPM120421P00020000	0.04	0.00	N/A	0.02	1	44
21.00	JPM120421P00021000	0.01	0.00	N/A	0.01	43	228
22.00	JPM120421P00022000	0.01	0.00	N/A	0.01	39	480
23.00	JPM120421P00023000	0.12	0.00	N/A	0.01	10	192
24.00	JPM120421P00024000	0.01	0.00	N/A	0.01	9	492
25.00	JPM120421P00025000	0.01	0.00	N/A	0.01	300	1,056
26.00	JPM120421P00026000	0.02	0.00	N/A	0.01	50	824
27.00	JPM120421P00027000	0.02	0.00	N/A	0.01	4	1,003
28.00	JPM120421P00028000	0.02	0.00	N/A	0.01	10	586
29.00	JPM120421P00029000	0.02	0.00	N/A	0.01	200	1,565
30.00	JPM120421P00030000	0.01	0.00	N/A	0.01	50	4,350
31.00	JPM120421P00031000	0.01	0.00	N/A	0.01	30	5,168
32.00	JPM120421P00032000	0.01	0.00	N/A	0.01	10	8,501
33.00	JPM120421P00033000	0.01	0.00	N/A	0.01	30	7,815
34.00	JPM120421P00034000	0.01	0.00	N/A	0.01	2	2,952
35.00	JPM120421P00035000	0.01	0.00	N/A	0.01	15	7,076
36.00	JPM120421P00036000	0.01	↓0.01	N/A	0.01	1	6,777
37.00	JPM120421P00037000	0.02	0.00	N/A	0.01	5	7,908
38.00	JPM120421P00038000	0.01	0.00	N/A	0.01	5	6,988
39.00	JPM120421P00039000	0.01	0.00	N/A	0.01	19	5,668
40.00	JPM120421P00040000	0.01	0.00	N/A	0.01	41	14,092
41.00	JPM120421P00041000	0.01	0.00	N/A	0.01	36	11,037
42.00	JPM120421P00042000	0.01	↓0.02	N/A	0.01	227	14,403
43.00	JPM120421P00043000	0.28	↑0.13	0.27	0.28	16,389	16,449
44.00	JPM120421P00044000	1.28	↑0.43	1.24	1.28	5,099	9,690
45.00	JPM120421P00045000	2.27	↑0.54	2.23	2.28	1,734	7,459
46.00	JPM120421P00046000	3.10	↑0.32	3.20	3.30	78	4,673
47.00	JPM120421P00047000	4.04	↑0.01	4.20	4.30	36	1,312
48.00	JPM120421P00048000	4.50	↓0.28	5.20	5.30	68	417
49.00	JPM120421P00049000	6.15	↑0.10	6.20	6.30	1	104
50.00	JPM120421P00050000	6.95	↑0.05	7.20	7.30	100	251
55.00	JPM120421P00055000	12.05	0.00	12.15	12.30	5	410
60.00	JPM120421P00060000	16.10	0.00	15.30	17.35	10	10

Highlighted options are in-the-money.

EXHIBIT 3.11 Put Options on JPMorgan Chase Common Stock

Source: www.yahoo.com, April 20, 2011.

management of many financial institutions. Many of the required outputs of the Federal Reserve's "Comprehensive Capital Analysis and Review 2012" stress tests are surprisingly expressed in net income terms.

- *Market value of portfolio equity*, which is "bank-speak" for the market value of the assets a financial institution owns today, less the market value of its liabilities. This is the common trading floor choice, aided by the luxury of having market prices for almost every asset and liability in full view. It is also increasingly popular among sophisticated total balance sheet risk managers, even when the balance sheet is dominated by assets and liabilities with no observable market price.
- *Market-based equity ratio*, which is the ratio of the mark-to-market value of the equity of the portfolio ("market value of portfolio equity" in bank-speak) divided by the market value of assets. This is most closely related to the capital ratio formulas of the primary capital era, Basel I, Basel II, Basel III, and Solvency II.
- *Default probability of the institution*, which is another strong candidate as a single measure of risk. The complexity here is over which time period the default probability should be for.

While these measures can move in conflicting directions under specific scenarios, the put option proposed by Merton and Jarrow behaves in a consistent way, which is why Jarrow labels the put option a “coherent risk measure” as noted in Chapter 1.

THE INTEREST RATE RISK SAFETY ZONE

The complex interaction between interest rate risk and credit risk shouldn't be allowed to obscure a key conclusion about interest rate risk. Assuming away the potential default of the assets that a financial institution may hold:

No matter which maturity of investment we make with the capital of the financial institution, this firm will never go bankrupt from interest rate risk as long as the remaining assets are financed on a matched maturity basis. This firm is in the “safety zone” and shareholders are indifferent to the choice of maturity on the investment of these equity funds. The reason they are indifferent is because, even at the retail investor level, the investor can form a model portfolio with the investor's desired risk. For example, what if management invested the equity in 10 year bonds and the retail investor wanted the funds invested in a one year maturity. The investor achieves that by buying x% of the firm's common stock, selling short x% of the 10-year bonds, and buying x% of the one-year bonds.¹

Once the firm begins to mismatch on the loan funding, the probability of bankruptcy creeps up from zero. For any mismatch, there is an interest rate scenario (often an incredibly unlikely but theoretically possible scenario), which can cause bankruptcy. Once this bankruptcy probability becomes significant from a practical point of view, shareholders begin to bear the expected cost of bankruptcy, which is the unhedgable loss of the perpetual ability of the financial institution to generate profitable assets that earn an excess above the return on traded securities of equivalent risk available in the market. It is this expected bankruptcy loss that causes the firm's stock price to decline when the firm moves out of the safety zone. As long as the firm is within the safety zone, the firm's stock price will always be indifferent to the level of interest rate risk,² all other things being equal, because the interest rate risk is hedgable even by retail investors for their own account.³

As we plunge into the interest rate analytics first and then credit risk analytics second, we should not lose sight of this simple truth about interest rate risk.

NOTES

1. More practically, the investor would short government bond futures.
2. This doesn't mean that the firm's stock price won't decline if Firm A has taken more risk than Firm B. It means only that Firm A and Firm B, all other things being equal, will have the same initial stock prices if both differ only in the degree of interest rate risk taken in the safety zone.
3. Lot sizes for financial futures are so small that the argument that this hedging can be done at the retail level is a practical one. At the institutional investor level, of course, this hedging could be much easier.

Fixed Income Mathematics

The Basic Tools

In the first three chapters, we discussed measures of risk and return in a general way. We have discussed the Jarrow-Merton put option as a measure of risk along with more traditional measures of risk, such as the sensitivity of net income to changes in risk factors and the sensitivity of the net market value of a portfolio (i.e., the value of equity in the portfolio) to changes in risk factors.

In the rest of this book, we discuss implementation of these concepts in a detailed and practical way. Implementation requires accurate and efficient modeling of market values of every transaction in the portfolio both at the current time and at any date in the future. Valuation and multiperiod simulation also require exact knowledge of cash flow dates and amounts, recognizing that these amounts may be random, such as an interest rate on a floating-rate mortgage, early prepayment on a callable bond, or payment on a first to default swap. We turn to that task now.

MODERN IMPLICATIONS OF PRESENT VALUE

The concept of present value is at the very heart of finance, and yet it can seem like the most mysterious and difficult element of risk management analytics even eight decades after the introduction of Macaulay's duration in 1938. It is safe to say, though, that no self-respecting finance person in a large financial institution should look forward to a pleasant stay in the finance area if he or she is uncomfortable with the present value concept and the basics of traditional fixed income mathematics. At the same time, the basic principles of present value have so many applications that a good understanding of them would be very beneficial to a wide variety of finance professionals. In this chapter, we present an overview of present value and fixed income mathematics. We touch on a wide variety of topics and leave a few to be covered in more detail in later chapters. Yield curve smoothing is covered in Chapter 5, and modern simulation of random interest rates is covered in Chapters 6 to 11. Finally, we cover "legacy" approaches to interest rate risk management and contrast them to more modern approaches in Chapters 12 and 14. We integrate this pure interest rate risk-focused section with credit modeling in Chapters 15 to 17.

The present value concept and related issues like yield-to-maturity and forward interest rates provide the building blocks for these more complex issues.

These concepts will be old hat to many readers, but they have great implications for the measurement of risk and return we've discussed in Chapters 1 to 3.

PRICE, ACCRUED INTEREST, AND VALUE

The accounting and economics professions have engaged in many wars during their history. Occasionally, there have been periods of peaceful coexistence and general agreement on what's important in life. That seems to be happening now with the implementation of the market value-based Financial Accounting Standards (FAS) 133, FAS 157, and similar International Financial Reporting Standards (IFRS) pronouncements.

There was one important battle lost by the economics profession that still causes finance professionals pain, however. That battle was fought over three concepts: *price*, *accrued interest*, and *value*.

In this chapter, we focus consistently on the concept of value—what a security is truly worth in the marketplace. In a rational world, a security's value and the risk-adjusted present value of future cash flows had better well be close to the same thing or the reader is wasting a lot of valuable time reading this book when he could be arbitraging the market. For the purposes of this book, when we say *value*, we really mean *present value*; that is, what a rational person would pay today for the risk-adjusted cash flows to be received in the future.

Unfortunately, this simple concept has been complicated by the noble idea that one “earns” interest on a bond even if you buy it just after the most recent coupon payment was paid and sell it before the next coupon payment is paid. This idea isn't harmful in and of itself, but, in the form the idea has been implemented in many markets, nothing could be farther from economic reality. The person who receives the interest on a bond is the person who is the owner of record on the record date that determines who receives the interest. For accounting (not economic) purposes, the accounting profession has decided that the value of a bond has to be split into two pieces: the “price,” which is intended by the accountants to be relatively stable, and “accrued interest,” which is an arbitrary calculation that determines who “earned” interest on the bond even if they received no cash from the issuer of the bond.

The basic accounting rule for splitting value isn't harmful on the surface:

$$\text{Value} = \text{Price plus accrued interest}$$

What causes the harm is the formula for accrued interest. In a calculation left over from the “BC (before calculators) era,” accrued interest is calculated as a proportion of the next coupon on a bond. While there are many variations on the calculation of this proportion, the simplest one divides the actual number of days between the settlement date and the last bond coupon payment date by the total number of days between coupon payments.

CALCULATION OF ACCRUED INTEREST

ABC Bank sells a 10-year, 10 percent coupon bond for an amount of money equal to 102 percent of par value, or \$1,020.00 per bond. The next semiannual coupon

of \$50.00 will be paid in 122 days, and it has been 60 days since the last coupon payment.

$$\begin{aligned}
 \text{Value} &= \$1,020.00 \\
 \text{Accrued interest} &= 60 (\$50) / [60 + 122] = \$16.48 \\
 \text{Price} &= \text{Value} - \text{Accrued interest} = \$1,020.00 - 16.48 \\
 &= \$1,003.52
 \end{aligned}$$

or 100.352 percent of par value.

There is one fundamental flaw in these simple rules. In economic terms, the amount of accrued interest is an arbitrary calculation that has no economic meaning.¹ Why? The amount of accrued interest bears no relationship to the current level of interest rates. The accrued interest on the bond above is calculated to be the same in an environment when interest rates are 3 percent as one where interest rates are 30 percent. The amount of accrued interest depends solely on the coupon on the bond, which reflects interest rate levels at the time of issue, not interest rates today. Since price is calculated as value minus an economically meaningless number, then price is an economically meaningless number as well. Unfortunately, the number referred to most often on a day-to-day basis in the financial industry is “price,” so those of us who focus on good numbers have to work backward to remain focused on what’s important: the value of the transaction, \$1,020.00.² Market slang labels true present value as the *dirty price*, meaning price plus accrued interest. *Clean price* means true present value minus accrued interest; that is, the price quoted in the market. In our minds, the market has applied the words “dirty” and “clean” to the wrong definitions of price.

From now on, when we say “value” we mean true present value, dirty price, and price plus accrued interest. To avoid confusion, we will always specify the concept of price intended between “clean price” and “dirty price.”

PRESENT VALUE

A dollar received in the future is almost always worth less than a dollar received today.³ This is more than just a reflection of inflation; it’s a recognition that market requires someone using the resources of others to pay “rent” until those resources are returned. The implications of this simple fact are very important but straightforward.

THE BASIC PRESENT VALUE CALCULATION

If the value of one dollar received at time t_i is written $P(t_i)$, then the present value of an investment that generates n varying cash flows $C(t_i)$ at n times t_i in the future (for i equal to 1 through n) can be written as follows:

$$\text{Present value} = \sum_{i=1}^n P(t_i) C(t_i)$$

Example

Cash Flow 1:	\$100.00
Date Received:	1 year from now
Cash Flow 2:	\$200.00
Date Received:	10 years from now
Value of \$1 dollar received in the future	
Received in 1 year:	\$0.90
Received in 10 years:	\$0.50
Value	$= 100(.9) + 200(.5) = 190.00$

This simple formula is completely general. It is the heart of most banking calculations and much of finance. Note that the present value formula has the following features:

- It is independent of the number of cash flows.
- It is independent of the method of interest compounding.
- It assumes that cash flows are known with certainty.

The $P(t_i)$ values, the value of one dollar received at time t_i , are called *discount factors*. They provide the basis for all yield curve calculations: bond valuation, forward rates, yield-to-maturity, forward bond prices, and so on. How they are determined and how they link with other concepts is the topic that makes up the heart of the rest of the book, and, in particular, we focus on them in Chapters 5 and 17. For the time being, we assume that these discount factors are known. That allows us to write down immediately a number of other formulas for valuation of securities where the contractual cash flows are known with certainty.

CALCULATING THE VALUE OF A FIXED COUPON BOND WITH PRINCIPAL PAID AT MATURITY

If the actual dollar amount of the coupon payment on a fixed coupon bond is C and principal is repaid at time t_n , the value (price plus accrued interest) of the bond is:

$$\text{Value of fixed coupon bond} = C \sum_{i=1}^n P(t_i) + P(t_n)\text{Principal}$$

Note that this formula applies regardless of how often coupon payments are made and regardless of whether the first coupon period is a full period in length. The calculation is nearly identical to the example given above.

CALCULATING THE COUPON OF A FIXED COUPON BOND WITH PRINCIPAL PAID AT MATURITY WHEN THE VALUE IS KNOWN

If the value (price plus accrued interest) of a fixed coupon bond is known and the discount factors are known, the dollar coupon payment that leads a bond to have such a value is calculated by rearranging the previous formulas:

$$\text{Dollar coupon amount of fixed coupon bond} = \frac{\text{Value} - P(t_n)\text{Principal}}{\sum_{i=1}^n P(t_i)}$$

Example

Principal amount:	\$1,000.00
Interest paid:	Semiannually
Value:	\$1,150.00
Periods to maturity:	4 semiannual periods
Days to next coupon:	40 days
Value of \$1 dollar received in the future	
Received in 40 days:	\$0.99
Received in 6 months plus 40 days:	\$0.94
Received in 1 year plus 40 days:	\$0.89
Received in 1.5 years plus 40 days:	\$0.83

$$\begin{aligned}\text{Coupon amount} &= (1,150 - .83 \times 1,000) / (.99 + .94 + .89 + .83) \\ &= 87.67\end{aligned}$$

$$\begin{aligned}\text{Coupon rate} &= [2 \text{ payments} \times \text{Amount}] / 1,000.00 \\ &= 17.53\%\end{aligned}$$

THE VALUE OF AN AMORTIZING LOAN

The value of an amortizing loan is the same calculation as that for a fixed coupon bond except that the variable C represents the constant level payment received in each period. There is no explicit principal amount at the end since all principal is retired in period-by-period payments included in the amount C . When the periodic payment dollar amount is known, the value of the loan is

$$\text{Value of amortizing bond} = C \left[\sum_{i=1}^n P(t_i) \right]$$

Note again that this formula holds even for a short first period and for any frequency of payments, be they daily, weekly, monthly, quarterly, semiannually, annually, or on odd calendar dates.

CALCULATING THE PAYMENT AMOUNT OF AN AMORTIZING BOND WHEN THE VALUE IS KNOWN

Like the fixed coupon bond case, the payment amount on an amortizing loan can be calculated using the known amount for value (principal plus accrued interest) and rearranging the equation above for amortizing loans.

$$C = \text{Payment amount on amortizing bond} = \frac{\text{Value}}{\sum_{i=1}^n P(t_i)}$$

Risk Management Implications

In Chapter 3, we stated that the interest rate risk of a financial institution is simply the sum of the present value of its assets less the present value of its liabilities. We also noted that the width of the interest rate risk safety zone is at least consistent with the capital of the organization being invested in any maturity from one day to infinitely long, provided all other assets are match funded. We now turn to the use of basic present value concepts for valuation of floating-rate securities.

CALCULATING THE VALUE OF A FLOATING-RATE BOND OR LOAN WITH PRINCIPAL PAID AT MATURITY

The calculation of the value of a floating-rate bond is more complicated than the valuation of a fixed rate bond because future coupon payments are unknown. In order to value floating-rate bonds, we divide the formula for determining the coupon on the bond into two pieces: an index plus a fixed dollar spread over the index. We assume the index is set at a level such that the value of a bond with one period (equal in length to $1/m$ years) to maturity and an interest rate equal to the index would be par value. This is equivalent to assuming that the yield curve (and its movements), which determine future values of the index, is the same as the yield curve we should use for valuing the floating-rate security. For example, we assume that the right yield curve for valuing all securities whose rate is reset based on the London interbank offered rate should be valued at the LIBOR yield curve. We relax this assumption below and in later chapters when we have more powerful tools to analyze more realistic assumptions about the value of a bond whose rate floats at the index plus zero spread. Using our simple assumptions for the time being, the present value of a floating-rate bond priced at the index level plus a fixed dollar amount (the spread) per period is equal to par value (the value of all future payments at an interest rate equal to the floating index plus the return of principal) plus the present value of the stream of payments equal to the spread received each period. The value of the stream of payments equal to the spread can be valued using the equation above for the value of an amortizing bond. Therefore, the value of a floating-rate bond is

$$\begin{aligned} \text{Value of floating-rate bond} = & P(t_1) \left[\left(1 + \frac{\text{Index}}{m} \right) \text{Principal} \right] \\ & + \text{Spread} \left[\sum_{i=1}^n P(t_i) \right] \end{aligned}$$

Note that the first term will be equal to the principal value, given our definition of the index, if the time to the first payment t_1 is equal to the normal length of the

period between coupons. The first term will generally not be equal to the principal value if the length of the first period is shorter than its normal length, such as 17 days when the interval between interest rate resets is six months. Note also that the index may not be the same as the interest rate (which we will call the *formula rate* from here on) used as the formula for setting the coupon on the bond. See the examples for valuation of bonds where the index rate and the formula rate are different.

Example

In this example, the coupon formula is based on LIBOR, which is a yield curve that was intended to be consistent with the credit risk of major international banks. Lately, this consistency has eroded, but we ignore that fact in this chapter. In this example, our counterparty has a higher level of credit risk than a major international bank because the index value that would cause a floating-rate note to have a present value equal to par value is LIBOR plus 1 percent, so the difference between the coupon rate on the floating-rate note and “par pricing” is 2.00 percent – 1.00 percent.

Coupon formula:	6-month LIBOR (adjusted) + 2.00 percent ⁴
Index value:	6-month LIBOR + 1.00 percent
Principal amount:	2,000.00
Spread in dollar terms over and above par pricing:	$(.02 - .01) \times 2,000/2 = 10.00$
Maturity:	2 years
Time to next coupon	Exactly 6 months
Value of \$1 received in the future	
Received in 6 months:	0.96
Received in 1 year:	0.92
Received in 1.5 years:	0.87
Received in 2.0 years:	0.82

$$\begin{aligned} \text{Value} &= .96 \left[\left(1 + \frac{\text{Libor} + .01}{2} \right) \text{Principal} \right] + 10.00[0.96 + 0.92 + 0.87 + 0.83] \\ &= 2,035.8 \end{aligned}$$

Risk Management Implications

We now have enough tools to analyze the interest rate risk embedded in a simple financial institution. Armed with these simple tools and the right set of present value factors, we can:

1. Mark the assets and liabilities to market
2. Calculate the implied value of the equity of the financial institution
3. Change the discount factors to stress test the calculations in 1 and 2 with respect to interest rate risk

We will illustrate how to do this in more detail in subsequent chapters. We now turn to another set of basic tools that will help us do this for a broad range of assets and liabilities.

COMPOUND INTEREST CONVENTIONS AND FORMULAS

No one who has been in the financial world for long expects interest rates to remain constant. Still, when financial analysts perform traditional yield-to-maturity calculations, the constancy of interest rates is the standard assumption simply because of the difficulty of making any other choice. Particularly when discussing compound interest, the assumption that interest rates are constant is a very common assumption. This section discusses those conventions as preparation for the yield-to-maturity discussion later in this chapter. Using the techniques in Chapter 5, we can relax these simple but inaccurate assumptions.

Future Value of an Invested Amount Earning at a Simple Interest Rate of y Compounded m Times per Year for n Periods

Almost all discussions of compound interest and future value depend on four factors:

- The constancy of interest rates
- The nominal annual rate of interest
- The number of times per year interest is paid
- The number of periods for which interest is compounded

$$\text{Future value of invested amount} = (\text{Invested amount}) \left(1 + \frac{y}{m}\right)^n$$

Future Value of an Invested Amount Earning at a Simple Interest Rate of y Compounded Continuously for n Years

What happens if the compounding period shrinks smaller and smaller, so that interest is compounded every instant, not every second or every day or every month? If one assumes a simple interest rate of y invested for n years, the corresponding formula to the equation above is

$$\text{Future amount} = \text{Invested amount}(e^{yn})$$

Many financial institutions actually use a continuously compounded interest rate formula for consumer deposits. The continuous compounding assumption is a very convenient mathematical shorthand for compound interest because derivatives of the constant $e = 2.7128 \dots$ to a power are much simpler (once one gets used to them) than the derivatives of a discrete compounding formula. The continuous time assumption also allows the use of the powerful mathematical tools called *stochastic processes*. (We use these in Chapter 6 through 11 and discuss explicitly in Chapter 13.)

Example

Invested amount: 100.00
 Continuously compounded interest rate: 12.00%
 Investment period: 2 years and 180 days

$$\begin{aligned}\text{Future value} &= 100.00e^{.12 \times \left(2 + \frac{180}{365}\right)} \\ &= 134.88\end{aligned}$$

Present Value of a Future Amount If Funds Are Invested at a Simple Interest Rate of y Compounded m Times per Year for n Periods

A parallel question to the compound interest investment question is this: If one seeks to have 100 dollars n periods in the future and funds had been invested at a constant rate y compounded m periods per year, how much needs to be invested initially? The answer is a rearrangement of our original formula for compounding of interest:

$$\text{Invested amount} = \frac{\text{Future value of invested amount}}{\left(1 + \frac{y}{m}\right)^n}$$

Present Value of a Future Amount If Funds Are Invested at a Simple Interest Rate of y Compounded Continuously for n Years

When interest is assumed to be paid continuously, the investment amount can be calculated by rearranging the equation for continuous compounding of interest:

$$\begin{aligned}\text{Invested amount} &= \frac{\text{Future amount}}{e^{yn}} \\ &= \text{Future amount}(e^{-yn})\end{aligned}$$

COMPOUNDING FORMULAS AND PRESENT VALUE FACTORS $P(t)$

In these sections, we focus on the calculation of present value factors on the assumption that interest rates are constant but compound at various frequencies. We could use this approach to calculate the present value factors that we used in all of the examples in the first part of this chapter, but, as President Richard M. Nixon famously said, “That would be wrong.” Why? Because the yield curve is almost never absolutely flat. We deal with this issue in Chapter 5. With these compounding formulas behind us, the yield-to-maturity calculations can be discussed using a combination of the formulas in the second and third section of this chapter. The yield-to-maturity formula no longer plays an important role in modern interest rate risk management, but it is an important piece of market slang and is a concept that is essential to communicating the “present value” of a wide array of financial instruments. Besides that usage, the yield-to-maturity concept and the related duration and

convexity formulas are legacy concepts that are now outmoded except as a means to communicate interest rate risk concepts in a conventional, albeit, dated way.

YIELDS AND YIELD-TO-MATURITY CALCULATIONS

The concept of yield to maturity is one of the most useful and one of the most deceptive calculations in finance. Take the example of a five-year bond with a current value equal to its par value and with coupons paid semiannually. The market is forecasting 10 different semiannual forward interest rates (which we discuss in the next section) that have an impact on the valuation of this bond. There are an infinite number of forward rate combinations that are consistent with the bond having a current value equal to its par value. What is the probability that the market is implicitly forecasting all 10 semiannual forward rates to be exactly equal? Except by extraordinary coincidence, the probability is a number only the slightest bit greater than zero. Nonetheless, this is the assumption implied when one asks “What’s the yield to maturity on ABC Company’s bond?” It’s an assumption that is no worse than any other—but no better either, in the absence of any other information.

What if ABC Company has two bonds outstanding—one of 10 years in maturity and another of five years of maturity? When an analyst calculates the yield to maturity on both bonds, the result is the implicit assumption that rates are constant at one level when analyzing the first bond and constant at a different level when analyzing the second bond. The implications of this are discussed in detail in Chapter 5. In later chapters, we discuss a number of ways to derive more useful information from a collection of outstanding value information on bonds of a comparable type. As an example, if ABC Company has bonds outstanding at all 10 semiannual maturities, it is both possible and very easy to calculate the implied forward rates, zero coupon bond prices (discount factors $P[t]$), and pricing on all possible combinations of forward bond issues by ABC Company. The calculation of yield to maturity is more complex than these calculations, and it yields less information. Nonetheless, it’s still the most often-quoted bond statistic other than clean price (present value less accrued interest).

THE FORMULA FOR YIELD TO MATURITY

For a settlement date (for either a secondary market bond purchase or a new bond issue) that falls exactly one period from the next coupon date, the formula for yield to maturity can be written as follows for a bond that matures in n periods and pays a fixed coupon amount C m times per year:

$$\text{Value} = C \left[\sum_{i=1}^n P(t_i) \right] + P(t_n)[\text{Principal}]$$

$$\text{with } P(t_i) = \frac{1}{\left(1 + \frac{y}{m}\right)^i}$$

Yield to maturity is the value of y that makes this equation true. This relationship is the present value equation for a fixed coupon bond, with the addition that the discount factors are all calculated by dividing one dollar by the future value of one dollar invested at a constant rate y for i periods with interest compounded m times per year. In terms of the present value concept, y is the internal rate of return on this bond. As recently as three decades ago, dedicated calculators were necessary to estimate y . Now it's easy to solve for y using the “solver” function or other functions in common spreadsheet software.

YIELD TO MATURITY FOR LONG OR SHORT FIRST COUPON PAYMENT PERIODS

For most outstanding bonds or securities, there is almost always a short (or occasionally a long) interest payment period remaining until the first coupon payment is received. The length of the short first period is equal to the number of days from settlement to the next coupon, divided by the total number of days from the last coupon (or hypothetical last coupon, if the bond is a new issue with a short coupon date) to the next coupon. The method for counting days is specified precisely by the interest accrual method appropriate for that bond. We call this ratio of remaining days to days between coupons the fraction x .⁵ Then the yield to maturity can be written as

$$\text{Value} = C \left[\sum_{i=1}^n P(t_i) \right] + P(t_n) [\text{Principal}]$$

$$\text{with } P(t_i) = \frac{1}{\left(1 + \frac{y}{m}\right)^{i-1+x}}$$

where, as above, y is the value that makes this equation true.

CALCULATING FORWARD INTEREST RATES AND BOND PRICES

What is the six-month interest rate that the market expects to prevail two years in the future?⁶ Straightforward questions like this are the key to funding strategy or investment strategy at many financial institutions. Once discount factors (or, equivalently, zero coupon bond prices) are known, making these calculations is simple.

IMPLIED FORWARD INTEREST RATES ON ZERO-COUPON BONDS

The forward simple interest rate t_i years in the future on a security that pays interest and principal at time t_{i+1} is

$$\text{Forward interest rate} = \frac{100}{t_{i+1} - t_i} \left[\frac{P(t_i)}{P(t_{i+1})} - 1 \right]$$

This forward rate is the simple interest rate consistent with $1/(t_{i+1} - t_i)$ interest payments per year. For example, if t_i is two years and t_{i+1} is 2.5 years, then the forward interest rate is expressed on the basis of semiannual interest payments.

Example

Value of \$1 dollar received in the future

Received in 350 days: \$0.88

Received in 1 year plus 350 days: \$0.80

$$\begin{aligned}\text{Forward rate} &= \frac{100}{1} \left[\frac{.88}{.80} - 1 \right] \\ &= 10.00\end{aligned}$$

This is the forward rate the market expects to prevail 350 days in the future on a one-year bond. The forward rate is expressed on the basis of annual compounding of interest.

IMPLIED FORWARD ZERO-COUPON BOND PRICES

A parallel question to the question of the level of forward interest rates is this query: What is the forward price of a six-month, zero-coupon bond two years in the future? The answer is a simple ratio. From the perspective of time 0 (the current time), the forward price at time t_i of a zero coupon bond with maturity at time t_{i+1} is

$$\text{Implied forward zero-coupon bond price} = \frac{P(t_{i+1})}{P(t_i)}$$

PRESENT VALUE OF FORWARD FIXED COUPON BOND

What is the present value of a bond to be issued on known terms some time in the future? The answer is a straightforward application of the basic present value formula. If the actual dollar amount of the coupon payment on a fixed coupon bond is C and principal is repaid at time t_n , the value (price plus accrued interest) of the bond is

$$\text{Present value of forward fixed coupon bond} = C \left[\sum_{i=1}^n P(t_i) \right] + P(t_n) [\text{Principal}]$$

This is exactly the same present value formula used earlier in the chapter.

IMPLIED FORWARD PRICE ON A FIXED COUPON BOND

There is another logical question to ask about a bond to be issued in the future on known terms. What would be the forward price, as of today, of the “when issued bond,” if we know its offering date? The answer is again a straightforward application of the basic present value equation. If the actual dollar amount of the coupon payment on a fixed coupon bond is C and principal is repaid at time t_n , the forward value (price plus accrued interest) of the bond at the issuance date t_0 is

$$\text{Forward value of fixed coupon bond} = \frac{C \left[\sum_{i=1}^n P(t_i) \right] + P(t_n) [\text{Principal}]}{P(t_0)}$$

IMPLIED FORWARD COUPON ON A FIXED COUPON BOND

Finally, a treasurer may often be asked the implied forward issue costs (i.e., forward coupon rates) for a bond to be issued at par (or any other value). This question can be answered using a minor modification of another basic present value equation. For a bond to be issued at time t_0 with n interest payments and principal repaid at time t_n , the dollar amount of the coupon is given by the formula

$$\begin{aligned} &\text{Dollar coupon amount of fixed coupon bond} \\ &= \frac{P(t_0)(\text{Value at issue}) - P(t_n)\text{Principal}}{\sum_{i=1}^n P(t_i)} \end{aligned}$$

OTHER FORWARD CALCULATIONS

The same types of calculations can be performed for a very wide variety of other instruments and statistics. For instance, the forward yield-to-maturity can be calculated on a fixed coupon bond. Forward amortizing bond prices and payment amounts can be calculated as well. The present values and forward prices for literally any form of cash flow can be analyzed using the formulas presented above.

SUMMARY

This chapter has summarized the building blocks of all financial market calculations: the concepts of present value, compound interest, yield to maturity, and forward bond yields and prices. All of these concepts can be expressed using algebra and nothing else. There are no mysteries here, yet the power of these simple concepts is great. Risk managers who make full use of present value and its implications truly do have a leg up on the competition. Those who don't bear a heavy burden in a very competitive world. As simple as these concepts are compared to the fixed income

derivatives calculations we tackle in later chapters, they are still widely misunderstood and their powerful implications are ignored by many financial institutions managers.

At this point, there are probably more than a few readers thinking “yes, these concepts aren’t so difficult as long as someone gives you the discount factors—but where am I supposed to get them?” We give the answers to that question in the next chapter and (on a credit-adjusted basis) in Chapter 17. There are so many potential ways to calculate the “correct” discount factors that we purposely have chosen to speak about them generally in this chapter so that no one technique would obscure the importance of the present value calculations.

NOTES

1. One exception to this comment is the differential tax impact of accrued interest and price.
2. There are exceptions to this concept of accrued interest. Australia and South Korea are two markets where there is no artificial division of value into accrued interest and price.
3. There have been a number of transactions in Japan at negative nominal interest rates in the late 1990s and early part of this decade, and the credit crisis of 2006–2011 has generated numerous examples of negative nominal and real interest rates. We return to this topic in later chapters.
4. Whenever LIBOR (the London interbank offered rate) is mentioned in the text that follows, we assume the rate has been converted from the market convention of actual days/360 days per year to actual days/365 days per year by multiplying nominal LIBOR by 365/360.
5. We reemphasize that the calculation of remaining days to days between coupons will vary depending on which of the common interest accrual methods are used to count the days.
6. In later chapters, we will discuss risk-neutral interest rates and their relationship to the true expected value of future interest rates.

Yield Curve Smoothing

In this chapter, we have two objectives. The first objective is to supply realistic zero-coupon bond prices, like those in Chapter 4, to a sophisticated user for analysis. The second objective for this chapter is to illustrate how to build a yield curve that meets the criteria for “best” that are defined by the user. Given a specific set of criteria, what yield curve results? By the criteria defined, how do we measure best and compare the best yield curve we derive with alternatives? There is a most important set of comparisons we need to make—among n different sets of criteria for best yield curve, which set of criteria is in fact best, and what does “best” mean in this context? The vast majority of literature on yield curve smoothing omits the definition of “best” and the related derivation of what yield curve-smoothing technique is consistent with the selected definition of best.

In this chapter, we report on the strengths and merits of this list of potential yield curve-smoothing techniques:¹

- Example A Yield step function
- Example B Linear yield function
- Example C Linear forward rate function
- Example D Quadratic yield function
- Example E Quadratic forward rate function
- Example F Cubic yield function
- Example G Cubic forward rate function
- Example H Quartic forward rate function

In this chapter, space permits us to focus only on selected techniques, examples A, D, F, and H. These examples include a yield step function, quadratic yield splines, cubic yield splines, and quartic forward rate (maximum smoothness) splines. (Readers interested in a more detailed treatment of these topics are referred to the original blogs listed in the notes to this chapter.)

For each of the techniques, we compare them with the well-known Nelson-Siegel (1987) yield curve approximation technique to show the stark contrasts between the results given the definition of “best” and what is produced by Nelson and Siegel.

In this section, we discuss how the definition of “best” yield curve or forward rate curve and the constraints one imposes on the resulting yield curve implies the mathematical function that is best. This is the right way to approach smoothing. The wrong way to approach smoothing is the exact opposite: Choose a mathematical function from the infinite number of functions one could draw and argue

qualitatively why your choice is the right one. We discuss the definition of best yield curve and the constraints commonly placed on the smoothing process.

The list of potential yield curve-smoothing methodologies given previously is not arbitrary. Instead, the mathematical yield curve or forward rate smoothing formulation can be derived as best after going through the following steps to define what “best” means and what constraints should be imposed upon the yield curve function (s) in deriving the best curve subject to those constraints:

Step 1: Should the smoothed curves fit the observable data exactly?

1a. Yes

1b. No

Except in the case where the data is known from the outset to be bad, there is no reason for yes not to be the standard choice here. If one does choose yes, the Nelson-Siegel approach to yield curve smoothing will be eliminated as a candidate because it is not flexible enough to fit five or more data points perfectly except by accident. We discuss the problems of the commonly used but flawed Nelson-Siegel approach at the end of this chapter.

Step 2: Select the element of the yield curve and related curves for analysis

2a. Zero coupon yields

2b. Forward rates

2c. Continuous credit spreads

2d. Forward continuous credit spreads

Smoothing choices 2a and 2b, smoothing with respect to either yields or forward rates, should only be selected when smoothing a curve with no credit risk. If the issuer of the securities being analyzed has a nonzero default probability, smoothing should only be done with respect to credit spreads (the risky zero-coupon yield minus the risk-free, zero-coupon yield of the same maturity) or forward credit spreads (the risky forward curve less the risk-free forward curve). (See Chapter 17 on why this distinction is important.)

Step 3: Define “best curve” in explicit mathematical terms

3a. Maximum smoothness

3b. Minimum length of curve

3c. Hybrid approach

The real objective of yield curve smoothing, broadly defined, is to estimate those points on the curve that are not visible in the marketplace with the maximum amount of realism. “Realism” in the original Adams and van Deventer (1994) paper on maximum smoothness forward rates (summarized later in this chapter) was defined as “maximum smoothness” for two reasons. First, it was a lack of smoothness in the forward rate curves produced by other techniques that was the criterion for their rejection by many analysts. Second, maximum smoothness was viewed as consistent with the feeling that economic conditions rarely change in a nonsmooth way and, even when they do (i.e., if the Central Bank manages short-term rates in such a way that they jump in 0.125 percent increments or multiples of that amount), the uncertainty about the *timing* of the jump makes future rate expectations smooth. The maximum

smoothness yield curve or forward rate curve $g(s)$ (where s is the years to maturity for the continuous yield or forward) is the one that minimizes this function Z :

$$Z = \int_{t_1}^{t_2} [g''(s)]^2 ds$$

There is another criterion that can be used for best curve, the curve that is the shortest in length. This again can be argued to be most realistic in the sense that one would not expect big swings in the yield, even if the swings are smooth. The length s of a yield curve or forward rate curve described by function $f(x)$ between maturities a and b is

$$s = \int_a^b \sqrt{1 + [f'(x)]^2} dx$$

where $f'(x)$ is the first derivative of the yield curve or forward rate curve. Some analysts, such as Andersen (2007), have suggested using tension splines that take a hybrid approach to these two criteria for best. Once the mathematical criterion for best is selected, we move on to specifications on curve fitting that will add to the realism of the fitted curves.

Step 4: Is the curve constrained to be continuous?

4a. Yes

4b. No

Why is this question being asked? It can be shown that the smoothest way to connect n points on the yield curve is *not* to fit a single function, like the Nelson-Siegel function following, to the data.

This was proven in an appendix contributed by Oldrich Vasicek in the Adams and van Deventer paper, corrected in van Deventer and Imai Financial Risk Analytics, 1996, and reproduced in the appendix that follows this chapter. The nature of the functional form that is consistent with the criterion for best and the constraints imposed can be derived using the calculus of variations, as Vasicek does in his proof. Any other functional form will not be best by the method selected by the analyst.

Given this, we now need to decide whether or not we should constrain the yield curve or forward rate curve to “connect” in a continuous fashion where each of the line segments meet. It is not obvious that yes is the only answer to this question. As an example, Jarrow and Yildirim (2003) use a four-step piecewise constant function for forward rates to model Treasury Inflation Protected Securities. Such a forward rate curve would not be continuous, but it is both perfectly smooth (except where the four segments join) and shorter in length than any nonconstant function. Most analysts, however, would answer yes in this step.

Step 5: Is the curve differentiable?

5a. Yes

5b. No

If one answers yes to this question, the first derivatives of the two line segments that join at “knot point” n have to be equal at that point in time. If one answers “no,” that allows the first derivatives to be different and the resulting curve will have kinks or elbows at each knot point, like you would get with linear yield curve smoothing—where each segment is linear and continuous over the full length of the curve but not differentiable at the points where the linear segments join. Most analysts would answer yes to this question.

Step 6: Is the curve twice differentiable?

6a. Yes

6b. No

This constraint means that the second derivatives of the line segments that join at maturity n have to be equal. If the answer to this question is yes, the segments will be neither linear nor quadratic.

Step 7: Is the curve thrice differentiable?

7a. Yes

7b. No

One can still further constrain the curve so that, even for the third derivative, the line segments paste together smoothly.

Step 8: At the spot date, time 0, is the curve constrained?

8a. Yes, the first derivative of the curve is set to zero or a nonzero value x .

8b. Yes, the second derivative of the curve is set to zero or a nonzero value y .

8c. No

This constraint is typically imposed in order to allow an explicit solution for the parameters of the curve for each line segment. See Appendix A as an example of how this assumption is used. The choice of nonzero values in 8a or 8b was suggested by Janosi and Jarrow (2002).

Step 9: At the longest maturity for which the curve is derived, time T , is the curve constrained?

9a. Yes, the first derivative of the curve is set to zero or a nonzero value j at time T .

9b. Yes, the second derivative of the curve is set to zero or a nonzero value k at time T .

9c. No

Step 9 is taken for two reasons. First, it is often necessary for an explicit, unique set of parameters to be derived for each segment as in step 8. Second and more importantly, it may be important for realistic curve fitting. If the longest maturity is 100 years, for instance, most analysts would expect the yield curve to be flat (choice 9a = 0), rather than rising or falling rapidly, at the 100 year point. If the longest maturity is only two years, this constraint would not be appropriate and one would probably make choice 9b. Again, the suggestion of nonzero values in 9a and 9b is useful in this case.

Once all of these choices have been made, both the functional form of the line segments and the parameters that are consistent with the data can be explicitly derived. The resulting forward rate curve or yield curve that is produced by this method has these attributes:

- Given the constraints imposed on the curve and the raw data, the curve is the best that can be drawn consistent with the analyst's definition of best.
- The data will be fit perfectly.
- All constraints will be adhered to exactly.

EXAMPLE A: STEPWISE CONSTANT YIELDS AND FORWARDS VS. NELSON-SIEGEL

In this section, we choose a common definition of best yield curve or forward rate curve and a simple set of constraints. We derive from our definition of best and the related constraints the fact that the best yield curve in this case is stepwise constant yields and forward rates. We then show that even this simplest of specifications is more accurate and “better” by our definition than the popular but flawed Nelson-Siegel approach.

For each of our smoothing techniques to be comparable, we will use the same set of data throughout the series. We will limit the number of data points to six points and smooth the five intervals between them. In order to best demonstrate how the definition of best affects the quality of the answer, we choose a simple set of sample data with two “humps” in the yield data shown in Exhibit 5.1.

Choosing a pattern of yields that falls on a perfectly straight line, for example, would lead a naïve analyst to the logical but naïve conclusion that a straight line is always the most accurate and best way to fit a yield curve. Such a conclusion would be nonsense. Similarly, restricting oneself to a yield curve functional form that allows only one “hump” would also lead one to simplistic conclusions about what is best.

We recently reviewed 4,435 days of “on the run” U.S. Treasury yields reported by the Federal Reserve in its H15 statistical release at maturities (in years) of 0.25, 0.5, 1, 2, 3, 5, 7, 10, and 30 years. In comparing same-day yield differences between maturities (i.e., the difference between yields at 0.25 and 0.5 years, 0.5 and one year, etc.), there were 813 days where there were at least two decreases in yields as maturities lengthened. For example, using the smoothing embedded in common spreadsheet software, Exhibit 5.2 is the yield curve for December 18, 1989.

EXHIBIT 5.1 Actual Input Data

Maturity in Years	Continuously Compounded	
	Zero-Coupon Yield	Zero-Coupon Bond Price
0.000	4.000%	1.000000
0.250	4.750%	0.988195
1.000	4.500%	0.955997
3.000	5.500%	0.847894
5.000	5.250%	0.769126
10.000	6.500%	0.522046




EXHIBIT 5.2 U.S. Treasury Yield Data, December 18, 1989

Source: Federal Reserve H15 Statistical Release.

Therefore, it is very important that our sample data have a similarly realistic complexity to it or we will avoid challenging the various yield curve-smoothing techniques sufficiently to see the difference in accuracy among them. It is important to note that the yields in these pictures are the simple yields to maturity on instruments that predominately are coupon-bearing instruments. These yields should *never* be smoothed directly as part of a risk management process. The graphs previously mentioned are included only as an introduction to how complex yield curve shapes can be. When we want to do accurate smoothing, we are very careful to do it in the following way:

1. We work with raw bond net present values (price plus accrued interest), not yields, because the standard yield-to-maturity calculation is riddled with inaccurate and inconsistent assumptions about interest rates over the life of the underlying bond.
2. We only apply smoothing to prices in the risk-free bond and money markets. If the underlying issuer of fixed income securities has a nonzero default probability, we apply our smoothing to credit spreads or forward credit spreads in such a way that bond mispricing is minimized. We never smooth yields or prices of a risky issuer directly, because this ignores the underlying risk-free yield curve.

We will illustrate these points later in this chapter and in Chapter 17. We now start with Example A, the simplest approach to yield curve smoothing.

From this point on in this chapter, unless otherwise noted, *yields* are always meant to be continuously compounded zero-coupon bond yields and *forwards* are the continuous forward rates that are consistent with the yield curve. The first step in exploring a yield curve-smoothing technique is to define our criterion for best and to specify what

constraints we impose on the best technique to fit our desired trade-off between simplicity and realism. We answer the nine questions posed earlier in this chapter.

Step 1: Should the smoothed curves fit the observable data exactly?

- 1a. Yes. With only six data points at six different maturities, it would be a poor exercise in smoothing if we could not fit this data exactly. We note later that the flawed Nelson-Siegel function is unable to fit this data and fails our first test.

Step 2: Select the element of the yield curve and related curves for analysis

- 2a. Zero-coupon yields is our choice. We find in the end that 2a and 2b are equivalent given our other answers to the following questions. If we were dealing with a credit-risky issuer of securities, we would have chosen 2c or 2d, but we have assumed our sample data is free of credit risk.

Step 3: Define “best curve” in explicit mathematical terms

- 3b. Minimum length of curve. This is the easiest definition of “best” to start with. We’ll try it and show its implications. We now move on to specifications on curve fitting that represent our desired trade-off between realism and ease of calculation.

Step 4: Is the curve constrained to be continuous?

- 4b. No. By choosing no, we are allowing discontinuities in the yield curve.

Step 5: Is the curve differentiable?

- 5b. No. Since the answer we have chosen above, 4b, does not require the curve to be continuous, it will not be differentiable at every point along its length.

Step 6: Is the curve twice differentiable?

- 6b. No. For the same reason, the curve will not be twice differentiable at some points on the full length of the curve.

Step 7: Is the curve thrice differentiable?

- 7b. No. Again, the reason is due to our choice of 4b.

Step 8: At the spot date, time 0, is the curve constrained?

- 8c. No. For simplicity, we answer no to this question. We relax this assumption in later posts in this series.

Step 9: At the longest maturity for which the curve is derived, time T , is the curve constrained?

- 9c. No. Again, we choose no for simplicity and relax this assumption later in the blog.

Now that all of these choices have been made, both the functional form of the line segments and the parameters that are consistent with the data can be explicitly derived from our sample data.

DERIVING THE FORM OF THE YIELD CURVE IMPLIED BY EXAMPLE A

The key question in the list of nine questions in the previous section is question 4. By our choice of 4b, we allow the “pieces” of the yield curve to be discontinuous. By virtue of our choices in questions 5 to 9, these yield curve pieces are also not subject

to any constraints. All we have to do to get the best yield curve is to apply our criterion for best—the curve that produces the yield curve with shortest length.

The length of a straight line between two points on the yield curve has a length that is known, thanks to Pythagoras, who rarely gets the credit he deserves in the finance literature. The length of a straight segment of the yield curve, which goes from maturity t_1 and yield y_1 to maturity t_2 and yield y_2 , is the square root of the square of $[t_2 - t_1]$ plus the square of $[y_2 - y_1]$. As we noted above, the general formula for the length of any segment of the yield curve between maturity a and maturity b is given by this formula, which is the function one derives as the difference between t_2 and t_1 becomes infinitely small:

$$s = \int_a^b \sqrt{1 + [f'(x)]^2} dx$$

where $f'(x)$ is the first derivative of the line segment at each maturity point, say x . If the line segment happens to be straight, the segment can be described as

$$y = f'(t) = mt + k$$

and the first derivative $f'(t)$ is, of course, m . How can we make this line segment as short as possible? By making $f'(t) = m$ as small as possible in absolute value, that is $m = 0$. Very simply, we have *derived* the fact that the yield curve segments that are best (have the shortest length) are flat line segments. We are allowed to use flat line segments to fit the yield curve because our answer 4b does not require the segments to join each other in continuous fashion. The functional form of the best yield curve given our definition and constraints can be derived more elegantly using the calculus of variations as Vasicek did in the proof of the maximum smoothness forward rate approach in Adams and van Deventer (1994), reproduced in the appendix at the end of this chapter.

Given that we have six data points, there are five intervals between points. We have taken advantage of our answer in 4b to treat the sixth interval from 0 years to maturity to 0 years to maturity as a separate segment. Given our original data, we know by inspection that our sixth segment has as a value the given value 4.000 percent for a maturity of zero.

In the rest of this book, we repeatedly use these four relationships between zero-coupon bond yields y , zero-coupon bond prices p , and continuous forward rates f . The relationship between forward rates and zero-coupon bond yields is given by the fourth relationship:

$$\begin{aligned} y(t) &= \frac{-1}{t} \ln[p(t)] \\ p(t) &= \exp\left[-\int_0^t f(s) ds\right] \\ p(t) &= \exp[-y(t)t] \\ f(t) &= y(t) + ty'(t) \end{aligned}$$

Since the first derivative of the yield curve in each case is zero in this Example A, the forward rates are identical to the zero yields.

How did we do in terms of minimizing the length of the yield curve over its 10-year span? We know the length of a flat line segment from t_1 to t_2 is just $t_2 - t_1$, so the total length of our discontinuous yield curve is

$$\begin{aligned} \text{Length} &= (0 - 0) + (0.25 - 0) + (0.5 - 0.25) + (1 - 0.5) + (3 - 1) \\ &\quad + (5 - 3) + (7 - 5) + (10 - 7) \\ &= 10.000 \end{aligned}$$

We now compare our results for our best Example A yield curve and constraints to the popular but flawed Nelson-Siegel approach.

FITTING THE NELSON-SIEGEL APPROACH TO SAMPLE DATA

We now want to compare Example A, the model we derived from our definition of best and related constraints, to the Nelson-Siegel approach. The table in Exhibit 5.3 emphasizes the stark differences between even the basic Example A model and Nelson-Siegel.

As we can see in Exhibit 5.3 even before we start this exercise, the Nelson-Siegel function will not fit the actual market data we have assumed because there are more data points than there are Nelson-Siegel parameters *and* the functional form of Nelson-Siegel is not flexible enough to handle the kind of actual U.S. Treasury data we reviewed earlier in this blog. The other thing that is important to note is that Nelson-Siegel will *never* be superior to a function that has the same constraints and is derived from either (1) the minimum length criterion for best or (2) the maximum smoothness criterion.

We have explicitly chosen to answer no to questions 4 to 9 in Example A, while the answers for Nelson-Siegel are yes. This is not a virtue of Nelson-Siegel; it is just a difference in modeling assumptions. In fitting Nelson-Siegel to our actual data, we have to make a choice of the function we are optimizing:

- 1. Minimize the sum of squared errors versus the actual zero-coupon yields.
- 2. Minimize the sum of squared errors versus the actual zero-coupon bond prices.

If we were using coupon bond price data, we would always optimize on the sum of squared pricing errors versus true net present value (price plus accrued interest),

EXHIBIT 5.3 Yield Curve–Smoothing Methodology

Name	Nelson-Siegel	Example A Yield Step Function
1. Fit observed data exactly?	No	Yes
2. Element of analysis (yields or forwards)	Yields	Yields
3. Max smoothness or minimum length?	Neither	Minimum length
4. Continuous?	Yes	No
5. Differentiable?	Yes	No
6. Twice differentiable?	Yes	No
7. Thrice differentiable?	Yes	No
8. Time zero constraint?	No	No
9. Longest maturity time <i>T</i> constraint?	No	No

because legacy yield quotations are distorted by inaccurate embedded forward rate assumptions. In this case, however, all of the assumed inputs are on a zero-coupon basis, and we have another issue to deal with. The zero-coupon bond price at a maturity = 0 is 1 for all yield values, so using the zero price at a maturity of zero for optimization is problematic. This means that we need to optimize in such a way that we minimize the sum of squared errors versus the zero-coupon yields at all of the input maturities, including the zero point. We need to make two other adjustments before we can do this optimization using common spreadsheet software:

- We optimize versus the sum of squared errors in yields times 1 million in order to minimize the effect of rounding error and tolerance settings embedded in the optimization routine.
- We note that, at maturity zero, the Nelson-Siegel function “blows up” because of a division by zero. Since $y(0)$ is equal to the forward rate function at time zero $f(0)$, we make that substitution to avoid dividing by zero.

That is, at the zero maturity point, instead of using the Nelson-Siegel yield function:

$$y(t) = \alpha + (\beta + \gamma) \frac{\delta}{t} \left(1 - \exp\left(\frac{-t}{\delta}\right) \right) - \gamma \exp\left(\frac{-t}{\delta}\right)$$

we use the forward rate function:

$$f(t) = \alpha + \exp\left[\frac{-t}{\delta}\right] \left[\beta + \frac{\gamma}{\delta} t \right]$$


We now choose the values of alpha, beta, delta, and gamma that minimize the sum of squared errors (times 1 million) in the actual and fitted zero yields. The results of that optimization are summarized in the table in Exhibit 5.4.

After two successive iterations, we reach the best fitting parameter values for alpha, beta, delta, and gamma. To those who have not used the Nelson-Siegel function before, the appallingly bad fit might come as a surprise. Even after the optimization, the errors in fitting zero yields are 32 basis points at the zero maturity, 36 basis points at 0.25, almost 12 basis points at 1 year, 36 basis points at 3 years, 33 basis points at 5 years, and 6 basis points at 10 years. The Nelson-Siegel formulation fails to meet the necessary condition for the consideration of a yield curve technique in this series: the technique *must* fit the observable data. Given this, why do people use the Nelson-Siegel technique? The only people who would use Nelson-Siegel are people who are willing to assume that the model is true and the market data is false. That is not a group likely to have a long career in risk management.

The graph in Exhibit 5.5 shows that our naïve model Example A, which has stepwise constant (and identical) forward rates and yields, fits the observable yields perfectly. The observable yields are plotted as black dots. The horizontal gray dots represent the stepwise constant forwards and yields of Example A. The lower smooth line is the Nelson-Siegel zero-coupon yield curve and the higher smooth line is the Nelson-Siegel forward rate curve as shown in Exhibit 5.5.

EXHIBIT 5.4 Nelson-Siegel Results

Maturity	Instantaneous Yield	Zero Price	NS Yield	NS Zero Price	NS Squared Price Error	NS Squared Yield Error
0	4.000%	1.000000	4.323%	1.000000	na	0.0000104
0.25	4.750%	0.988195	4.399%	0.989062	0.0000008	0.0000123
1	4.500%	0.955997	4.618%	0.954872	0.0000013	0.0000014
3	5.500%	0.847894	5.140%	0.857110	0.0000849	0.0000130
5	5.250%	0.769126	5.584%	0.756384	0.0001624	0.0000112
10	6.500%	0.522046	6.436%	0.525396	0.0000112	0.0000004
Sum of Squared Errors times 1 Million					260.5679209	48.70567767
Nelson Siegel Segment Length (After Optimization)					10.23137293	
Nelson Siegel Parameters (After Optimization)						
Alpha		0.0926906				
Beta		−0.0494584				
Gamma		0.0000039				
Delta		8.0575982				

**EXHIBIT 5.5** Example A: Continuous Yields and Forward Rates for Stepwise Constant Yields vs. Nelson-Siegel Smoothing

Now we pose a different question: Given that we have defined the “best” yield curve as the one with the shortest length, how does the length of the Nelson-Siegel curve compare with Example A’s length of 10 units?

There are two ways to answer this question. First, we could evaluate this integral using the first derivative of the Nelson-Siegel yield formula to evaluate the length of the curve, substituting y' for f' below:

$$s = \int_a^b \sqrt{1 + [f'(x)]^2} dx$$

The second alternative is suggested by the link on length of an arc (given earlier): We can approximate the Nelson-Siegel length calculation by using a series of straight line segments and the insights of Pythagoras to evaluate length numerically. When we do this at monthly time intervals, the first 12 months’ contributions to length are as follows:

Zero Yield Curve		
Maturity	NS Yield $\times 100$	NS Segment Length
0.000	4.323222	0.000000
0.083	4.348709	0.087134
0.167	4.374023	0.087093
0.250	4.399164	0.087043
0.333	4.424133	0.086994
0.417	4.448931	0.086945
0.500	4.473558	0.086896
0.583	4.498018	0.086849
0.667	4.522311	0.086802
0.750	4.546437	0.086756
0.833	4.570399	0.086710
0.917	4.594198	0.086665
1.000	4.617835	0.086621

We sum up the lengths of each segment, 120 months in total, to get a total line length of 10.2314, compared to a length of 10.0000 for Example A’s derived “best” curve, a step function of forward rates and yields. Note that the length of the curve, when the segments are not flat, depends on whether yields are displayed in percent (4.00) or decimal (0.04). To make the differences in length more obvious, our length calculations are based on yields in percent.

In this section, we have accomplished three things:

1. We have shown that the functional form used for yield curve fitting can and should be derived from a mathematical definition of “best” rather than being asserted for qualitative reasons.
2. We have shown that the Nelson-Siegel yield curve fitting technique fails to fit yield data with characteristics often found in the U.S. Treasury bond market.

3. We have shown that the Nelson-Siegel technique is inferior to the step-wise constant forward rates and yields that were derived from Example A's specifications: that "best" means the shortest length and that a continuous yield function is not required, consistent with the Jarrow and Yildirim paper cited earlier.

EXAMPLE D: QUADRATIC YIELD SPLINES AND RELATED FORWARD RATES

In this section, we turn to Example D, in which we seek to create smoother yield and forward rate functions by requiring that the yield curve be continuous and where the first derivatives of the curve segments we fit be equal at the knot points where the segments are joined together. We use a quadratic spline of yields to achieve this objective, and we optimize to produce the maximum tension/minimum length yields and forwards consistent with the quadratic splines. Finally, we compare the results to the popular but flawed Nelson-Siegel approach and gain still more insights on how to further improve the realism of our smoothing techniques.

In this section, we make two more modifications in our answers to the previous nine questions and derive, not assert, the best yield curve consistent with the definition of "best" given the constraints we impose.

Step 1: Should the smoothed curves fit the observable data exactly?

- 1a. Yes. Our answer is unchanged. With only six data points at six different maturities, a technique that cannot fit this simple data (like Nelson-Siegel cannot) is too simplistic for practical use.

Step 2: Select the element of the yield curve and related curves for analysis

- 2a. Zero-coupon yields is our choice for Example D. We continue to observe that we would never choose 2a or 2b to smooth a curve where the underlying securities issuer is subject to default risk. In that case, we would make the choices of either 2c or 2d. We do that later in this series.

Step 3: Define "best curve" in explicit mathematical terms

- 3b. Minimum length of curve. We continue with this criterion for best.

Step 4: Is the curve constrained to be continuous?

- 4b. Yes. We now insist on continuous yields and see what this implies for forward rates.

Step 5: Is the curve differentiable?

- 5a. Yes. This is another major change in Example D. We seek to create smoother yields and forward rates by requiring that the first derivatives of two curve segments be equal at the knot point where they meet.

Step 6: Is the curve twice differentiable?

- 6b. No. The curve will not be twice differentiable at some points on the full length of the curve given this answer, and we evaluate whether this assumption should be changed given the results we find.

Step 7: Is the curve thrice differentiable?

- 7b. No. The reason is due to our choice of 6b.

Step 8: At the spot date, time 0, is the curve constrained?

- 8c. No. For simplicity, we again answer no to this question. We wait until later in this chapter to explore this option.

Step 9: At the longest maturity for which the curve is derived, time T , is the curve constrained?

9a. Yes. This is the third major change in Example D. As we explain below, our other constraints and our definition of “best” put us in a position where the parameters of the best yield curve are not unique without one more constraint. We impose this constraint to obtain unique coefficients and then optimize the parameter used in this constraint to achieve the best yield curve.

Now that all of these choices have been made, both the functional form of the line segments for forward rates and the parameters that are consistent with the data can be explicitly derived from our sample data.

DERIVING THE FORM OF THE YIELD CURVE IMPLIED BY EXAMPLE D

Our data set has observable yield data at maturities of 0, 0.25 years, 1, 3, 5, and 10 years. For each segment of the yield curve, we have two constraints per segment (that the segment produces the market yields at the left side and right side of the segment). In addition to this, our imposition of the requirement that the full curve be differentiable means that there are other constraints. We have six knot points in total and four interior knot points (0.25, 1, 3, and 5) where we require the adjacent yield curve segments to produce the same first derivative of yields at a given knot point. That means we have 14 constraints (not including the constraint in step 9), but a linear yield spline would only have two parameters per segment and 10 parameters in total.

That means that we need to step up to a functional form for each segment that has more parameters. Our objective is not to reproduce the intellectual history of splines, because it is too voluminous, not to mention too hard. Indeed, a Google search on “quadratic splines” on one occasion produced applications to automotive design (where much of the original work on splines was done), rocket science (really), engineering, computer graphics, and astronomy. Instead, we want to focus on the unique aspect of using splines and other smoothing techniques in finance, where both yields and related forwards are important. The considerations for “best” in finance are dramatically different than they would be for the designer of the hood for a new model of the Maserati.

In the current Example D, we are still defining the best yield curve as the one with the shortest possible length or maximum tension. When we turn to cubic splines, we will discuss the proof that cubic splines produce the smoothest curve. In this example, we numerically force the quadratic splines we derive to be the shortest possible quadratic splines that are consistent with the constraints we impose in our continuing search for greater financial realism.

We can derive the shortest possible quadratic spline implementation in two ways. First, we could evaluate the function s given in Step 3, since for the first time the full yield curve will be continuously differentiable, even at the knot points. Alternatively, we can again measure the length of the curve by approximating the curve with a series of line segments, thanks to Pythagoras, whom we credited in Example A.

We chose the latter approach for consistency with our earlier examples. We break the 10-year yield curve into 120 monthly segments to numerically measure length. We now derive the set of five quadratic spline coefficients that are consistent with the constraints we have imposed.

Each yield curve segment has the quadratic form

$$y_i(t) = a_i + b_{i1}t + b_{i2}t^2$$

The subscript i refers to the segment number. The segment from 0 to 0.25 years is segment 1, the segment from 0.25 years to 1 year is segment 2, and so on. As mentioned previously, we have 10 constraints that require the left-hand side and the right hand of each segment to equal the actual yield at that maturity, which we call y^* . For the j th line segment, we have two constraints, one at the left-hand side of the segment, where the years to maturity are t_j , and one at the right-hand side of the segment t_{j+1} .

$$\begin{aligned} y^*(t_j) &= a_j + b_{j1}t_j + b_{j2}t_j^2 \\ y^*(t_{j+1}) &= a_j + b_{j1}t_{j+1} + b_{j2}t_{j+1}^2 \end{aligned}$$

In addition we have four constraints at the interior knot points. At each of these knot points, the first derivatives of the two segments that join at that point must be equal:

$$b_{j1} + 2b_{j2}t_{j+1} = b_{j+1,1} + 2b_{j+1,2}t_{j+1}$$

When we solve for the coefficients, we rearrange these four constraints in this manner:

$$b_{j1} + 2b_{j2}t_{j+1} - b_{j+1,1} - 2b_{j+1,2}t_{j+1} = 0$$

Our last constraint is imposed at the right-hand side of the yield curve, time $T = 10$ years. We will discuss alternative specifications later, but for now we assume that we want to constrain the first derivative of the yield curve at $T = 10$ to take a specific value x :

$$b_{j1} + 2b_{j2}T = x$$

In the equation above, $j = 5$, the fifth line segment, and $T = 10$. The most commonly used value for x is zero, the constraint that the yield curve be flat where $T = 10$. Our initial implementation will use $x = 0$.

In matrix form, our constraints look like the one in Exhibit 5.6.

Note that it is the last element of the y vector matrix where we have set the constraint that the first derivative of the yield curve be zero at 10 years. When we invert the coefficient matrix, we get the result shown in Exhibit 5.7.

EXHIBIT 5.6 Coefficient Matrix for Quadratic Yield Curve Line Segments

Equation Number	a1	b11	b12	a2	b21	b22	a3	b31	b32	a4	b41	b42	a5	b51	b52	Coefficient Vector	y Vector
1	1	0.000	0.000	0	0	0	0	0	0	0	0	0	0	0	0	a1	4.000%
2	1	0.25	0.0625	0	0	0	0	0	0	0	0	0	0	0	0	b11	4.750%
3	0	0	0	1	0.250	0.063	0	0	0	0	0	0	0	0	0	b12	4.750%
4	0	0	0	1	1	1	0	0	0	0	0	0	0	0	0	a2	4.500%
5	0	0	0	0	0	0	1	1.000	1.000	0	0	0	0	0	0	b21	4.500%
6	0	0	0	0	0	0	1	3	9	0	0	0	0	0	0	b22	5.500%
7	0	0	0	0	0	0	0	0	0	1	3.000	9.000	0	0	0	a3	5.500%
8	0	0	0	0	0	0	0	0	0	1	5	25	0	0	0	b31	= 5.250%
9	0	0	0	0	0	0	0	0	0	0	0	0	1	5.000	25.000	b32	5.250%
10	0	0	0	0	0	0	0	0	0	0	0	0	1	10	100	a4	6.500%
11	0	1	0.5	0	-1	-0.5	0	0	0	0	0	0	0	0	0.000	b41	0.000%
12	0	0	0	0	1	2	0	-1	-2	0	0	0	0	0	0.000	b42	0.000%
13	0	0	0	0	0	0	0	1	6	0	-1	-6	0	0	0.000	a5	0.000%
14	0	0	0	0	0	0	0	0	0	0	1	10	0	-1	-10	b51	0.000%
15	0	0	0	0	0	0	0	0	0	0	0	0	0	1	20.000	b52	0.000%

EXHIBIT 6.7 Inverted Coefficient Matrix

Row	a1	b11	b12	a2	b21	b22	a3	b31	b32	a4	b41	b42	a5	b51	b52
1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	-8	8	2.666666667	-2.666666667	-1	1	1	-1	-0.4	0.4	-1	1	-1	1	-1
3	16	-16	-10.66666667	10.66666667	4	-4	-4	4	1.6	-1.6	4	-4	4	-4	4
4	0	0	1.777777778	-0.777777778	-0.333333333	0.333333333	0.333333333	-0.333333333	-0.133333333	0.133333333	0	-0.333333333	-0.333333333	0.333333333	-0.333333333
5	0	0	-3.555555556	3.555555556	1.666666667	-1.666666667	-1.666666667	1.666666667	0.666666667	-0.666666667	0	-1.666666667	1.666666667	-1.666666667	1.666666667
6	0	0	1.777777778	-1.777777778	-1.333333333	1.333333333	1.333333333	-1.333333333	-0.533333333	0.533333333	0	1.333333333	-1.333333333	1.333333333	-1.333333333
7	0	0	0	0	2.25	-1.25	-1.5	1.5	-0.6	-0.6	0	0	1.5	-1.5	1.5
8	0	0	0	0	-1.5	1.5	2	-2	-0.8	0.8	0	0	-2	2	-2
9	0	0	0	0	0.25	-0.25	-0.5	0.5	-0.2	-0.2	0	0	0.5	-0.5	0.5
10	0	0	0	0	0	0	6.25	-5.25	-3	3	0	0	0	7.5	-7.5
11	0	0	0	0	0	0	-2.5	2.5	1.6	-1.6	0	0	0	-4	4
12	0	0	0	0	0	0	0.25	-0.25	-0.2	0.2	0	0	0	0.5	-0.5
13	0	0	0	0	0	0	0	0	4	-3	0	0	0	0	10
14	0	0	0	0	0	0	0	0	-0.8	0.8	0	0	0	0	-3
15	0	0	0	0	0	0	0	0	0.04	-0.04	0	0	0	0	0.2

We solve for the following coefficient values:

Coefficient Vector	Values
a1	0.04
b11	0.08416667
b12	-0.21666667
a2	0.05527778
b21	-0.03805556
b22	0.02777778
a3	0.02125
b31	0.03
b32	-0.00625
a4	0.105625
b41	-0.02625
b42	0.003125
a5	0.015
b51	0.01
b52	-0.0005

These coefficients give us the five quadratic functions that make up the yield curve. We use the fourth relationship below to derive the forward rate curve:

$$y(t) = \frac{-1}{t} \ln[p(t)]$$

$$p(t) = \exp\left[-\int_0^t f(s)ds\right]$$

$$p(t) = \exp[-y(t)t]$$

$$f(t) = y(t) + ty'(t)$$

The forward rate function for any segment j is

$$f_j(t) = a_j + 2b_{j1}t + 3b_{j2}t^2$$

We then plot the yield and forward rate curves to see whether or not our definition of “best yield curve” and the related constraints we imposed produced a realistic yield curve and forward rate curve as shown in Exhibit 5.8.

The answer for Example D so far is mixed. The yield curve itself fits the data perfectly, our standard necessary condition. The yield curve is also reasonably smooth and looks quite plausible given where the observable yields (solid black dots) are. The bad news, as usual, relates to forward rates. As in prior examples, it is much easier to spot flaws in the assumptions when looking at forward rates than it is when looking at yields alone. This is one of the reasons why best practice in the financial application of splines is so different than it is in computer graphics or automotive design, where there is no equivalent of the forward rate curve.




EXHIBIT 5.8 Example D: Quadratic Y with $Y'(T) = 0$ and Related Forwards

The lack of realism in the forward rates can be summarized as follows:

- The forward curve segments do not have first derivatives that are equal at the knot points where the forward segments meet; the forward curve has serious “kinks” in it at the knot point
- The movements of the forward curve are not smooth by any definition

It's possible that our fifteenth constraint, $y'(10) = 0$, should be changed to some other value $y'(10) = x$. We explore that below. Before doing so, however, we overlay the yield and forward curves on the best fitting (but not perfectly fitting) Nelson-Siegel curve (see Exhibit 5.9).

The graph makes two key points. First of all, the Nelson-Siegel curves are simply wrong and do not pass through the actual data points (in black). We know that already. Why, then, does anyone use Nelson-Siegel? For some, the smoothness of the forward curves is so seductive that the fact that valuation using the curve is wrong seems not to concern these analysts.

We clearly need to do better than a quadratic yield spline to get the smooth forward rates that we want. Let's first ask the question: “Can we improve the quadratic spline approach by optimizing the value of x in the fifteenth constraint, $y'(10) = x$?” This insight comes from a working paper by Janosi and Jarrow (2002), who noted that this constraint and others like it could have important implications for the best curve. Our first implementation was $x = 0$, a flat yield curve on the long end of the curve. We now look at alternative values for x .




EXHIBIT 5.9 Example D: Quadratic Y with $Y'(T) = 0$ and Related Forwards vs. Nelson-Siegel Smoothing

There are two criteria for best that we could use given our previous statement that the best curve is the one that is the shortest, the one with maximum tension. We could choose the value of x that produces the shortest yield curve or the value of x that produces the shortest forward rate curve. We do both.

If we use common spreadsheet software and iterate x so that the length (calculated numerically at 120 monthly intervals) of the yield curve is minimized, we get this set of coefficients:

Coefficient Vector	Values
a1	0.04
b11	0.07776339
b12	-0.19105355
a2	0.05314335
b21	-0.02738342
b22	0.01924007
a3	0.03085492
b31	0.01719344
b32	-0.00304836
a4	0.05760041
b41	-0.00063689
b42	-7.6639E-05
a5	0.073903279
b51	-0.00920984
b52	0.00078066

If we iterate x so that the length of the forward rate curve is minimized, we get still another set of coefficients:

Coefficient Vector	Values
a1	0.04
b11	0.080529
b12	-0.20211601
a2	0.05406522
b21	-0.03199278
b22	0.02292756
a3	0.0267065
b31	0.02272467
b32	-0.00443117
a4	0.07834252
b41	-0.01169935
b42	0.00130617
a5	0.05137664
b51	-0.00091299
b52	0.00022753

The value of x that produces the shortest yield curve is 0.00640. The value of x that produces the shortest forward rate curve is 0.00364. We can compare these three variations on the yield curve produced by the quadratic spline approach (see Exhibit 5.10).




EXHIBIT 5.10 Example D: Yields for Quadratic Y with $Y'(T) = 0$, $Y'(T)$ to Minimize Yield Length, and $Y'(T)$ to Minimize Forward Length

Our base case is the one that is highest on the right-hand side of the graph, and the flatness of the yield curve on the right-hand side of the curve is readily apparent. Relaxing that constraint to produce the shortest yield curve produces the curve that is lowest on the right-hand side of the graph, where the yield curve is rising sharply at the 10-year point. The shortest forward rate curve produces the yield curve in the middle. When we look at forward rates associated with these yield curves, we get an improvement. Nonetheless, it is clear that we can remove the kinks in the forward rate curves like those seen previously and generate more realistic movements. For that reason we move on to the next example.

EXAMPLE F: CUBIC YIELD SPLINES AND RELATED FORWARDS

In this section, we make an important change in our definition of the best yield curve. Here we change the definition of “best” to be the yield curve with “maximum smoothness,” which we define mathematically. Given the other constraints we impose for realism, we find that this definition of best implies that yield curve should be formed from a series of cubic polynomials. We then turn to Example F, in which we apply cubic splines to yields and derive the related forward rates.

We make a number of changes in our answers to the nine questions we have posed in each example so far.

Step 1: Should the smoothed curves fit the observable data exactly?

1a. Yes. Our answer is unchanged.

Step 2: Select the element of the yield curve and related curves for analysis

2a. Zero-coupon yields are our choice for Example F. We continue to observe that we would never choose 2a or 2b to smooth a curve where the underlying securities issuer is subject to default risk. In that case, we would make the choices in either 2c or 2d. We do that later in this book.

Step 3: Define “best curve” in explicit mathematical terms

3a. Maximum smoothness. For the first time, we choose maximum smoothness as an intuitive criterion for best, since in past installments we have critiqued the results from some assumptions because of the lack of smoothness in either yields, forward rates, or both. Smoothness is defined as the variable z such that

$$z = \int_0^T g''(s)^2 ds$$

In this case, the function $g(s)$ is defined as the yield curve. In order to evaluate z over the full maturity spectrum of the yield curve, the yield segments must be twice differentiable at each point, including the knot points. We want to minimize z over the full length of the yield curve. Our answers in Steps 4 to 6 make the analytical valuation of s possible just as when we evaluated length in Example D.

Step 4: Is the curve constrained to be continuous?

4b. Yes. We insist on continuous yields and see what this implies for forward rates.

Step 5: Is the curve differentiable?

5a. Yes. This is a major change imposed in Example D. We seek to take the spikes out of yields and forward rates by requiring that the first derivatives of two yield curve segments be equal at the knot point where they meet. This constraint also means that a linear “curve” isn’t sufficiently rich to satisfy the constraints we’ve imposed.

Step 6: Is the curve twice differentiable?

6a. Yes. For the first time, we insist that the yield curve be twice differentiable everywhere along the curve, including the knot points. This would allow us to evaluate smoothness analytically if we wished to do so, and it insures that the yield curve segments will join in a visibly smooth way at the knot points.

Step 7: Is the curve thrice differentiable?

7b. No. We do not need to specify this constraint as yet.

Step 8: At the spot date, time 0, is the curve constrained?

8b. Yes. For the first time, we find it necessary to constrain the yield curve at its left-hand side in order to derive a unique set of coefficients for each segment of the yield curve.

Step 9: At the longest maturity for which the curve is derived, time T , is the curve constrained?

9a. Yes. As in question 8, our other constraints and our definition of best put us in a position where the parameters of the best yield curve are not unique without one more constraint. We impose this constraint to obtain unique coefficients and then optimize the parameter used in this constraint to achieve the best yield curve. We explore choice 9b later in this chapter. The constraint here is imposed on the yield curve, not the forward rate curve.

Now that all of these choices have been made, both the functional form of the line segments for yields and the parameters that are consistent with the data can be explicitly derived from our sample data.

DERIVING THE FORM OF THE YIELD CURVE IMPLIED BY EXAMPLE F ASSUMPTIONS

Our data set has observable yield data at maturities of 0, 0.25 years, 1, 3, 5, and 10 years. We have six knot points in total and four interior knot points (0.25, 1, 3, and 5) where we require (1) the adjacent yield curve segments to produce the same yield values; (2) the same first derivatives of yields; and (3) the same second derivatives of yields.

That means that we need to step up to a functional form for each yield curve segment that has more parameters than the quadratic segments we used in Examples D and E.²

From our definition of best, maximum smoothness, there is no question what mathematical function is the best: it is a spline of cubic polynomials. This can be proven using the calculus of variations, as was done in the case of maximum smoothness forward rates in the proof contributed by Oldrich Vasicek, reproduced in the appendix at the end of this chapter. As is the case throughout the series, we are not *asserting* that a functional form is best like advocates of the Nelson-Siegel approach. Instead, we are defining “best” mathematically, imposing constraints that are essential to realism in financial applications and *deriving* the fact that, in Example F, it is the cubic spline that is best.

As in prior examples where we measured tension or length, we can measure the smoothness of a yield curve in two ways. First, we could explicitly evaluate the integral that defines the smoothness statistic z since our constraints for Example F will result in a yield curve that is twice differentiable over its full length. The other alternative is to use a discrete approximation to evaluation of the integral. This is attractive because it allows us to calculate smoothness for the other examples in this series for which the relevant functions are not twice differentiable. For that reason, we use a discrete measure of smoothness at 1/12-year maturity intervals over the 120 months of the yield curve we derive.

For any function x , the first difference between points i and $i - 1$ is

$$x(i) - x(i - 1)$$

The first derivative can be approximated by

$$\frac{x(i) - x(i - 1)}{\Delta i}$$

The second derivative is

$$\frac{x'(i) - x'(i - 1)}{\Delta i} = \frac{x(i) - 2x(i - 1) + x(i - 2)}{(\Delta i)^2}$$

To evaluate smoothness numerically, we calculate the second derivative at 1/12-year intervals, square it, and sum over the full length of the yield curve. With this background out of the way, we now derive the maximum smoothness yield curve.

Each yield curve segment has the cubic form

$$y_i(t) = a_i + b_{i1}t + b_{i2}t^2 + b_{i3}t^3$$

The subscript i refers to the segment number. The segment from 0 to 0.25 years is segment 1, the segment from 0.25 years to one year is segment 2, and so on. The first 10 constraints require the yield curve to be equal to the observable values of y at either the left-hand or right-hand side of that segment of the yield

curve, so for each of the five segments there are two constraints like this where y^* denotes observable data:

$$y^*(t_j) = y(t_j) = a_j + b_{j1}t_j + b_{j2}t_j^2 + b_{j3}t_j^3$$

In addition we have four constraints that require the first derivatives of the yield curve segments to be equal at the four interior knot points:

$$y'_j(t_{j+1}) = b_{j1} + 2b_{j2}t_{j+1} + 3b_{j3}t_{j+1}^2 = y'_{j+1}(t_{j+1}) = b_{j+1,1} + 2b_{j+1,2}t_{j+1} + 3b_{j+1,3}t_{j+1}^2$$

at the interior knot points. We rearrange these four constraints so that there is zero on the right hand side of the equation:

$$b_{j1} + 2b_{j2}t_{j+1} + 3b_{j3}t_{j+1}^2 - b_{j+1,1} - 2b_{j+1,2}t_{j+1} - 3b_{j+1,3}t_{j+1}^2 = 0$$

At each of these knot points, the second derivatives must also be equal:

$$2b_{j2} + 6b_{j3}t_{j+1} = 2b_{j+1,2} + 6b_{j+1,3}t_{j+1}$$

When we solve for the coefficients, we will rearrange these four constraints in this manner:

$$2b_{j2} + 6b_{j3}t_{j+1} - 2b_{j+1,2} - 6b_{j+1,3}t_{j+1} = 0$$

So far, we have 18 constraints to solve for $5 \times 4 = 20$ coefficients. Our nineteenth constraint was imposed in our answer to question 8, where we required that the second derivative of the yield curve was zero when maturity t_j is zero or some other value x_1 :

$$2b_{j2} + 6b_{j3}t_j = x_1$$

Our last constraint is imposed by our answer to question 9 at the right-hand side of the forward rate curve, time $T = 10$ years. We discuss the alternative specifications later in the chapter, but for now assume we want to constrain the first derivative of the yield curve at $t_{j+1} = 10$ to take a specific value x_2 , a value commonly selected to be zero:

$$y'_j(t_{j+1}) = b_{j1} + 2b_{j2}t_{j+1} + 3b_{j3}t_{j+1}^2 = x_2$$

Our initial implementation will use x_1 and $x_2 = 0$.
In matrix form, our constraints look Exhibit 5.11.

EXHIBIT 5.11 Coefficient Matrix for Cubic Yield Curve Line Segments

Equation Number	a1	b11	b12	b13	a2	b21	b22	b23	a3	b31	b32
1	1.00000	0	0	0	0	0	0	0	0	0	0
2	1.00000	0.25000	0.06250	0.01563	0	0	0	0	0	0	0
3	0	0	0	0	1.00000	0.25000	0.06250	0.01563	0	0	0
4	0	0	0	0	1.00000	1.00000	1.00000	1.00000	0	0	0
5	0	0	0	0	0	0	0	0	1.00000	1.00000	1.00000
6	0	0	0	0	0	0	0	0	1.00000	3.00000	9.00000
7	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0	0
11	0	1.00000	0.50000	0.18750	0	-1.00000	-0.50000	-0.18750	0	0	0
12	0	0	0	0	0	1.00000	2.00000	3.00000	0	-1.00000	-2.00000
13	0	0	0	0	0	0	0	0	0	1.00000	6.00000
14	0	0	0	0	0	0	0	0	0	0	0
15	0	0.00000	2.00000	1.50000	0.00000	0.00000	-2.00000	-1.50000	0	0	0
16	0	0	0	0	0	0	2.00000	6.00000	0	0	-2.00000
17	0	0	0	0	0	0	0	0	0	0	2.00000
18	0	0	0	0	0	0	0	0	0	0	0
19	0	0	2.00000	0	0	0	0	0	0	0	0
20	0	0	0	0	0	0	0	0	0	0	0

Note that it is the last two elements of the y vector matrix where we have set the constraints that the second derivative at time zero and the first derivative of the yield curve at 10 years be zero. We then invert the coefficient matrix and solve for the coefficients, which are shown here:

Coefficient	Values
a1	0.04
b11	0.03462
b12	-5.6E-19
b13	-0.07392
b2	0.038359
b21	0.054307
b22	-0.07875
b23	0.031082
a3	0.072955
b31	-0.04948
b32	0.025037
b33	-0.00351
a4	-0.06253
b41	0.086005
b42	-0.02012
b43	0.001505
a5	0.162438
b51	-0.04898
b52	0.006872
b53	-0.00029

b33	a4	b41	b42	b43	a5	b51	b52	b53	Coefficient Vector	y Vector
0	0	0	0	0	0	0	0	0	a1	4.000%
0	0	0	0	0	0	0	0	0	b11	4.750%
0	0	0	0	0	0	0	0	0	b12	4.750%
0	0	0	0	0	0	0	0	0	b13	4.500%
1.00000	0	0	0	0	0	0	0	0	a2	4.500%
27.00000	0	0	0	0	0	0	0	0	b21	5.500%
0	1.00000	3.00000	9.00000	27.00000	0	0	0	0	b22	5.500%
0	1.00000	5.00000	25.00000	125.00000	0	0	0	0	b23	= 5.250%
0	0	0	0	0	1.00000	5.00000	25.00000	125.00000	a3	5.250%
0	0	0	0	0	1.00000	10.00000	100.00000	1000.00000	b31	6.500%
0	0	0	0	0	0	0	0	0	b32	0.000%
-3.00000	0	0	0	0	0	0	0	0	b33	0.000%
27.00000	0	-1.00000	-6.00000	-27.00000	0.00000	0.00000	0.00000	0.00000	a4	0.000%
0	0	1.00000	10.00000	75.00000	0.00000	-1.00000	-10.00000	-75.00000	b41	0.000%
0	0	0	0	0	0	0	0	0	b42	0.000%
-6.00000	0	0	0	0	0	0	0	0	b43	0.000%
18.00000	0	0	-2.00000	-18.00000	0	0	0	0	a5	0.000%
0	0	0	2.00000	30.00000	0	0	-2.00000	-30.00000	b51	0.000%
0	0	0	0	0	0	0	0	0	b52	0.000%
0	0	0	0	0	0	1.00000	20.00000	300.00000	b53	0.000%

These coefficients give us the five cubic functions that make up the yield curve. We use the fourth relationship below to derive the forward rate curve implied by the derived coefficients for the yield curve:

$$y(t) = \frac{-1}{t} \ln[p(t)]$$

$$p(t) = \exp[-\int_0^t f(s)ds]$$

$$p(t) = \exp[-y(t)t]$$

$$f(t) = y(t) + ty'(t)$$

We can now plot the yield and forward rate curves to see whether or not our definition of “best yield curve” and the related constraints we imposed have produced a realistic yield curve and forward rate curve (see Exhibit 5.12).

Here in Example F, the yield curve itself fits the data perfectly, our standard necessary condition. The yield curve is extremely smooth given the variations in it that are required by observable yields, noted as solid black dots. The forward rate curve, the more volatile line, is also reasonably smooth but it shows some more dramatic movements than expected. Is it possible to do even better by optimizing the discrete smoothness statistic by changing x_1 and x_2 in our nineteenth and twentieth constraints? Yes. There are two criteria for best that we could use given our previous statement that the best curve is the one that is the smoothest. We could choose the values of x_1 and x_2 that produce the smoothest yield curve or the values of x_1 and x_2 that produce the smoothest forward rate curve. We do both, as before.




EXHIBIT 5.12 Example F: Cubic Yields with $y'(T) = 0$ and $y''(0) = 0$ and Related Forwards

If we use common spreadsheet software and iterate x_1 and x_2 so that the smoothness statistic z (calculated numerically at 120 monthly intervals) of the yield curve is minimized, we get the following set of coefficients.

Coefficient Vector	Values
a1	0.04
b11	0.034517
b12	0.000644
B13	−0.07485
a2	0.038345
b21	0.054379
b22	−0.0788
b23	0.031078
a3	0.072892
b31	−0.04926
b32	0.024841
b33	−0.00347
a4	−0.05723
b41	0.080857
b42	−0.01853
b43	0.00135
a5	0.125941
b51	−0.02904
b52	0.003448
b53	−0.00012

If we iterate x_1 and x_2 so that the smoothness statistic z of the forward rate curve is minimized, we get still another set of coefficients:

Coefficient Vector	Values
a1	0.04
b11	0.041167
b12	-0.04209
B13	-0.0103
a2	0.039399
b21	0.048377
b22	-0.07093
b23	0.028158
a3	0.070823
b31	-0.04589
b32	0.023337
b33	-0.00327
a4	-0.05044
b41	0.075366
b42	-0.01708
b43	0.001226
a5	0.104982
b51	-0.01789
b52	0.001567
b53	-1.8E-05

The smoothest yield curve is produced by $x_1 = 0.00129$ and $x_2 = 0.00533$. The smoothest forward rate curve is produced by $x_1 = -0.08419$ and $x_2 = 0.00811$. We graph the differences that result when we review Example H, the maximum smoothness forward rate approach.

EXAMPLE H: MAXIMUM SMOOTHNESS FORWARD RATES AND RELATED YIELDS³

We present the final installment in yield curve and forward rate smoothing techniques before later moving on to smoothing credit spreads in Chapter 17. We introduce the maximum smoothness forward rate technique introduced by Adams and van Deventer (1994) and corrected in van Deventer and Imai (1996), which we call Example H. We explain why a quartic function is needed to maximize smoothness of the forward rate function over the full length of the forward rate curve. Finally, we compare 23 different techniques for smoothing yields and forward rates that have been discussed in this series and show why the maximum smoothness forward rate approach is the best technique by multiple criteria.

Here are the answers to our normal nine questions for Example H, quartic splines of forward rates:

Step 1: Should the smoothed curves fit the observable data exactly?

- 1a. Yes. Our answer is unchanged. Since the Nelson-Siegel specification cannot meet this minimum standard, it is not included in the following comparisons.

Step 2: Select the element of the yield curve and related curves for analysis

- 2b. Forward rates are the choice for Example H. In addition to making this choice for empirically improving smoothness by changing parameter values, we add other constraints that improve our ability to maximize smoothness over the full length of the curve. We continue to observe that we would never choose 2a or 2b to smooth a curve where the underlying securities issuer is subject to default risk. In that case, we would make the choices in either 2c or 2d. (We do that in Chapter 17.)

Step 3: Define “best curve” in explicit mathematical terms

- 3a. Maximum smoothness. We unambiguously choose maximum smoothness as the primary criterion for best. Within that class of functions that are maximum smoothness, we also take the opportunity to iterate parameter values to show which of the infinite number of maximum smoothness forward curves have minimum length and maximum smoothness for each set of constraints we impose. The two measures are very consistent.

As noted earlier in this chapter, we have critiqued the results from some yield curve-smoothing techniques because of the lack of smoothness in either yields, forward rates, or both. Smoothness is defined as the variable z such that

$$z = \int_0^T g''(s)^2 ds$$

The smoothest possible function has the minimum z value. Since a straight line has a second derivative of zero, z is zero in that case and a straight line is perfectly smooth. In this case, the function $g(s)$ is defined as the forward rate curve. In order to evaluate z over the full maturity spectrum of the forward curve, the forward rate segments must be at least twice differentiable at each point, including the knot points. We want to minimize z over the full length of the yield curve. Our answers in Steps 4 to 6 make the analytical valuation of z possible just as when we evaluated length in example D.

Recall that forward rates are related to zero-coupon yields and bond prices by the same four relationships that we have been using throughout this book:

$$y(t) = \frac{-1}{t} \ln[p(t)]$$

$$p(t) = \exp\left[-\int_0^t f(s) ds\right]$$

$$p(t) = \exp[-y(t)t]$$

$$f(t) = y(t) + ty'(t)$$

Given that forward rates are in effect a derivative of zero yields and zero prices, Oldrich Vasicek uses the calculus of variations to

prove that the smoothest forward rate function that can be derived is a thrice differentiable function, the quartic spline of forward rates. This proof, reproduced in the appendix to this chapter, was corrected in van Deventer and Imai (1996) from the original version in Adams and van Deventer (1994) thanks to helpful comments from Volf Frishling of the Commonwealth Bank of Australia and further improved by the insights of Kamakura's Robert A. Jarrow. The choice of the quartic form is a *derivation* shown in the proof. Similarly, in order to achieve maximum smoothness, the same proof shows that the curves fitted together *must* be thrice differentiable over the entire length of the curve, which obviously includes the knot points. For that reason, we impose this constraint below.

Step 4: Is the curve constrained to be continuous?

4b. Yes. We insist on continuous forwards and see what this implies for yields.

Step 5: Is the curve differentiable?

5a. Yes. Again, this is the change first imposed in Example D.

Step 6: Is the curve twice differentiable?

6a. Yes. We insist that the forward rate curve be twice differentiable everywhere along the curve, including the knot points. This would allow us to evaluate smoothness analytically if we wished to do so, and it insures that the forward rate curve segments will join in a visibly smooth way at the knot points.

Step 7: Is the curve thrice differentiable?

7b. Yes. For the first time, we impose this constraint because we have a mathematical proof that we will not obtain a unique maximum smoothness forward curve unless we do impose it.

Step 8: At the spot date, time 0, is the curve constrained?

8a and 8b. Both approaches will be used and compared. Like Example F, we find it necessary to constrain the forward rate curve at its left-hand side and right-hand side in order to derive a unique set of coefficients for each segment of the forward rate curve. We compare results using both answers 8a and 8b and reach conclusions about which is best for the sample data we are analyzing. The answers to question 8 can have a critical impact on the reasonableness of splines for forward rates in a financial context. In fact, in this example, we need to use 3 of the four constraints listed in questions 8 and 9.

Step 9: At the longest maturity for which the curve is derived, time T , is the curve constrained?

9a and 9b. Both approaches will be used and compared. For uniqueness of the parameters of the quartic forward rate segment coefficients, we again have to choose three of the four constraints in questions 8 and 9. We then optimize the parameters used in these constraints to achieve the best forward rate curve as suggested by Janosi and Jarrow (2002). The constraints here are again imposed on the forward rate curve, not the yield curve. We will use both maximum smoothness and minimum

length as criterion for best, conditional on our choice of quartic forward curve segments and the other constraints we impose. We are reassured to see that the maximum smoothness forward rate technique produces results that are both the smoothest of any twice differentiable functions and the shortest length/maximum tension over the full length of the curve.

DERIVING THE PARAMETERS OF THE QUARTIC FORWARD RATE CURVES IMPLIED BY EXAMPLE H ASSUMPTIONS

Our data set has observable yield data at maturities of 0, 0.25 years, 1, 3, 5, and 10 years. We have six knot points in total and four interior knot points (0.25, 1, 3, and 5 years) where the curves that join must have equal first, second, and third derivatives.

As in Example F, we can measure the smoothness of a forward rate curve in two ways. First, we could explicitly evaluate the integral that defines the smoothness statistic z since our constraints for Example H will result in a forward rate curve that is twice differentiable over its full length. The other alternative is to use a discrete approximation to evaluation of the integral, like we did in Example F. This is attractive because it allows us to calculate smoothness for the other examples in this series for which the relevant functions are not twice differentiable. For that reason, we use the same discrete measure of smoothness at 1/12-year maturity intervals over the 120 months of the forward curve and yield curves that we derive here.

Each forward rate curve segment has the quartic form

$$f_i(t) = c_i + d_{i1}t + d_{i2}t^2 + d_{i3}t^3 + d_{i4}t^4$$

The subscript i refers to the segment number. The segment from 0 to 0.25 years is segment 1, the segment from 0.25 years to one year is segment 2, and so on. The first constraint requires the first forward rate curve segment to be equal to the observable value of y at time zero since $y(0) = f(0)$.

$$y * (t_j) = f * (t_j) = c_j + d_{j1}t_j + d_{j2}t_j^2 + d_{j3}t_j^3 + d_{j4}t_j^4$$

where for this first constraint $t_j = 0$. In addition we have four constraints that require the forward rate curves to be equal at the four interior knot points:

$$\begin{aligned} f_j(t_{j+1}) &= c_j + d_{j1}t_{j+1} + d_{j2}t_{j+1}^2 + d_{j3}t_{j+1}^3 + d_{j4}t_{j+1}^4 = f_{j+1}(t_{j+1}) \\ &= c_{j+1} + d_{j+1,1}t_{j+1} + d_{j+1,2}t_{j+1}^2 + d_{j+1,3}t_{j+1}^3 + d_{j+1,4}t_{j+1}^4 \end{aligned}$$

at the interior knot points. We rearrange these four constraints, for $j = 1, 4$ like this:

$$\begin{aligned} c_j + d_{j1}t_{j+1} + d_{j2}t_{j+1}^2 + d_{j3}t_{j+1}^3 + d_{j4}t_{j+1}^4 - c_{j+1} - d_{j+1,1}t_{j+1} \\ - d_{j+1,2}t_{j+1}^2 - d_{j+1,3}t_{j+1}^3 - d_{j+1,4}t_{j+1}^4 = 0 \end{aligned}$$

At each of these interior knot points, the first derivatives of the two forward rate segments that join at that point must also be equal:

$$d_{j1} + 2d_{j2}t_{j+1} + 3d_{j3}t_{j+1}^2 + 4d_{j4}t_{j+1}^3 = d_{j+1,1} + 2d_{j+1,2}t_{j+1} + 3d_{j+1,3}t_{j+1}^2 + 4d_{j+1,4}t_{j+1}^3$$

When we solve for the coefficients, we will rearrange these four constraints in this manner:

$$d_{j1} + 2d_{j2}t_{j+1} + 3d_{j3}t_{j+1}^2 + 4d_{j4}t_{j+1}^3 - d_{j+1,1} - 2d_{j+1,2}t_{j+1} - 3d_{j+1,3}t_{j+1}^2 - 4d_{j+1,4}t_{j+1}^3 = 0$$

Next, we need to impose the constraint that the second derivatives of the joining forward rate curve segments are equal at each of the four interior knot points. This requires

$$2d_{j2} + 6d_{j3}t_{j+1} + 12d_{j4}t_{j+1}^2 = 2d_{j+1,2} + 6d_{j+1,3}t_{j+1} + 12d_{j+1,4}t_{j+1}^2$$

As usual, we rearrange this as follows:

$$2d_{j2} + 6d_{j3}t_{j+1} + 12d_{j4}t_{j+1}^2 - 2d_{j+1,2} - 6d_{j+1,3}t_{j+1} - 12d_{j+1,4}t_{j+1}^2 = 0$$

Finally, the new constraint that the third derivatives be equal at each knot point requires that

$$6d_{j3} + 24d_{j4}t_{j+1} = 6d_{j+1,3} - 24d_{j+1,4}t_{j+1}$$

We arrange this constraint to read as follows:

$$6d_{j3} + 24d_{j4}t_{j+1} - 6d_{j+1,3} - 24d_{j+1,4}t_{j+1} = 0$$

So far, we have $1 + 4 + 4 + 4 + 4 = 17$ constraints to solve for $5 \times 5 = 25$ coefficients. We have five more constraints that require that the forward curve segments produce observable zero-coupon bond prices:

$$\begin{aligned} -\ln \left[\frac{p(t_k)}{p(t_j)} \right] &= \left[\int_{t_j}^{t_k} f_j(s) ds \right] = c_j(t_k - t_j) + \frac{1}{2}d_{j1}(t_k^2 - t_j^2) + \frac{1}{3}d_{j2}(t_k^3 - t_j^3) \\ &\quad + \frac{1}{4}d_{j3}(t_k^4 - t_j^4) + \frac{1}{5}d_{j4}(t_k^5 - t_j^5) \end{aligned}$$

For our twenty-third, twenty-fourth, and twenty-fifth constraints, we can choose any three of the constraints in 8a, 8b, 9a, and 9b. We choose three combinations that are very common in financial applications where it is logical and realistic to expect the forward rate curve to be flat at the right-hand side of the yield curve, $f'(10) = 0$ in this example, where $t_{j+1} = 10$.

EXHIBIT 5.13 Coefficient Matrix for Quartic Forward Curve Line Segments

Equation Number	c1	d11	d12	d13	d14	c2	d21	d22	d23	d24	c3
1	1	0.000	0.000	0.000	0	0	0	0	0	0	0
2	1	0.25	0.0625	0.015625	0.00390625	-1	-0.25	-0.0625	-0.015625	-0.00390625	0
3	0	0	0	0	0	1	1	1	1	1	-1
4	0	0	0	0	0	0	0	0	0	0	1
5	0	0	0	0	0	0	0	0	0	0	0
6	0	1	0.5	0.1875	0.0625	0	-1	-0.5	-0.1875	-0.0625	0
7	0	0	0	0	0	0	1	2	3	4	0
8	0	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	0
10	0	0	2	1.5	0.75	0	0	-2	-1.5	-0.75	0
11	0	0	0	0	0	0	0	2	6	12	0
12	0	0	0	0	0	0	0	0	0	0	0
13	0	0	0	0	0	0	0	0	0	0	0
14	0	0	0	6	6	0	0	0	-6	-6	0
15	0	0	0	0	0	0	0	0	6	24	0
16	0	0	0	0	0	0	0	0	0	0	0
17	0	0	0	0	0	0	0	0	0	0	0
18	0.25	0.03125	0.005208333	0.000976563	0.000195313	0	0	0	0	0	0
19	0	0	0	0	0	0.75	0.46875	0.328125	0.249023438	0.199804688	0
20	0	0	0	0	0	0	0	0	0	0	2
21	0	0	0	0	0	0	0	0	0	0	0
22	0	0	0	0	0	0	0	0	0	0	0
23	0	0	2	0	0	0	0	0	0	0	0
24	0	0	0	0	0	0	0	0	0	0	0
25	0	0	0	0	0	0	0	0	0	0	0

$$f'_j(t_{j+1}) = c_{j1} + 2d_{j2}t_{j+1} + 3d_{j3}t_{j+1}^2 + 4d_{j4}t_{j+1}^3 = 0$$

Our twenty-fourth and twenty-fifth constraints come from setting the second derivatives at time $t_{j+1} = 0$ to x_1 on the left-hand side of the curve and x_2 on the right-hand side of the curve where $t_{j+1} = 10$:

$$f''_j(t_{j+1}) = 2d_{j2} + 6d_{j3}t_{j+1} + 12d_{j4}t_{j+1}^2 = x_k$$

Our initial implementation, which we label Example H-Qf1a (quadratic forwards 1a), will use x_1 and $x_2 = 0$. Our second implementation Example H-Qf1b uses the insights of Janosi and Jarrow (2002) to optimize x_1 and x_2 to minimize the length of the forward curve. The third implementation optimizes x_1 and x_2 to maximize the smoothness of the curve by minimizing the function z given above on a discrete 1/12 of a year basis.

In matrix form, with apologies to those readers with less than perfect vision, our constraints for example H-Qf1a look like Exhibit 5.13.

Note that it is the last three elements of the “y Vector” matrix where we have set the constraints that the second derivatives at time zero and 10 and the first derivative of the yield curve at 10 years be zero. Then invert the matrix and solve for these coefficients:

d31	d32	d33	d34	c4	d41	d42	d43	d44	c5	d51	d52	d53	d54	Coefficient Vector	y Vector
0	0	0	0	0	0	0	0	0	0	0	0	0	0	c1	4.000%
0	0	0	0	0	0	0	0	0	0	0	0	0	0	d11	0
-1	-1	-1	-1	0	0	0	0	0	0	0	0	0	0	d12	0
3	9	27	81	-1	-3	-9	-27	-81	0	0	0	0	0	d13	0
0	0	0	0	1	5	25	125	625	-1	-5	-25	-125	-625	d14	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	c2	0
-1	-2	-3	-4	0	0	0	0	0	0	0	0	0	0	d21	0
1	6	27	108	0	-1	-6	-27	-108	0	0	0	0	0	d22	= 0
0	0	0	0	0	1	10	75	500	0	-1	-10	-75	-500	d23	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	d24	0
0	-2	-6	-12	0	0	0	0	0	0	0	0	0	0	c3	0
0	2	18	108	0	0	-2	-18	-108	0	0	0	0	0	d31	0
0	0	0	0	0	0	2	30	300	0	0	-2	-30	-300	d32	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	d33	0
0	0	-6	-24	0	0	0	0	0	0	0	0	0	0	d34	0
0	0	6	72	0	0	0	-6	-72	0	0	0	0	0	c4	0
0	0	0	0	0	0	0	6	120	0	0	0	-6	-120	d41	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	d42	0.01188
0	0	0	0	0	0	0	0	0	0	0	0	0	0	d43	0.03313
4	8.666666667	20	48.4	0	0	0	0	0	0	0	0	0	0	d44	0.12000
0	0	0	0	2	8	32.66666667	136	576.4	0	0	0	0	0	c5	0.09750
0	0	0	0	0	0	0	0	0	5	37.5	291.6666667	2343.75	19375	d51	0.38750
0	0	0	0	0	0	0	0	0	0	0	0	0	0	d52	0.00000
0	0	0	0	0	0	0	0	0	0	1	20	300	4000	d53	0.00000
0	0	0	0	0	0	0	0	0	0	0	2	60	1200	d54	0.00000

Coefficient Vector	Values
C1	0.0400000000
D11	0.0744582074
D12	0.0000000000
D13	-0.6332723810
D14	0.8530487205
C2	0.0363357405
D21	0.1330863589
D22	-0.3517689090
D23	0.3047780430
D24	-0.0850017035
C3	0.1272995768
D31	-0.2307689862
D32	0.1940141086
D33	-0.0590773021
D34	0.0059621327
C4	-0.5156595261
D41	0.6265098177
D42	-0.2346252934

(Continued)

Coefficient Vector	Values
D43	0.0361758984
D44	-0.0019756340
C5	0.8790434876
D51	-0.4892525933
D52	0.1001034299
D53	-0.0084545981
D54	0.0002558909

These coefficients give us the five quartic functions that make up the forward rate curve. We use the second and third relationships listed above to solve for the derived coefficients for the zero-coupon yield curve. The yield function for any segment j is

$$y_j(t) = \frac{1}{t} \left[y^*(t_j)t_j + c_j(t - t_j) + \frac{1}{2}d_{j1}(t^2 - t_j^2) + \frac{1}{3}d_{j2}(t^3 - t_j^3) + \frac{1}{4}d_{j3}(t^4 - t_j^4) + \frac{1}{5}d_{j4}(t^5 - t_j^5) \right]$$

We note that y^* denotes the observable value of y at the left-hand side of the line segment where the maturity is t_j . Within the segment, y is a quintic function of t , divided by t .

We can now plot the yield and forward rate curves to see the realism of our assumptions about maximum smoothness and the related constraints we have imposed. The results show something very interesting. When the infinite number of forward rate curves that are maximum smoothness subject to constraints 23, 24, and 25 are compared, it is seen that the real trade-off comes from sacrificing smoothness at the left-hand side of the curve in order to get shorter length overall, with a smaller rise in the forward curve at the 10-year point. Example H-Qf1a, where we have set all three assumptions 23, 24, and 25 to zero falls between the two optimized results at the 10-year point (see Exhibit 5.14).

The results for all three examples are far superior to what we saw in earlier examples. When we compare the three yield curves from our three Example H alternatives, again the results are very reasonable (see Exhibit 5.15).

Next we compare our base case, with assumptions 23, 24, and 25 all set to zero, with the arbitrary Nelson-Siegel formulation. The Nelson-Siegel functions (the two lines in the middle, with the forward rates above the zero yields) are simply two lines drawn on a page with no validity because they don't match the observable data, which are the solid black dots on the yield curve derived from our base case (see Exhibit 5.16).

We continue to be surprised that the Nelson-Siegel technique is used in academia given the fact that it is both inaccurate and more complex in calculation: the Nelson-Siegel formulation is incapable of fitting an arbitrary set of data and requires a nonlinear optimization. The maximum smoothness forward rate approach, by contrast, fits the data perfectly in every case and requires (in our base case) just a matrix inversion.




EXHIBIT 5.14 Example H Comparison of Alternative Forward Rates




EXHIBIT 5.15 Example H Comparison of Alternative Yields




EXHIBIT 5.16 Example H-1: Quartic Forwards with $f'(0) = 0$ and $f'(T) = 0$ Related Yields vs. Nelson-Siegel Smoothing

For example H-Qf1b, when we use simple spreadsheet optimize to iterate on x_1 and x_2 to derive the maximum smoothness forward rate curve of shortest length, we get these coefficients:

Coefficient Vector	Values
C1	0.04
D11	0.232654282
D12	-1.80184349
D13	5.124447912
D14	-5.19776484
C2	0.060389164
D21	-0.09357235
D22	0.155516289
D23	-0.09517817
D24	0.021861242
C3	0.040261445
D31	-0.01306147
D32	0.034749971
D33	-0.01466729
D34	0.001733522
C4	-0.18940957
D41	0.293166555
D42	-0.011836404

Coefficient Vector	Values
D43	0.019358042
D44	-0.00110192
C5	0.653581615
D51	-0.3812264
D52	0.083953844
D53	-0.00761768
D54	0.000246863

When we iterate using the Example H-Qf1c objective, maximum smoothness, we get the following set of coefficients for our five quartic forward rate segments:

Coefficient Vector	Values
C1	0.04
D11	0.077396762
D12	-0.03383111
D13	-0.52363325
D14	0.736847176
C2	0.036799797
D21	0.128600013
D22	-0.34105061
D23	0.295618762
D24	-0.08240483
C3	0.124892139
D31	-0.22376935
D32	0.187503439
D33	-0.05675061
D34	0.005687509
C4	-0.4755168
D41	0.576775895
D42	-0.21276919
D43	0.032198867
D44	-0.00172495
C5	0.672902511
D51	-0.34195955
D52	0.062851448
D53	-0.00455055
D54	0.000112524

COMPARING YIELD CURVE AND FORWARD RATE SMOOTHING TECHNIQUES

As we have discussed throughout this chapter, the standard process for evaluating yield curve and forward rate curves should involve the mathematical statement of the

objective of smoothing, the imposition of constraints that are believed to generate realistic results, and then the testing of the results to establish whether they are reasonable or not. In the last section of this chapter, we discuss the Shimko test used in Adams and van Deventer (1994) for measuring the accuracy of yield curve and forward rate smoothing techniques on large amounts of data. For expositional purposes in this section, we analyze only one case, our base case, and reach some tentative conclusions whose validity can be confirmed or denied on huge masses of data using the Shimko test.

Ranking 23 Smoothing Techniques by Smoothness of the Forward Rate Curve

The chart in Exhibit 5.17 ranks 23 smoothing techniques on their effectiveness in lowering the discrete smoothing statistic as much as possible as explained above.

As we would expect, the best result was the maximum smoothness forward rate approach where the optimization of the twenty-fourth and twenty-fifth constraints was done with respect to smoothness. The second-best approach was maximum smoothness base case, Example H-Qf1a, where we set the derivatives with respect to constraints 23, 24, and 25 to zero. The technique ranked 17th was also one of the maximum smoothness alternatives, where we intentionally sacrificed the smoothness of the curve (most notably on the short end of the curve) to minimize length. Linear yield smoothing produced the worst results when smoothness is the criterion for best.

Ranking 23 Smoothing Techniques by Length of the Forward Curve

Next we report on the discrete approximation for length of the forward curve if its length is the sole criterion. The results for 23 techniques are shown in Exhibit 5.18.

Not surprisingly, if one does not insist on continuity or smoothness at all, one gets a short forward curve. Ranks 1 to 4 are taken by a yield/forward step function and quadratic curve fitting. Among all of the techniques that are at least twice differentiable, the maximum smoothness forward rate technique, optimized to minimize length, was the winner. There is no need to sacrifice one attribute (length) for another (smoothness). In this example, the same functional form (quartic forward rates) can produce the winner by either criterion.

If we look at both attributes, which techniques best balance the trade-offs between smoothness and length? We turn to that question next.

TRADING OFF SMOOTHNESS VS. THE LENGTH OF THE FORWARD RATE CURVE

In the graph shown in Exhibit 5.19, we plot the 23 smoothing techniques on an XY graph where both length and smoothness of the forward curve are relevant. We truncated the graph to eliminate some extreme outliers as shown in Exhibit 5.19.

A reasonable person would argue that the best techniques by some balance of smoothness and length would fall in the lower left-hand corner of the graph. There are five techniques with a forward rate curve with smoothness below 7,000 and length below 30.00 (see Exhibit 5.20).

EXHIBIT 5.17 Ranking by Smoothness of Forward Rate Curve

Rank	Example	Smoothing Technique	Length of Yield Curve Fitted	Length of Forward Curve Fitted	Smoothness of Yield Curve Fitted	Smoothness of Forward Curve Fitted
1	Example H-Qf1c	Quartic Forward Rate with $f''(0)=x1$, $f''(T)=x2$ and $f'(T)=0$ to Max f Smoothness	11.58	21.01	505.76	4,287.35
2	Example H-Qf1a	Quartic Forward Rate with $f''(0)=0$, $f''(T)=0$ and $f'(T)=0$	11.51	19.91	495.93	4,309.97
3	Example F	Cubic y , Max f Smoothness	11.75	24.94	510.10	4,564.23
4	Example F	Cubic y , Max y Smoothness	11.73	22.62	473.48	4,811.18
5	Example F	Cubic y , $y'(T)=0$ and $y''(0)=0$	11.64	23.14	475.92	5,038.87
6	Example G-4Cfb	Cubic Forward Rate with $f'(0)=x1$ and $f'(T)=x2$ to Max f Smoothness	14.15	48.89	687.21	5,931.25
7	Example G-3Cfb	Cubic Forward Rate with $f''(0)=x1$ and $f'(T)=x2$ to Max f Smoothness	14.15	48.89	687.21	5,931.25
8	Example G-3Cfe	Cubic Forward Rate with $f''(0)=x1$ and $f'(T)=0$ to Max f Smoothness	13.59	40.52	690.46	6,179.14
9	Example E	Quadratic f , $f'(T)=0$	11.07	15.28	1,061.01	7,767.53
10	Example E	Quadratic f , Min f Length	11.07	15.26	1,061.76	7,779.04
11	Example E	Quadratic f , Min y Length	11.07	15.28	1,062.55	7,791.63
12	Example G-3Cfa	Cubic Forward Rate with $f''(0)=0$ and $f'(T)=0$	15.36	52.23	597.41	8,118.72
13	Example G-4Cfc	Cubic Forward Rate with $f'(0)=x1$ and $f'(T)=x2$ to Min f Length	11.38	15.79	4,111.96	24,073.20
14	Example G-3Cfc	Cubic Forward Rate with $f''(0)=x1$ and $f'(T)=x2$ to Min f Length	11.38	15.79	4,111.96	24,073.23
15	Example G-3Cfd	Cubic Forward Rate with $f''(0)=x1$ and $f'(T)=0$ to Min f Length	11.40	15.87	4,200.81	24,615.22
16	Example G-4Cfa	Cubic Forward Rate with $f'(0)=0$ and $f'(T)=0$	19.83	78.05	2,128.84	25,398.14
17	Example H-Qf1b	Quartic Forward Rate with $f''(0)=x1$, $f''(T)=x2$ and $f'(T)=0$ to Min Length of f	11.23	15.52	5,267.60	25,629.13
18	Example C	Linear forwards	12.15	31.59	619.35	27,834.72
19	Example D	Quadratic y , Min y Length	11.46	18.70	3,346.49	35,901.19
20	Example D	Quadratic y , Min f Length	11.56	17.21	3,780.67	43,108.27
21	Example D	Quadratic y , $y'(T)=0$	12.00	20.48	4,409.57	56,443.05
22	Example A	Yield step function	13.12	13.12	123,120.00	123,120.00
23	Example B	Linear yields	10.99	16.87	1,776.50	319,346.00

EXHIBIT 5.18 Ranking by Length of Forward Rate Curve

Rank	Example Description	Smoothing Technique	Length of Yield Curve Fitted	Length of Forward Curve Fitted	Smoothness of Yield Curve Fitted	Smoothness of Forward Curve Fitted
1	Example A	Yield step function	13.12	13.12	123,120.00	123,120.00
2	Example E	Quadratic f , Min f Length	11.07	15.26	1,061.76	7,779.04
3	Example E	Quadratic f , $f'(T)=0$	11.07	15.28	1,061.01	7,767.53
4	Example E	Quadratic f , Min y Length	11.07	15.28	1,062.55	7,791.63
5	Example H-Qf1b	Quartic Forward Rate with $f''(0)=x1$, $f''(T)=x2$ and $f'(T)=0$ to Min Length of f	11.23	15.52	5,267.60	25,629.13
6	Example G-3Cfc	Cubic Forward Rate with $f''(0)=x1$ and $f'(T)=x2$ to Min f Length	11.38	15.79	4,111.96	24,073.23
7	Example G-4Cfc	Cubic Forward Rate with $f'(0)=x1$ and $f'(T)=x2$ to Min f Length	11.38	15.79	4,111.96	24,073.20
8	Example G-3Cfd	Cubic Forward Rate with $f''(0)=x1$ and $f'(T)=0$ to Min f Length	11.40	15.87	4,200.81	24,615.22
9	Example B	Linear yields	10.99	16.87	1,776.50	319,346.00
10	Example D	Quadratic y , Min f Length	11.56	17.21	3,780.67	43,108.27
11	Example D	Quadratic y , Min y Length	11.46	18.70	3,346.49	35,901.19
12	Example H-Qf1a	Quartic Forward Rate with $f''(0)=0$, $f''(T)=0$ and $f'(T)=0$	11.51	19.91	495.93	4,309.97
13	Example D	Quadratic y , $y'(T)=0$	12.00	20.48	4,409.57	56,443.05
14	Example H-Qf1c	Quartic Forward Rate with $f''(0)=x1$, $f''(T)=x2$ and $f'(T)=0$ to Max f Smoothness	11.58	21.01	505.76	4,287.35
15	Example F	Cubic y , Max y Smoothness	11.73	22.62	473.48	4,811.18
16	Example F	Cubic y , $y'(T)=0$ and $y''(0)=0$	11.64	23.14	475.92	5,038.87
17	Example F	Cubic y , Max f Smoothness	11.75	24.94	510.10	4,564.23
18	Example C	Linear forwards	12.15	31.59	619.35	27,834.72
19	Example G-3Cfe	Cubic Forward Rate with $f''(0)=x1$ and $f'(T)=0$ to Max f Smoothness	13.59	40.52	690.46	6,179.14
20	Example G-3Cfb	Cubic Forward Rate with $f''(0)=x1$ and $f'(T)=x2$ to Max f Smoothness	14.15	48.89	687.21	5,931.25
21	Example G-4Cfb	Cubic Forward Rate with $f'(0)=x1$ and $f'(T)=x2$ to Max f Smoothness	14.15	48.89	687.21	5,931.25
22	Example G-3Cfa	Cubic Forward Rate with $f''(0)=0$ and $f'(T)=0$	15.36	52.23	597.41	8,118.72
23	Example G-4Cfa	Cubic Forward Rate with $f'(0)=0$ and $f'(T)=0$	19.83	78.05	2,128.84	25,398.14




EXHIBIT 5.19 Trade-Off between Smoothness and Length of Forward Rate Curve for Various Smoothing Techniques

EXHIBIT 5.20 Best Techniques: Smoothness Under 7,000, Length Under 30, Ranked by Length of Forward Curve

Example Description	Smoothing Technique	Length of Yield Curve Fitted	Length of Forward Curve Fitted	Smoothness of Yield Curve Fitted	Smoothness of Forward Curve Fitted
Example H-Qf1a	Quartic Forward Rate with $f''(0)=0$, $f''(T)=0$ and $f'(T)=0$	11.51	19.91	495.93	4,309.97
Example H-Qf1c	Quartic Forward Rate with $f''(0)=x1$, $f''(T)=x2$ and $f'(T)=0$ to Max f Smoothness	11.58	21.01	505.76	4,287.35
Example F	Cubic y , Max y Smoothness	11.73	22.62	473.48	4,811.18
Example F	Cubic y , $y'(T)=0$ and $y''(0)=0$	11.64	23.14	475.92	5,038.87
Example F	Cubic y , Max f Smoothness	11.75	24.94	510.10	4,564.23

It is interesting to see that the best five consist of two quartic forward rate smoothing approaches, ranked first and second by smoothness, and three cubic yield spline approaches. Of this best five group, the quartic forward rate approach, optimized for smoothness, was the smoothest. The quartic forward rate approach, with all derivatives in constraints 23 to 25 set to zero, was the shortest. We now turn to a more comprehensive approach to performance measurement of smoothing techniques.

THE SHIMKO TEST FOR MEASURING ACCURACY OF SMOOTHING TECHNIQUES

In 1994, Adams and van Deventer (1994) published “Fitting Yield Curves and Forward Rate Curves with Maximum Smoothness.” In 1993, our thoughtful friend David Shimko responded to an early draft of the Adams and van Deventer paper by saying, “I don’t care about a mathematical proof of ‘best,’ I want something that would have best estimated a data point that I intentionally leave out of the smoothing process—this to me is proof of which technique is most realistic.” A statistician would add, “And I want something that is most realistic on a very large sample of data.” We agreed that Shimko’s suggestion was the ultimate proof of the accuracy and realism of any smoothing technique. The common academic practice of using one set of fake data and then judging which technique “looks good” or “looks bad” is as ridiculous as it is common. Therefore, we atone for the same sin, which we have used in the first 10 installments of this series, by explaining how to perform the Shimko test as in Adams and van Deventer (1994).

The Shimko test works as follows. First, we assemble a large data set, which in the Adams and van Deventer case was 660 days of swap data. Next, we select one of the maturities in that data set and leave it out of the smoothing process. For purposes of this example, say we leave out the seven-year maturity because that leaves a wide five-year gap in the swap data to be filled by the smoothing technique. We smooth the 660 yield curves one by one. Using the smoothing results and the zero-coupon bond yields associated with the 14 semiannual payment dates of the seven-year interest rate swap, we calculate the seven-year swap rate implied by the smoothing process. We have 660 observations of this estimated seven-year swap rate, and we compare it to the actual seven-year swap rates that we left out of the smoothing process. The best smoothing technique is the one that most accurately estimates the omitted data point over the full sample.

This test can be performed on any of the maturities that were inputs to the smoothing process, and we strongly recommend that all maturities be used one at a time. Because this test suggested by Shimko is a powerful test applicable to any contending smoothing techniques, we strongly recommend that no assertion of superior performance be made without applying the Shimko test on a large amount of real data.⁴

SMOOTHING YIELD CURVES USING COUPON-BEARING BOND PRICES AS INPUTS

For expositional purposes, we have assumed in this chapter that the raw inputs to this process are zero-coupon bond yields. When, instead, the inputs are the prices and terms on coupon-bearing bonds, the analysis changes in a minor way, which we illustrate in Chapter 17. The initial zero-coupon bond yields are guessed, and an iteration of zero-coupon bond yields is performed that minimizes the sum of squared bond-pricing errors. The Dickler, Jarrow, and van Deventer analysis discussed in Chapter 3 was done in that manner using Kamakura Risk Manager version 8.0. Kamakura Corporation produces a weekly forecast of implied forward rates derived in this manner. The implied forward U.S. Treasury yields as of May 3, 2012, are shown in the graph in Exhibit 5.21.




EXHIBIT 5.21 Kamakura Corporation, 10-Year Forecast of U.S. Treasury Yield Curve Implied by Forward Rates Using Maximum Smoothness Forward Rate Smoothing

We now turn to interest rate simulation using the smoothed yield curves that we have generated in this chapter.

APPENDIX: PROOF OF THE MAXIMUM SMOOTHNESS FORWARD RATE THEOREM

Schwartz (1989) demonstrates that cubic splines produce the maximum smoothness discount functions or yield curves if the spline is applied to discount bond prices or yields respectively. In this appendix, we derive by a similar argument the functional form that produces the forward rate curve with maximum smoothness. Let $f(t)$ be the current forward rate function, so that

$$P(t) = \exp \left(- \int_0^t f(s) ds \right) \quad (5.A1)$$

is the price of a discount bond maturing at time t . The maximum smoothness term structure is a function f with a continuous derivative that satisfies the optimization problem

$$\min \int_0^T f''^2(s) ds \quad (5.A2)$$

subject to the constraints

$$\int_0^{t_i} f(s) ds = -\log P_i, \text{ for } i = 1, 2, \dots, m. \quad (5.A3)$$

Here the $P_i = P(t_i)$, for $i = 1, 2, \dots, m$ are given prices of discount bonds with maturities $0 < t_1 < t_2 < \dots < t_m < T$.

Integrating twice by parts we get the following identity:

$$\int_0^t f(s) ds = \frac{1}{2} \int_0^t (t-s)^2 f''(s) ds + tf(0) + \frac{1}{2} t^2 f'(0) \quad (5.A4)$$

Put

$$g(t) = f''(t), \quad 0 \leq t \leq T \quad (5.A5)$$

and define the step function

$$\begin{aligned} u(t) &= 1 \text{ for } t \geq 0 \\ &= 0 \text{ for } t < 0. \end{aligned}$$

The optimization problem can then be written as

$$\min \int_0^T g^2(s) ds \quad (5.A6)$$

subject to

$$\frac{1}{2} \int_0^T (t_i - s)^2 u(t_i - s) g(s) ds = -\log P_i - t_i f(0) - \frac{1}{2} t_i^2 f'(0) \quad (5.A7)$$

for $i = 1, 2, \dots, m$. Let λ_i for $i = 1, 2, \dots, m$ be the Lagrange multipliers corresponding to the constraints (5.A7). The objective then becomes

$$\begin{aligned} \min Z[g] &= \int_0^T g^2(s) ds \\ &+ \sum_{i=1}^m \lambda_i \left(\frac{1}{2} \int_0^T (t_i - s)^2 u(t_i - s) g(s) ds + \log P_i + t_i f(0) + \frac{1}{2} t_i^2 f'(0) \right) \end{aligned} \quad (5.A8)$$

According to the calculus of variations, if the function g is a solution to equation (5.A8), then

$$\frac{d}{d\varepsilon} Z[g + \varepsilon h]_{\varepsilon=0} = 0 \quad (5.A9)$$

for any function $h(t)$ identically equal to $w''(t)$ where $w(t)$ is any twice differentiable function defined on $[0, T]$ with $w'(0) = w(0) = 0^2$. We get

$$\frac{d}{d\varepsilon} Z[g + \varepsilon h]_{\varepsilon=0} = 2 \int_0^T \left[g(s) + \frac{1}{4} \sum_{i=1}^m \lambda_i (t_i - s)^2 u(t_i - s) \right] h(s) ds$$

In order that this integral is zero for any function h , we must have

$$g(t) + \frac{1}{4} \sum_{i=1}^m \lambda_i (t_i - t)^2 u(t_i - t) = 0 \quad (5.A10)$$

for all t between 0 and T . This means that

$$g(t) = 12e_i t^2 + 6d_i t + 2c_i \quad \text{for } t_{i-1} < t \leq t_i, i = 1, 2, \dots, m+1, \quad (5.A11)$$

where

$$\begin{aligned} e_i &= -\frac{1}{48} \sum_{j=1}^m \lambda_j \\ d_i &= \frac{1}{12} \sum_{j=1}^m \lambda_j t_j \\ c_i &= -\frac{1}{8} \sum_{j=i}^m \lambda_j t_j^2 \end{aligned} \quad (5.A12)$$

and we define $t_0 = 0, t_{m+1} = T$. Moreover, equation (5.A10) implies that g and g' (and therefore f'' and f''') are continuous. From equation (5.A4) we get

$$f(t) = e_i t^4 + d_i t^3 + c_i t^2 + b_i t + a_i, \quad t_{i-1} < t \leq t_i, i = 1, 2, \dots, m+1 \quad (5.A13)$$

Continuity of f, f', f'' and f''' then implies that

$$\begin{aligned} e_i t_i^4 + d_i t_i^3 + c_i t_i^2 + b_i t_i + a_i &= e_{i+1} t_i^4 + d_{i+1} t_i^3 + c_{i+1} t_i^2 + b_{i+1} t_i + a_{i+1}, \\ &= 1, 2, \dots, m \end{aligned} \quad (5.A14)$$

$$\begin{aligned} 4e_i t_i^3 + 3d_i t_i^2 + 2c_i t_i + b_i &= 4e_{i+1} t_i^3 + 3d_{i+1} t_i^2 + 2c_{i+1} t_i + b_{i+1}, i = 1, 2, \dots, m \\ 12e_i t_i^2 + 6d_i t_i + 2c_i &= 12e_{i+1} t_i^2 + 6d_{i+1} t_i + 2c_{i+1} \\ 24e_i t_i + 6d_i &= 24e_{i+1} t_i + 6d_{i+1} \end{aligned} \quad (5.A15)$$

The constraints (5.A3) become

$$\begin{aligned} & \frac{1}{5}e_i(t_i^5 - t_{i-1}^5) + \frac{1}{4}d_i(t_i^4 - t_{i-1}^4) + \frac{1}{3}c_i(t_i^3 - t_{i-1}^3) + \frac{1}{2}b_i(t_i^2 - t_{i-1}^2) + a_i(t_i - t_{i-1}) \\ & = -\log \left[\frac{P_i}{P_{i-1}} \right], i = 1, 2, \dots, m \end{aligned} \quad (5.A16)$$

where we define $P_0 = 1$. This proves the theorem.

NOTES

1. These techniques are reviewed in great detail, with worked examples, in a series of blogs on the Kamakura Corporation website at www.kamakuraco.com:

Basic Building Blocks of Yield Curve Smoothing, Part 1, November 2, 2009. www.kamakuraco.com/Blog/tabid/231/EntryId/150/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-1.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 2: A Menu of Alternatives, November 17, 2009. www.kamakuraco.com/Blog/tabid/231/EntryId/152/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-2-A-Menu-of-Alternatives.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 3: Stepwise Constant Yields and Forwards versus Nelson-Siegel, November 18, 2009. www.kamakuraco.com/Blog/tabid/231/EntryId/156/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-3-Stepwise-Constant-Yields-and-Forwards-versus-Nelson-Siegel.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 4: Linear Yields and Forwards versus Nelson-Siegel, November 20, 2009. www.kamakuraco.com/Blog/tabid/231/EntryId/157/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-4-Linear-Yields-and-Forwards-versus-Nelson-Siegel.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 5: Linear Forward Rates and Related Yields versus Nelson-Siegel, November 30, 2009. www.kamakuraco.com/Blog/tabid/231/EntryId/158/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-5-Linear-Forward-Rates-and-Related-Yields-versus-Nelson-Siegel.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 6: Quadratic Yield Splines and Related Forwards versus Nelson-Siegel, December 3, 2009. www.kamakuraco.com/Blog/tabid/231/EntryId/159/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-6-Quadratic-Yield-Splines-and-Related-Forwards-versus-Nelson-Siegel.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 7: Quadratic Forward Rate Splines and Related Yields versus Nelson-Siegel, December 8, 2009. www.kamakuraco.com/Blog/tabid/231/EntryId/161/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-7-Quadratic-Forward-Rate-Splines-and-Related-Yields-versus-Nelson-Siegel.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 8: Cubic Yield Splines and Related Forwards versus Nelson-Siegel, December 10, 2009. www.kamakuraco.com/Blog/tabid/231/EntryId/162/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-8-Cubic-Yield-Splines-and-Related-Forwards-versus-Nelson-Siegel.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 9: Cubic Forward Rate Splines and Related Yields versus Nelson-Siegel, December 30, 2009. www.kamakuraco.com/Blog/tabid/231/EntryId/167/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-9-Cubic-Forward-Rate-Splines-and-Related-Yields-versus-Nelson-Siegel-Updated-January-6-2010.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 10: Maximum Smoothness Forward Rates and Related Yields versus Nelson-Siegel, January 5, 2010. www.kamakuraco.com/Blog/tabid/231/EntryId/168/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-10-Maximum-Smoothness-Forward-Rates-and-Related-Yields-versus-Nelson-Siegel-Revised-May-8-2012.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 11: The Shimko Test for Measuring Accuracy of Smoothing Techniques,” January 13, 2010. www.kamakuraco.com/Blog/tabid/231/EntryId/170/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-11-The-Shimko-Test-for-Measuring-Accuracy-of-Smoothing-Techniques.aspx.

Basic Building Blocks of Yield Curve Smoothing, Part 12: Smoothing with Bond Prices as Inputs,” January 20, 2010. www.kamakuraco.com/Blog/tabid/231/EntryId/172/Basic-Building-Blocks-of-Yield-Curve-Smoothing-Part-12-Smoothing-with-Bond-Prices-as-Inputs.aspx.

2. www.wikipedia.com has generally been a very useful reference for the topics “spline interpolation” and “spline.” We encourage interested readers to review the current version of Wikipedia on these topics.
3. The authors wish to thank Keith Luna of Western Asset Management Company for helpful comments on this section of Chapter 5.
4. As an example of the volumes of data that can be analyzed, Dickler, Jarrow, and van Deventer present a history of forward rate curves, zero-coupon yield curves, and par coupon bond yield curves for every business day (more than 12,000 business days) from January 2, 1962, to August 22, 2011 in these volumes available on www.kamakuraco.com:

Daniel T. Dickler, Robert A. Jarrow, and Donald R. van Deventer, “Inside the Kamakura Book of Yields: A Pictorial History of 50 Years of U.S. Treasury Forward Rates,” Kamakura Corporation memorandum, September 13, 2011. kamakuraco.com/Portals/0/PDF/KamakuraYieldBook-ForwardRates-FinalVersion-20110913.pdf.

Daniel T. Dickler and Donald R. van Deventer, “Inside the Kamakura Book of Yields: An Analysis of 50 Years of Daily U.S. Treasury Forward Rates,” Kamakura blog, www.kamakuraco.com, September 14, 2011. www.kamakuraco.com/Blog/tabid/231/EntryId/333/Inside-the-Kamakura-Book-of-Yields-An-Analysis-of-50-Years-Of-Daily-U-S-Treasury-Forward-Rates.aspx.

Daniel T. Dickler, Robert A. Jarrow and Donald R. van Deventer, “Inside the Kamakura Book of Yields: A Pictorial History of 50 Years of U.S. Treasury Zero-coupon Bond Yields,” Kamakura Corporation memorandum, September 26, 2011. kamakuraco.com/Portals/0/PDF/KamakuraYieldBook-ZeroCouponBondYields-Final-20110930.pdf.

Daniel T. Dickler and Donald R. van Deventer, “Inside the Kamakura Book of Yields: An Analysis of 50 Years of Daily U.S. Treasury Zero-coupon Bond Yields,” Kamakura blog, www.kamakuraco.com, September 26, 2011. www.kamakuraco.com/Blog/tabid/231/EntryId/337/Inside-the-Kamakura-Book-of-Yields-An-Analysis-of-50-Years-of-Daily-U-S-Treasury-Zero-Coupon-Bond-Yields.aspx.

Daniel T. Dickler, Robert A. Jarrow and Donald R. van Deventer, “Inside the Kamakura Book of Yields: A Pictorial History of 50 Years of U.S. Treasury Par Coupon Bond Yields,” Kamakura Corporation memorandum, October 5, 2011. kamakuraco.com/Portals/0/PDF/KamakuraYieldBook-ParCouponBondYields-Final-20111005.pdf.

Daniel T. Dickler and Donald R. van Deventer, "Inside the Kamakura Book of Yields: An Analysis of 50 Years of Daily U.S. Par Coupon Bond Yields," Kamakura blog, www.kamakuraco.com, October 6, 2011. www.kamakuraco.com/Blog/tabid/231/EntryId/339/Inside-the-Kamakura-Book-of-Yields-An-Analysis-of-50-Years-of-Daily-U-S-Treasury-Par-Coupon-Bond-Yields.aspx.

5. This proof was kindly provided by Oldrich Vasicek. We also appreciate the comments of Volf Frishling, who pointed out an error in the proof in Adams and van Deventer (1994). Robert Jarrow also made important contributions to this proof.
6. We are grateful to Robert Jarrow for his assistance on this point.

Introduction to Heath, Jarrow, and Morton Interest Rate Modeling

In Chapter 5, we mentioned the review by Dickler, Jarrow, and van Deventer of 50 years of daily U.S. Treasury interest rate movements. They concluded that U.S. interest rates are driven by 5 to 10 random factors, a much larger number of factors than major financial institutions typically use for interest rate risk management. In a series of three papers, David Heath, Robert A. Jarrow, and Andrew Morton (1990a, 1990b, 1992) introduced the Heath, Jarrow, and Morton (HJM) framework for modeling interest rates driven by a large number of random factors while preserving the standard “no arbitrage” assumptions of modern finance. The HJM framework is a very powerful general solution for modeling interest rates driven by a large number of factors and for the valuation and simulation of cash flows on securities of all types. The authors believe that an understanding of the HJM approach is a fundamental requirement for a risk manager. In this chapter and the following three chapters, we lay out four worked examples of the use of the HJM approach with increasing realism. A full enterprise-wide risk management infrastructure relies on a sophisticated implementation of the HJM framework with the 5 to 10 driving risk factors and a very large number of time steps. Such an implementation should make use of both the HJM implications for Monte Carlo simulation and the “bushy tree” approach for valuing securities like callable bonds and home mortgages. For expositional purposes, this and the following chapters focus on the bushy tree approach, but we explain why an HJM Monte Carlo simulation is essential for the simulation of cash flows, net income, and liquidity risk in Chapter 10.

The authors wish to thank Robert A. Jarrow for his encouragement and advice on this series of worked examples of the HJM approach. What follows is based heavily on Jarrow’s classic book, *Modeling Fixed Income Securities and Interest Rate Options*, 2nd ed. (2002), particularly Chapters 4, 6, 8, 9, and 15. In our Chapter 9, we incorporate some modifications of Chapter 15 in Jarrow’s book that would have been impossible without Jarrow’s advice and support.

In Chapters 6 through 9, we use data from the Federal Reserve statistical release H15, published on April 1, 2011. U.S. Treasury yield curve data was smoothed using Kamakura Corporation’s Kamakura Risk Manager version 7.3 to create zero-coupon bonds via the maximum smoothness technique of Adams and van Deventer as described in Chapter 5. Most of the applications of the HJM approach by financial market participants have come in the form of the LIBOR market model and




EXHIBIT 6.1 Kamakura Corporation 10-Year Forecast of U.S. Treasury Yield Curve Implied by Forward Rates Using Maximum Smoothness Forward Rate Smoothing

Sources: Kamakura Corporation; Federal Reserve.

applications related to the interest rate swap market and deposits issued at the London interbank offered rate. In light of recent lawsuits alleging manipulation of LIBOR rates, this application of HJM is a complex special case that we will defer to Chapter 17. In this chapter, we assume away such complexities and assume away the potential default of the U.S. government. In this chapter, as in the original HJM publications, it is assumed that the underlying yield curve is “risk free” from a credit risk perspective.

The smoothed U.S. Treasury yield curve and the implied forward yield curve monthly for 10 years looked like the one in Exhibit 6.1 as of the March 31, 2011, data reported by the Federal Reserve on April 1, 2011.

The continuous forward rate curve and zero-coupon bond yield curve that prevailed as of the close of business on March 31, 2011, were as shown in Exhibit 6.2, with the forward rate curve being the higher of the two lines.

OBJECTIVES OF THE EXAMPLE AND KEY INPUT DATA

Following Jarrow (2002), we want to show how to use the HJM framework to model random interest rates and value fixed income securities and other securities using the simulated random rates. For all of the examples in Chapters 6 through 9, we assume the following:




EXHIBIT 6.2 U.S. Treasury Forward Rates and Zero-Coupon Yields Derived from the Federal Reserve H15 Statistical Release Using Maximum Smoothness Forward Rate Smoothing

Sources: Kamakura Corporation; Federal Reserve.

- Zero-coupon bond prices for the U.S. Treasury curve on March 31, 2011, are the basic inputs.
- Interest rate volatility assumptions are based on the Dickler, Jarrow, and van Deventer papers cited previously on daily U.S. Treasury yields and forward rates from 1962 to 2011.
- The modeling period is four equal length periods of one year each.
- The HJM implementation is that of a “bushy tree,” which we describe below.

The HJM framework is usually implemented using Monte Carlo simulation or a bushy tree approach where a lattice of interest rates and forward rates is constructed. This lattice, in the general case, does not recombine like the popular binomial or trinomial trees used to replicate Black-Scholes options valuation or simple legacy one-factor term structure models that we discuss in Chapter 13. In general, the bushy tree does not recombine because the interest rate volatility assumptions imply a path-dependent interest rate model, not one that is path independent like the simplest one-factor term structure model implementations. Instead, the bushy tree of upshifts and downshifts in interest rates looks like Exhibit 6.3.

At each of the points in time on the lattice (times 0, 1, 2, 3, and 4), there are sets of zero-coupon bond prices and forward rates. At time 0, there is one set of data. At time one, there are two sets of data, the “up-set” and “down-set.” At time two, there are four sets of data (up up, up down, down up, and down down), and at time three there are $8 = 2^3$ sets of data.

In order to build a bushy tree like this, we need to specify how many factors we are modeling and what the interest volatility assumptions are. The number of ending nodes on a bushy tree for n periods is 2^n for a 1 factor model, 3^n for a 2 factor model, and 4^n for a 3 factor model. For this example, we build a one-factor bushy tree.




EXHIBIT 6.3 Example of Bush Tree for HJM Modeling of No-Arbitrage, Zero-Coupon Bond Price Movements

“One factor” implies some important unrealistic restrictions on simulated yield curve movements, which we discuss below and in Chapters 7 through 9 and 13.

The simplest interest rate volatility function one can assume is that the one-year spot U.S. Treasury rate and the three one-year forward rates (the one-year forwards maturing at times 2, 3, and 4) all have an identical interest rate volatility σ . This is the discrete time equivalent of the Ho and Lee model (1986). We start by examining the actual standard deviations of the daily changes from 1962 to 2011 for 10 interest rates: The continuously compounded one-year, zero-coupon bond yield, the continuously compounded one-year forward rate maturing in year 2, the continuously compounded one-year forward rate maturity in year 3, and so on for each one-year forward rate out to the one maturing in year 10. The daily volatilities are as shown in Exhibit 6.4.

The one-year spot rate’s volatility was 0.0832 percent over this period. The volatilities for the forward rates maturing in years 2, 3, and 4 were 0.0911 percent, 0.0981 percent, and 0.0912 percent, respectively. We can reach two conclusions immediately:

1. Interest rate volatility is not constant over the maturity horizon (contrary to the Ho and Lee 1986 assumptions).
2. Interest rate volatility is not declining over the maturity horizon (contrary to the Vasicek 1977 assumptions).

EXHIBIT 6.4 Analysis of Daily Changes in Spot and One-Year Forward Rates

	Maturity in Years									
	1	2	3	4	5	6	7	8	9	10
Standard Deviation	0.0832%	0.0911%	0.0981%	0.0912%	0.1065%	0.0989%	0.1160%	0.0954%	0.1018%	0.1524%
Maximum Daily Change	1.0124%	1.0011%	1.0470%	0.7269%	0.8183%	1.0634%	1.6645%	1.0765%	1.0964%	1.7287%
Maximum Daily Change	-1.0012%	-0.9360%	-0.9856%	-0.9428%	-1.2709%	-1.3032%	-1.0771%	-0.8747%	-1.0972%	-2.3316%
Width of Range	2.0135%	1.9371%	2.0325%	1.6696%	2.1092%	2.3665%	2.7416%	1.9512%	2.1936%	4.0603%

The HJM approach is so general that the Ho and Lee and Vasicek models can be replicated as special cases. We discuss that replication in Chapter 13. There is no need to limit ourselves to such unrealistic volatility assumptions, however, because of the power of the HJM approach.

Instead, we assume that forward rates have their actual volatilities over the period from 1962–2011, but we need to adjust for the fact that we have measured volatility over a time period of one day (which we assume is $1/250$ of a year based on 250 business days per year) and we need a period for which the length of the period Δ is 1, not $1/250$. We know that the variance changes as the length of the time period:

$$\text{Measured variance} = \Delta\sigma^2$$

and, therefore,

$$[\text{Measured volatility}] \frac{1}{\sqrt{\Delta}} = \sigma$$

Therefore, the one-year forward rate, which has a volatility of 0.00091149 over a one-day period has an estimated annual volatility of $0.00091149/[(1/250)^{1/2}] = 0.01441187$. We revisit this method of adjustment for changes in the time period repeatedly in this book. It's an adjustment that should be avoided whenever possible.

For convenience, we use the annual volatilities from this chart in our HJM implementation:

Interest Rate Volatility Assumptions for an Annual Period Model

	Daily	Annual
1-year forward rate maturing in year 2	0.00091149	0.01441187
1-year forward rate maturing in year 3	0.00098070	0.01550626
1-year forward rate maturing in year 4	0.00091243	0.01442681

We will use the zero-coupon bond prices prevailing on March 31, 2011, as our other inputs:

Number of periods:	4
Length of periods (years)	1
Number of risk factors:	1
Volatility term structure:	Empirical

Period Start	Period End	Parameter Inputs	
		Zero-Coupon Bond Prices	Forward Rate Volatility
0	1	0.99700579	
1	2	0.98411015	0.01441187
2	3	0.96189224	0.01550626
3	4	0.93085510	0.01442681

Our final assumption is that there is one random factor driving the forward rate curve. This implies (since all of our volatilities are positive) that all forward rates move up or down together. This implication of the model, as we will see in Chapters 7 through 9, is grossly unrealistic. We relax this assumption when we move to two- and three-factor examples.

KEY IMPLICATIONS AND NOTATION OF THE HJM APPROACH

The HJM conclusions are very complex to derive but their implications are very straightforward. Once the zero-coupon bond prices and volatility assumptions are made, the mean of the distribution of forward rates (in a Monte Carlo simulation) and the structure of a bushy tree are completely determined by the constraints that there be no arbitrage in the economy. Modelers who are unaware of this insight would choose means of the distributions for forward rates such that Monte Carlo or bushy tree valuation would provide different prices for the zero-coupon bonds on March 31, 2011, than those used as input. This would create the appearance of an arbitrage opportunity, but it is in fact a big error that calls into question the validity of the calculation, as it should. Even more than two decades after the publication of the HJM conclusions, financial market participants perform risk analysis that contains such errors.

We show in this example that the zero-coupon bond valuations using a bushy tree are 100 percent consistent with the inputs. We now introduce our notation:

Δ	= length of time period, which is 1 in this example
$r(t, s_t)$	= the simple one-period, uncompounded risk-free interest rate as of current time t in state s_t
$R(t, s_t)$	= $1 + r(t, s_t)$, the total return, the value of \$1 dollar invested at the risk-free rate for one period with no compounding
$\sigma(t, T, s_t)$	= forward rate volatility at time t for forward maturing at T in state s (a sequence of ups and downs)
$P(t, T, s_t)$	= zero-coupon bond price at time t maturing at time T in state s_t (i.e., up or down)
Index	= 1 if the state is up and = -1 if the state is down
$U(t, T, s_t)$	= the total return that leads to bond price $P(t + \Delta, T, s_t = \text{up})$ on a T maturity bond in an upshift state one period from current time t
$D(t, T, s_t)$	= the total return that leads to bond price $P(t + \Delta, T, s_t = \text{down})$ on a T maturity bond in a downshift state one period from current time t
$K(t, T, s_t)$	= the sum of the forward volatilities for the one-period forward rates from $t + \Delta$ to $T - \Delta$, as shown here:

$$K(t, T, s_t) = \sum_{j=t+\Delta}^{T-\Delta} \sigma(t, j, s_t) \sqrt{\Delta}$$

We also see the rare appearance of a trigonometric function in finance, one found in common spreadsheet software:

$$\text{Cosh}(x) = \frac{1}{2}[e^x + e^{-x}]$$

Note that the current times that will be relevant in building a bushy tree of zero-coupon bond prices are current times $t = 0, 1, 2$, and 3 . We'll be interested in maturity dates $T = 2, 3$, and 4 . We know that at time zero, there are four zero-coupon bonds outstanding. At time 1, only the bonds maturity at $T = 2, 3$, and 4 will remain outstanding. At time 2, only the bonds maturing at times $T = 3$ and 4 will remain, and at time 3 only the bond maturing at time 4 will remain. For each of the boxes below, we need to fill in the relevant bushy tree (one for each of the four zero-coupon bonds) with each upshift and downshift of the zero-coupon bond price as we step forward one more period (by $\Delta = 1$) on the tree. In the interests of saving space, we arrange the tree to look like a table by stretching the bushy tree as shown in Exhibit 6.5.

A completely populated zero-coupon bond price tree, in this case for the zero coupon bond maturing at time $T = 4$, would then be summarized like the one in Exhibit 6.6, where the final shift in price is shaded for up- or downshifts.

The mapping of the sequence of up- and downshifts is shown in Exhibit 6.7, consistent with the stretched tree in Exhibit 6.6.

In order to populate the trees with zero-coupon bond prices and forward rates, there is one more piece of information that we need to supply.




EXHIBIT 6.5 Example of Bushy Tree for HJM Modeling of No-Arbitrage, Zero-Coupon Bond Price Movements

EXHIBIT 6.6 Four-Year, Zero-Coupon Bond

Row Number	Current Time			
	0	1	2	3
1	0.930855	0.975026	1.002895	1.010619
2		0.892275	0.944645	0.981905
3			0.944617	0.979758
4			0.889752	0.951921
5				0.981876
6				0.953979
7				0.951893
8				0.924847

EXHIBIT 6.7 Map of Sequence of States

Row Number	Current Time			
	0	1	2	3
1	Time 0	up	up-up	up-up-up
2		dn	up-dn	up-up-dn
3			dn-up	up-dn-up
4			dn-dn	up-dn-dn
5				dn-up-up
6				dn-up-dn
7				dn-dn-up
8				dn-dn-dn

PSEUDO-PROBABILITIES

In Chapter 7 of Jarrow (2002), he shows that a necessary and sufficient condition for no arbitrage is that, at every node in the tree, the one-period return on a zero-coupon bond neither dominates nor is dominated by a one-period investment in the risk-free rate. Written as the total return on a \$1 investment, an upshift U must be greater than the risk-free total return and the downshift D must be less than the risk-free total return R :

$$U(t, T, s_t) > R(t, T, s_t) > D(t, T, s_t)$$

for all points in time t and, states s_t .

Jarrow adds (Jarrow, 2002, p. 126) “This condition is equivalent to the following statement: There exists a unique number $\pi(t, s_t)$ strictly between 0 and 1 such that

$$R(t, s_t) = \pi(t, s_t)U(t, T, s_t) + (1 - \pi(t, s_t))D(t, T, s_t)$$

Solving for $\pi(t, s_t)$ gives us the relationship

$$\pi(t, s_t) = \frac{[R(t, s_t) - D(t, T, s_t)]}{[U(t, T, s_t) - D(t, T, s_t)]}$$

If the computed $\pi(t, s_t)$ values are between 0 and 1 everywhere on the bushy tree, then the tree is arbitrage free.”

Jarrow continues as follows (Jarrow, p. 127):

Note that each $\pi(t, s_t)$ can be interpreted as a pseudo-probability of the up state occurring over the time interval $(t, t + \Delta)$. . . We call these $\pi(t, s_t)$ s pseudo-probabilities because they are “false” probabilities. They are not the actual probabilities generating the evolution of the money market account and the $[T]$ period zero-coupon bond prices. Nonetheless, these pseudo-probabilities have an important use [in valuation].

Jarrow goes on to explain that risk-neutral valuation is computed by “taking the expected cash flow, using the pseudo-probabilities, and discounting at the spot rate of interest.” He adds “this is called risk-neutral valuation because it is the value that the random cash flow ‘x’ would have in an economy populated by risk-neutral investors, having the pseudo-probabilities as their beliefs”(Jarrow, p. 129).

We now demonstrate how to construct the bushy tree and use it for risk-neutral valuation.

THE FORMULA FOR ZERO-COUPON BOND PRICE SHIFTS

Jarrow demonstrates that the values for the risk-neutral probability of an upshift $\pi(t, s_t)$ are determined by a set of mathematical limit conditions for an assumed evolution of interest rates. These limit conditions allow multiple solutions. Without loss of generality, one can always construct a tree whose limit is the assumed interest rate evolution with $\pi(t, s_t) = 1/2$ for all points in current time t and all up and down states s_t . Then, using our notation above and remembering that the variable “Index” is 1 in an upshift and -1 in a downshift, the shifts in zero-coupon bond prices can be written as follows:

$$P(t + \Delta, T, s_{t+\Delta}) = \frac{[P(t, T, s_t)R(t, s_t)e^{K(t, T, s_t)\Delta(Index)}]}{[\cosh(K(t, T, s_t)\Delta)]} \quad (6.1)$$

where for notational convenience $\cosh[K(t, T, s_t)\Delta] \equiv 1$ when $T - \Delta < t + \Delta$. This is equation 15.17 in Jarrow (2002, 286). We now put this formula to use.

BUILDING THE BUSHY TREE FOR ZERO-COUPON BONDS MATURING AT TIME $T = 2$

We now populate the bushy tree for the two-year, zero-coupon bond. We calculate each element of equation (6.1). When $t = 0$ and $T = 2$, we know $\Delta = 1$ and

$$P(0, 2, s_t) = 0.98411015$$

The one-period, risk-free return (remember, this is one plus the risk-free rate) is

$$R(0, s_t) = 1/P(0, 1, s_t) = 1/0.997005786 = 1.003003206$$

The scaled sum of sigmas $K(t, T, s_t)$ for $t = 0$ and $T = 2$ becomes

$$K(t, T, s_t) = \sum_{j=t+\Delta}^{T-\Delta} \sigma(t, j, s_t) \sqrt{\Delta} = \sum_{j=1}^1 \sigma(0, j, s_t) \sqrt{\Delta}$$

and, therefore, $K(0, T, s_t) = (\sqrt{1})(0.01441187) = 0.01441187$

Using equation (6.1) with these inputs and the fact that the variable Index = 1 for an upshift gives

$$P(1, 2, s_t = \text{up}) = 1.001290$$

For a downshift we set Index = -1 and recalculate equation (6.1) to get

$$P(1, 2, s_t = \text{down}) = 0.972841$$

We have fully populated the bushy tree for the zero-coupon bond maturing at $T = 2$ (note values have been rounded to six decimal places for display only), since all of the up and down states at time $t = 2$ result in a riskless payoff of the zero-coupon bond at its face value of 1 (Exhibit 6.8). Note also that, as a result of our volatility assumptions, the price of the bond maturing at time $T = 2$ after an upshift to time $t = 1$ has resulted in a value higher than the par value of 1. This implies negative interest rates. Negative nominal interest rates have appeared in Switzerland, Japan, Hong Kong, and (selectively) in the United States. The academic argument is that this cannot occur, because one could “costlessly” store cash in order to avoid a negative interest rate. The assumption that cash can be stored costlessly is an

EXHIBIT 6.8 Two-Year, Zero-Coupon Bond

Row Number	Current Time			
	0	1	2	3
1	0.98411	1.00129	1	
2		0.972841	1	
3			1	
4			1	
5				
6				
7				
8				

assumption that is not true. In fact, on August 5, 2011, Bank of New York Mellon announced that it would charge a fee (i.e., creating a negative return) on demand deposit balances in excess of \$50 million), according to the *Wall Street Journal*. We will return to this issue in later chapters.

BUILDING THE BUSHY TREE FOR ZERO-COUPON BONDS MATURING AT TIME $T = 4$

We now populate the bushy tree for the four-year, zero-coupon bond (see Exhibit 6.9). We calculate each element of equation (6.1). When $t = 0$ and $T = 4$, we know $\Delta = 1$ and

$$P(0, 4, s_t) = 0.930855099$$

The one-period, risk-free return (one plus the risk-free rate) is

$$R(0, s_t) = 1/P(0, 1, s_t) = 1/0.997005786 = 1.003003206$$

The scaled sum of sigmas $K(t, T, s_t)$ for $t = 0$ and $T = 4$ becomes

$$K(t, T, s_t) = \sum_{j=t+\Delta}^{T-\Delta} \sigma(t, j, s_t) \sqrt{\Delta} = \sum_{j=1}^3 \sigma(0, j, s_t) \sqrt{\Delta}$$

and, therefore,

$$K(0, T, s_t) = (\sqrt{1})(0.01441187 + 0.015506259 + 0.01442681) = 0.044344939.$$

Using equation (6.1) with these inputs and the fact that the variable Index = 1 for an upshift gives

EXHIBIT 6.9 Four-Year, Zero-Coupon Bond

Row Number	Current Time			
	0	1	2	3
1	0.930855	0.975026	1.002895	1.010619
2		0.892275	0.944645	0.981905
3			0.944617	0.979758
4			0.889752	0.951921
5				0.981876
6				0.953979
7				0.951893
8				0.924847

$$P(1, 4, s_t = \text{up}) = 0.975026$$

For a downshift, we set Index = -1 and recalculate equation (6.1) to get

$$P(1, 4, s_t = \text{down}) = 0.892275$$

We have correctly populated the first two columns of the bushy tree for the zero-coupon bond maturing at $T = 4$ (note values have been rounded to six decimal places for display only).

Now we move to the third column, which displays the outcome of the $T = 4$ zero-coupon bond price after four scenarios: up-up, up-down, down-up, and down-down. We calculate $P(2, 4, s_t = \text{up})$ and $P(2, 4, s_t = \text{down})$ after the initial downstate as follows. When $t = 1$, $T = 4$, and $\Delta = 1$ then

$$P(1, 4, s_t = \text{down}) = 0.892275$$

as shown in the second row of the second column.

The one-period risk-free return (one plus the risk-free rate) is

$$R(1, s_t = \text{down}) = 1/P(1, 2, s_t = \text{down}) = 1/0.972841 = 1.027917$$

The scaled sum of sigmas $K(t, T, s_t)$ for $t = 1$ and $T = 4$ becomes

$$K(t, T, s_t) = \sum_{j=t+\Delta}^{T-\Delta} \sigma(t, j, s_t) \sqrt{\Delta} = \sum_{j=2}^3 \sigma(0, j, s_t) \sqrt{\Delta}$$

and, therefore,

$$K(1, 4, s_t) = (\sqrt{1})(0.01441187 + 0.015506259) = 0.029918$$

Using equation (6.1) with these inputs and the fact that the variable Index = 1 for an upshift gives

$$P(2, 4, s_t = \text{up}) = 0.944617$$

For a downshift, we set Index = -1 and recalculate equation (6.1) to get

$$P(2, 4, s_t = \text{down}) = 0.889752$$

We have correctly populated the third and fourth rows of column 3 ($t = 2$) of the bushy tree for the zero-coupon bond maturing at $T = 4$ (note values have been rounded to six decimal places for display only) shown previously in Exhibit 6.9. The remaining calculations are left to the reader.

In a similar way, the bushy tree for the zero-coupon bond price maturing at time $T = 3$ can be calculated as shown in Exhibit 6.10.

If we combine all of these tables, we can create a table of the term structure of zero-coupon bond prices in each scenario as in Exhibit 6.11.

At any point in time t , the continuously compounded yield to maturity at time T can be calculated as $y(T - t) = -\ln[P(t, T)]/(T - t)$. When we display the term structure of zero-coupon yields in each scenario, we can see that our volatility assumptions have indeed produced some implied negative yields in this particular simulation (Exhibit 6.12). We discuss how to remedy that problem in the next chapter.

Finally, in Exhibit 6.13 we can display the one-year U.S. Treasury spot rates and the associated term structure of one-year forward rates in each scenario.

EXHIBIT 6.10 Three-Year, Zero-Coupon Bond

Row Number	Current Time			
	0	1	2	3
1	0.961892	0.993637	1.006657	1
2		0.935925	0.978056	1
3			0.975917	1
4			0.948189	1
5				1
6				1
7				1
8				1

EXHIBIT 6.11 Zero-Coupon Bond Prices

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.997006	0.98411	0.961892	0.930855
Time 1 up	1	1.00129	0.993637	0.975026	
Time 1 down	2	0.972841	0.935925	0.892275	
Time 2 up-up	3	1.006657	1.002895		
Time 2 up-down	4	0.978056	0.944645		
Time 2 down-up	5	0.975917	0.944617		
Time 2 down-down	6	0.948189	0.889752		
Time 3 up-up-up	7	1.010619			
Time 3 up-up-down	8	0.981905			
Time 3 up-down-up	9	0.979758			
Time 3 up-down-down	10	0.951921			
Time 3 down-up-up	11	0.981876			
Time 3 down-up-down	12	0.953979			
Time 3 down-down-up	13	0.951893			
Time 3 down-down-down	14	0.924847			

EXHIBIT 6.12 Continuously Compounded Zero Yields

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.2999%	0.8009%	1.2951%	1.7913%
Time 1 up	1	−0.1289%	0.3192%	0.8430%	
Time 1 down	2	2.7534%	3.3110%	3.7994%	
Time 2 up-up	3	−0.6635%	−0.1445%		
Time 2 up-down	4	2.2188%	2.8473%		
Time 2 down-up	5	2.4377%	2.8488%		
Time 2 down-down	6	5.3201%	5.8406%		
Time 3 up-up-up	7	−1.0563%			
Time 3 up-up-down	8	1.8261%			
Time 3 up-down-up	9	2.0449%			
Time 3 up-down-down	10	4.9273%			
Time 3 down-up-up	11	1.8290%			
Time 3 down-up-down	12	4.7114%			
Time 3 down-down-up	13	4.9303%			
Time 3 down-down-down	14	7.8127%			

EXHIBIT 6.13 One-Year Forward Rates, Simple Interest

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.3003%	1.3104%	2.3098%	3.3343%
Time 1 up	1	−0.1288%	0.7702%	1.9087%	
Time 1 down	2	2.7917%	3.9443%	4.8920%	
Time 2 up-up	3	−0.6613%	0.3752%		
Time 2 up-down	4	2.2436%	3.5368%		
Time 2 down-up	5	2.4677%	3.3136%		
Time 2 down-down	6	5.4642%	6.5678%		
Time 3 up-up-up	7	−1.0508%			
Time 3 up-up-down	8	1.8428%			
Time 3 up-down-up	9	2.0660%			
Time 3 up-down-down	10	5.0507%			
Time 3 down-up-up	11	1.8459%			
Time 3 down-up-down	12	4.8242%			
Time 3 down-down-up	13	5.0539%			
Time 3 down-down-down	14	8.1260%			

VALUATION IN THE HJM FRAMEWORK

Jarrow in a previous quote described valuation as the expected value of cash flows using the risk-neutral probabilities. Note that column 1 in the exhibit above denotes the riskless one-period interest rate in each scenario. For the scenario number 14

(three consecutive downshifts in zero-coupon bond prices), cash flows at time $T = 4$ would be discounted by the one-year spot rates at time $t = 0$, by the one-year spot rate at time $t = 1$ in scenario 2 ("down"), by the one-year spot rate in scenario 6 (down down) at time $t = 2$, and by the one-year spot rate at time $t = 3$ in scenario 14 (down down down). The discount factor is

$$\text{Discount Factor}(0, 4, \text{down down down}) = 1 / (1.003003)(1.027917)(1.054642) \times (1.081260)$$

These discount factors are displayed in Exhibit 6.14 for each potential cash flow date.

When taking expected values, we can calculate the probability of each scenario coming about since the probability of an upshift is one-half (see Exhibit 6.15).

It is convenient to calculate the probability-weighted discount factors for use in calculating the expected present value of cash flows (Exhibit 6.16).

We now use the HJM bushy trees we have generated to value representative securities.

EXHIBIT 6.14 Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1	0.997006	0.998292	1.004938	1.01561
2		0.997006	0.998292	1.004938	0.986754
3			0.969928	0.976385	0.956622
4			0.969928	0.976385	0.929442
5				0.94657	0.929414
6				0.94657	0.903007
7				0.919676	0.875433
8				0.919676	0.85056

EXHIBIT 6.15 Probability of Each State

Row Number	Current Time				
	0	1	2	3	4
1	100.00%	50.00%	25.00%	12.50%	12.50%
2		50.00%	25.00%	12.50%	12.50%
3			25.00%	12.50%	12.50%
4			25.00%	12.50%	12.50%
5				12.50%	12.50%
6				12.50%	12.50%
7				12.50%	12.50%
8				12.50%	12.50%

EXHIBIT 6.16 Probability-Weighted Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1	0.498503	0.249573	0.125617	0.126951
2		0.498503	0.249573	0.125617	0.123344
3			0.242482	0.122048	0.119578
4			0.242482	0.122048	0.11618
5				0.118321	0.116177
6				0.118321	0.112876
7				0.114959	0.109429
8				0.114959	0.10632

EXHIBIT 6.17 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0	0	0	0	1
2		0	0	0	1
3			0	0	1
4			0	0	1
5				0	1
6				0	1
7				0	1
8				0	1
Risk-Neutral Value =					0.93085510

VALUATION OF A ZERO-COUPON BOND MATURING AT TIME $T = 4$

A riskless zero-coupon bond pays \$1 in each of the eight nodes of the bushy tree that prevail at time $T = 4$ (Exhibit 6.17).

When we multiply this vector of 1s times the probability-weighted discount factors in the time $T = 4$ column in the previous table and add them, we get a zero-coupon bond price of 0.93085510, which is the value we should get in a no-arbitrage economy, the value observable in the market and used as an input to create the tree. While this is a very important confirmation of the accuracy of our interest rate simulation, very few interest rate analysts take this important step in the “model audit” process. The valuation of security whose net present value (price plus accrued interest) is used as an input should replicate the input value. If this is not the case, the calculation is wrong.

EXHIBIT 6.18 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0	3	3	3	103
2		3	3	3	103
3			3	3	103
4			3	3	103
5				3	103
6				3	103
7				3	103
8				3	103
Risk-Neutral Value =					104.70709974

VALUATION OF A COUPON-BEARING BOND PAYING ANNUAL INTEREST

Next we value a bond with no credit risk that pays \$3 in interest at every scenario at times $T = 1, 2, 3$, and 4 plus principal of 100 at time $T = 4$. The valuation is calculated by multiplying each cash flow by the matching probability weighted discount factor, to get a net present value (which is price plus accrued interest) of 104.70709974 (Exhibit 6.18).

VALUATION OF A DIGITAL OPTION ON THE ONE-YEAR U.S. TREASURY RATE

Now we value a digital option that pays \$1 at time $T = 3$ if (at that time) the one-year U.S. Treasury rate (for maturity at $T = 4$) is over 8 percent. If we look at the table of the term structure of one-year spot and forward rates above, this happens in only one scenario, scenario 14 (down down down). The cash flow can be input in the table in Exhibit 6.19 and multiplied by the probability-weighted discount factors to find that this option has a value of 0.11495946:

CONCLUSION

The Dickler, Jarrow, and van Deventer studies of movements in U.S. Treasury yields and forward rates from 1962 to 2011 confirm that 5 to 10 factors are needed to accurately model interest rate movements. Popular one-factor models (Ho and Lee; Vasicek; Hull and White; Black, Derman, and Toy) cannot replicate the actual movements in yields that have occurred. The interest rate volatility assumptions in these models (constant, constant proportion, declining, etc.) are also inconsistent with observed volatility.

EXHIBIT 6.19 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0	0	0	0	0
2		0	0	0	0
3			0	0	0
4			0	0	0
5				0	0
6				0	0
7				0	0
8				1	0
Risk-Neutral Value =					0.11495946

In order to handle a large number of driving factors and complex interest rate volatility structures, the HJM framework is necessary. This chapter shows how to simulate zero-coupon bond prices, forward rates, and zero-coupon bond yields in an HJM framework with maturity-dependent interest rate volatility. Monte Carlo simulation, an alternative to the bushy tree framework, can be done in a fully consistent way.

In the next chapter, we enrich the realism of the simulation by allowing for interest rate volatility that is also a function of the level of interest rates.

HJM Interest Rate Modeling with Rate and Maturity-Dependent Volatility

In Chapter 6, we provided a worked example of interest rate modeling and valuation using the Heath, Jarrow, and Morton (HJM) framework and the assumption that forward rate volatility was dependent on years to maturity using actual U.S. Treasury yields prevailing on March 31, 2011, and historical volatility from 1962 to 2011. In this chapter, we increase the realism of the forward rate volatility assumption. Actual data makes it very clear that forward rate volatility depends not just on the time to maturity of the forward rate but also on the level of spot interest rates. This chapter shows how to incorporate that assumption into the analysis for enhanced accuracy and realism. This is a much more general assumption than that used by Ho and Lee (1986) for constant volatility or by Vasicek (1977) and Hull and White (1990b) for declining volatility.

We again use data from the Federal Reserve statistical release H15 published on April 1, 2011, for yields prevailing on March 31, 2011. U.S. Treasury yield curve data was smoothed using Kamakura Risk Manager version 7.3 to create zero-coupon bonds via the maximum smoothness forward rate technique of Adams and van Deventer discussed in Chapter 5.

OBJECTIVES OF THE EXAMPLE AND KEY INPUT DATA

Following Jarrow (2002), we make the same modeling assumptions for our worked example here as in Chapter 6:

- Zero-coupon bond prices for the U.S. Treasury curve on March 31, 2011 are the basic inputs.
- Interest rate volatility assumptions are based on the Dickler, Jarrow, and van Deventer series on daily U.S. Treasury yields and forward rates from 1962 to 2011. In this chapter, we make the volatility assumptions more accurate than in Chapter 6.
- The modeling period is four equal length periods of one year each.
- The HJM implementation is that of a bushy tree, which we describe below.

The bushy tree looks the same as in Chapter 6, but the modification we make in volatility assumptions will change the “height” of the branches in the tree (Exhibit 7.1).

At each of the points in time on the lattice (time 0, 1, 2, 3 and 4), there are again sets of zero-coupon bond prices and forward rates. At time 0, there is one set of data. At time one, there are two sets of data, the “up-set” and “down-set.” At time two, there are four sets of data (up up, up down, down up, and down down), and at time three there are $8 = 2^3$ data sets.

In Chapter 6, we calculated the standard deviation of the daily changes in three one-year forward rates (the one-year U.S. Treasury rates maturing at time $T = 2, 3$, and 4) over the period from 1962–2011. We then adjusted for the fact that we have measured volatility over a time period of one day (which we assume is $1/250$ of a year based on 250 business days per year) and we need a volatility for which the length of the period Δ is 1, not $1/250$. The resulting yield curve shifts implied by the HJM approach resulted in negative interest rates in three of the scenarios. In this chapter, we reexamine our volatility assumptions to see whether volatility is dependent on the level of interest rates as well as on maturity. We hypothesize that interest rate volatility is dependent on the starting level of the spot one-year U.S. Treasury yield (continuously compounded) and the maturity of the forward rate (the first one-year forward rate matures at time $t + 2$, the second one-year forward rate matures at time $t + 3$, and so on). We categorize the data by interest rate levels, grouping data for spot one-year rates under 0.09 percent, between 0.09 percent and 0.25 percent, between 0.25 percent and 0.50 percent, between 0.50 percent and 0.75 percent, between 0.75 percent and 1.00 percent, and then in one percent intervals to 10




EXHIBIT 7.1 Example of Bushy Tree for HJM Modeling of No-Arbitrage, Zero-Coupon Bond Price Movements




EXHIBIT 7.2 Standard Deviation of Daily Changes in One-Year U.S. Treasury Forward Rates Maturing in Years 2, 3, and 4, 1962–2011

Sources: Kamakura Corporation; Federal Reserve.

percent. The final group of data includes all of those observations in which the starting spot one-year U.S. Treasury yield was over 10 percent.

The graph in Exhibit 7.2 shows the need for incorporating the level of interest rates in our selection of forward rate volatility. We plot the actual daily volatility of one-year forward rates maturing in years 2, 3, and 4 by interest rate groups, which are determined by the highest level of starting one-year U.S. Treasury yields in that group.

The graph makes it obvious that interest rate volatility depends, in a complex but smooth way, on both forward rate maturity and the starting level of the one-year U.S. Treasury yield.

As suggested by Kamakura Senior Research Fellow Sean Patrick Klein, we then annualized these daily volatilities as we did in Chapter 6 and compared them to volatilities of the same forward rates calculated based on one-year changes (not annualized daily changes) in the forward rate from 1963 to 2011. The results will surprise many. We start by plotting the forward rate volatilities using annual time intervals and find a striking difference in how volatilities change as rates rise (Exhibit 7.3).

Forward rate volatility rises in a smooth but complex way as the level of interest rates rises. When we plot the volatilities calculated based on annual time intervals against the annualized volatilities derived from daily time intervals, the results are very surprising (Exhibit 7.4).

We used a very standard financial methodology for annualizing the daily volatilities, but the result substantially overestimated the actual forward rate annual volatilities for forward rates when the one-year U.S. Treasury was 3.00 percent or below. We conclude that this contributed strongly to the fact that our first worked example produced negative interest rates. The results also confirm that the level of interest rates does impact the level of forward rate volatility, but in a more complex




EXHIBIT 7.3 Standard Deviation of Annual Changes in One-Year U.S. Treasury Forward Rates Maturing in Years 2, 3, and 4, 1962–2011
Sources: Kamakura Corporation; Federal Reserve.




EXHIBIT 7.4 Comparison of Annualized Standard Deviation of Daily Changes in One-Year Spot and Forward Rates with Standard Deviation of Annual Changes U.S. Treasury Continuously Compounded Spot and Forward Rates, 1962–2011
Sources: Kamakura Corporation; Federal Reserve.

way than the lognormal interest rate movements assumed by Fischer Black with Derman and Toy (1990) and with Karasinski (1991).

In this chapter, we use the exact volatilities for the one-year changes in continuously compounded forward rates as shown in this table. There were no observations at the time this data was compiled for starting one-year U.S. Treasury yields below 0.002499 percent, so we have set those volatilities by assumption. With the passage of time, these assumptions can be replaced with facts or with data from other low-rate countries like Japan. The volatilities used were selected in lookup table fashion based on U.S. Treasury spot rate levels at that node on the bushy tree from this table (Exhibit 7.5):

We will continue to use the zero-coupon bond prices prevailing on March 31, 2011, as our other inputs:

Number of periods:	4
Length of periods (years)	1
Number of risk factors:	1
Volatility term structure:	Empirical

Period Start	Period End	Parameter Inputs
		Zero-Coupon Bond Prices
0	1	0.99700579
1	2	0.98411015
2	3	0.96189224
3	4	0.93085510

Our final assumption is that there is one random factor driving the forward rate curve. This implies (since all of our volatilities are positive) that all forward rates move up or down together. We relax this assumption when we move to two- and three-factor examples.

KEY IMPLICATIONS AND NOTATION OF THE HJM APPROACH

We now reintroduce the same notation used in Chapter 6:

Δ	= length of time period, which is 1 in this example
$r(t, s_t)$	= the simple one-period, uncompounded risk-free interest rate as of current time t in state s_t
$R(t, s_t)$	= $1 + r(t, s_t)$, the total return, the value of \$1 dollar invested at the risk-free rate for one period with no compounding
$\sigma(t, T, s_t)$	= forward rate volatility at time t for forward maturing at T in state s (a sequence of ups and downs). Because σ is now dependent on the spot U.S. Treasury rate, the state s_t is very important in determining its level.

$P(t, T, s_t)$	= zero-coupon bond price at time t maturing at time T in state s_t (i.e., up or down)
Index	= 1 if the state is up and = -1 if the state is down
$U(t, T, s_t)$	= the total return that leads to bond price $P(t + \Delta, T, s_t = \text{up})$ on a T maturity bond in an upshift
$D(t, T, s_t)$	= the total return that leads to bond price $P(t + \Delta, T, s_t = \text{down})$ on a T maturity bond in a downshift
$K(t, T, s_t)$	= the sum of the forward volatilities as of time t for the one-period forward rates from $t + \Delta$ to $T - \Delta$, as shown here

$$K(t, T, s_t) = \sum_{j=t+\Delta}^{T-\Delta} \sigma(t, j, s_t) \sqrt{\Delta}$$

We again use this trigonometric function, which is found in common spreadsheet software:

$$\text{Cosh}(x) = \frac{1}{2} [e^x + e^{-x}]$$

In the interests of saving space, we again arrange the bushy tree to look like a table by stretching it in the same way shown in Chapter 6.

A completely populated zero-coupon bond price tree would then be summarized like the one in Exhibit 7.6 for the zero coupon bond maturing at time $T = 4$ where the final shift in price is shaded for up- or downshifts.

The mapping of the sequence of up and down states is shown in Exhibit 7.7, consistent with the previous stretched tree and with Chapter 6.

PSEUDO-PROBABILITIES

We use the same pseudo-probabilities, consistent with no arbitrage, as in Chapter 6. There is an upshift probability in each state of 1/2 and a downshift probability in each state of 1/2. We now demonstrate how to construct the bushy tree with our new volatility assumptions and how to use it for risk-neutral valuation.

THE FORMULA FOR ZERO-COUPON BOND PRICE SHIFTS

The formula for the shift in the array of zero-coupon bond prices is unchanged from Chapter 6:

$$P(t + \Delta, T, s_{t+\Delta}) = \frac{[P(t, T, s_t) R(t, s_t) e^{K(t, T, s_t) \Delta (\text{Index})}]}{[\cosh(K(t, T, s_t) \Delta)]} \quad (7.1)$$

EXHIBIT 7.5 Standard Deviation of Annual Changes in One-Year Rate

Data Group	1-Year Rate in Data Group		Number of Observations	1 Year	2 Years	3 Years	4 Years
	Minimum Spot Rate	Maximum Spot Rate					
1	0.00%	0.09%	Assumed	0.01000000%	0.02000000%	0.10000000%	0.20000000%
2	0.09%	0.25%	Assumed	0.04000000%	0.10000000%	0.20000000%	0.30000000%
3	0.25%	0.50%	346	0.06363856%	0.38419072%	0.67017474%	0.86385093%
4	0.50%	0.75%	86	0.08298899%	0.46375394%	0.61813347%	0.80963400%
5	0.75%	1.00%	20	0.94371413%	1.08296308%	0.96789760%	0.74672896%
6	1.00%	2.00%	580	1.01658509%	0.84549720%	0.63127387%	0.66276260%
7	2.00%	3.00%	614	1.30986113%	1.28856305%	1.07279248%	0.80122998%
8	3.00%	4.00%	1492	1.22296401%	1.11359303%	1.02810902%	0.90449832%
9	4.00%	5.00%	1672	1.53223304%	1.21409679%	1.00407389%	0.83492621%
10	5.00%	6.00%	2411	1.22653783%	1.09067681%	0.97034198%	0.91785966%
11	6.00%	7.00%	1482	1.67496041%	1.40755286%	1.20035444%	1.04528208%
12	7.00%	8.00%	1218	1.70521091%	1.34675823%	1.30359062%	1.29177866%
13	8.00%	9.00%	765	2.16438588%	1.77497758%	1.72454046%	1.69404480%
14	9.00%	10.00%	500	2.28886301%	2.14441843%	2.02862693%	2.02690844%
15	10.00%	16.60%	959	2.86758369%	2.64068702%	2.55664943%	2.32042833%
Grand Total			12145	1.73030695%	1.49119219%	1.37014684%	1.26364341%

EXHIBIT 7.6 Four-Year, Zero-Coupon Bond

Row Number	Current Time			
	0	1	2	3
1	0.930855099	0.951557892	0.980029882	0.989916622
2		0.915743405	0.940645037	0.980777516
3			0.945073689	0.97890891
4			0.917568724	0.954003494
5				0.977728127
6				0.952852752
7				0.963793904
8				0.942565694

EXHIBIT 7.7 Map of Sequence of States

Row Number	Current Time			
	0	1	2	3
1	Time 0	up	up-up	up-up-up
2		dn	up-dn	up-up-dn
3			dn-up	up-dn-up
4			dn-dn	up-dn-dn
5				dn-up-up
6				dn-up-dn
7				dn-dn-up
8				dn-dn-dn

where for notational convenience $\cosh[K(t, T, s_t)\Delta] \equiv 1$ when $T - \Delta < t + \Delta$. This is equation 15.17 in Jarrow (2002, 286). We now put this formula to use.

BUILDING THE BUSHY TREE FOR ZERO-COUPON BONDS MATURING AT TIME $T = 2$

We now populate the bushy tree for the two-year, zero-coupon bond. We calculate each element of equation (7.1). When $t = 0$ and $T = 2$, we know $\Delta = 1$ and

$$P(0, 2, s_t) = 0.98411015$$

The one-period, risk-free rate is again

$$R(0, s_t) = 1/P(0, 1, s_t) = 1/0.997005786 = 1.003003206$$

The one-period spot rate for U.S. Treasuries is $r(0, s_t) = R(0, s_t) - 1 = 0.3003206$ percent. At this level of the spot rate for U.S. Treasuries, volatilities are selected from data group 3 in the previous look-up table. The volatilities for the one-year forward rates maturing in years 2, 3, and 4 are 0.003841907, 0.006701747, and 0.008638509.

The scaled sum of sigmas $K(t, T, s_t)$ for $t = 0$ and $T = 2$ becomes

$$K(t, T, s_t) = \sum_{j=t+\Delta}^{T-\Delta} \sigma(t, j, s_t) \sqrt{\Delta} = \sum_{j=1}^1 \sigma(0, j, s_t) \sqrt{\Delta}$$

and, therefore,

$$K(0, T, s_t) = (\sqrt{1})(0.003841907) = 0.003841907$$

Using equation (7.1) with these inputs and the fact that the variable Index = 1 for an upshift gives

$$P(1, 2, s_t = \text{up}) = 0.990858$$

For a downshift we set Index = -1 and recalculate equation (7.1) to get

$$P(1, 2, s_t = \text{down}) = 0.983273$$

We have fully populated the bushy tree for the zero-coupon bond maturing at $T = 2$ (note values have been rounded to six decimal places for display only), since all of the up and down states at time $t = 2$ result in a riskless pay-off of the zero-coupon bond at its face value of one. Note also that, as a result of our volatility assumptions, the price of the bond maturing at time $T = 2$ after an upshift to time $t = 1$ no longer results in a value higher than the par value of 1. The change in our volatility assumptions has eliminated the negative rates found in example one, since all zero coupon bond prices in Exhibit 7.8 are less than one.

BUILDING THE BUSHY TREE FOR ZERO-COUPON BONDS MATURING AT TIME $T = 4$

We now populate the bushy tree for the four-year, zero-coupon bond. We calculate each element of equation (7.1). When $t = 0$ and $T = 4$, we know $\Delta = 1$ and

$$P(0, 4, s_t) = 0.930855099$$

The one-period, risk-free rate is

$$R(0, s_t) = 1/P(0, 1, s_t) = 1/0.997005786 = 1.003003206$$

EXHIBIT 7.8 Two-Year, Zero-Coupon Bond

Row Number	Current Time			
	0	1	2	3
1	0.98411015	0.990857831	1	
2		0.983273439	1	
3			1	
4			1	
5				
6				
7				
8				

The spot U.S. Treasury rate $r(0, s_t)$ is still 0.003003206 and the same volatilities from data group 3 are relevant. The scaled sum of sigmas $K(t, T, s_t)$ for $t = 0$ and $T = 4$ becomes

$$K(t, T, s_t) = \sum_{j=t+\Delta}^{T-\Delta} \sigma(t, j, s_t) \sqrt{\Delta} = \sum_{j=1}^3 \sigma(0, j, s_t) \sqrt{\Delta}$$

and, therefore,

$$K(0, T, s_t) = (\sqrt{1})(0.003841907 + 0.006701747 + 0.008638509) = 0.019182164$$

Using equation (7.1) with these inputs and the fact that the variable Index = 1 for an upshift gives

$$P(1, 4, s_t = \text{up}) = 0.951558$$

For a downshift we set Index = -1 and recalculate equation (7.1) to get

$$P(1, 4, s_t = \text{down}) = 0.915743$$

We have correctly populated the first two columns of the bushy tree for the zero-coupon bond maturing at $T = 4$ (note values have been rounded to six decimal places for display only) as shown in Exhibit 7.9.

Now we move to the third column, which displays the outcome of the $T = 4$ zero-coupon bond price after four scenarios: up up, up down, down up, and down down. We calculate $P(2, 4, s_t = \text{up})$ and $P(2, 4, s_t = \text{down})$ after the initial downstate as follows. When $t = 1$, $T = 4$, and $\Delta = 1$ then the one-period, risk-free rate is

$$R(1, s_t = \text{down}) = 1/P(1, 2, s_t = \text{down}) = 1/0.983273 = 1.017011098$$

EXHIBIT 7.9 Four-Year, Zero-Coupon Bond

Row Number	Current Time			
	0	1	2	3
1	0.930855099	0.951557892	0.980029882	0.989916622
2		0.915743405	0.940645037	0.980777516
3			0.945073689	0.97890891
4			0.917568724	0.954003494
5				0.977728127
6				0.952852752
7				0.963793904
8				0.942565694

The U.S. Treasury spot rate is $r(1, s_t = \text{down}) = R(1, s_t = \text{down}) - 1 = 1.7011098$ percent, which is in data group 6. The volatilities for the two remaining one-period forward rates that are relevant are taken from the lookup table for data group 6: 0.008454972 and 0.006312739.

The scaled sum of sigmas $K(t, T, s_t)$ for $t = 1$ and $T = 4$ becomes

$$K(t, T, s_t) = \sum_{j=t+\Delta}^{T-\Delta} \sigma(t, j, s_t) \sqrt{\Delta} = \sum_{j=2}^3 \sigma(0, j, s_t) \sqrt{\Delta}$$

and, therefore,

$$K(1, 4, s_t) = (\sqrt{1})(0.008454972 + 0.006312739) = 0.014767711$$

Using equation (7.1) with these inputs and the fact that the variable Index = 1 for an upshift gives

$$P(2, 4, s_t = \text{up}) = 0.945074$$

For a downshift, we set Index = -1 and recalculate equation (7.1) to get

$$P(2, 4, s_t = \text{down}) = 0.917569$$

We have correctly populated the third and fourth rows of column 3 ($t = 2$) of the bushy tree for the zero-coupon bond maturing at $T = 4$ (note values have been rounded to six decimal places for display only) in Exhibit 7.10. The remaining calculations are left to the reader.

In a similar way, the bushy tree for the price of the zero-coupon bond maturing at time $T = 3$ can be calculated as shown in Exhibit 7.11.

If we combine all of these tables, we can create a table of the term structure of zero-coupon bond prices in each scenario as in Exhibit 7.12.

At any point in time t , the continuously compounded yield to maturity at time T can be calculated as $y(T - t) = -\ln[P(t, T)]/(T - t)$. Unlike Chapter 6, no negative rates appear in the bushy tree because of the revised approach to interest rate volatility (Exhibit 7.13).

We can graph yield curve movements as shown in Exhibit 7.14. Note that at time $T = 3$, only the bond maturing in time $T = 4$ remains, so there is only one data point observable on the yield curve.

Finally, we can display the simple interest one-year U.S. Treasury spot rates and the associated term structure of one-year forward rates in each scenario (Exhibit 7.15).

VALUATION IN THE HJM FRAMEWORK

Valuation using our revised volatility assumptions is identical in methodology to Chapter 6. Note that column 1 denotes the riskless one-period interest rate in each scenario. For the scenario number 14 (three consecutive downshifts in zero-coupon

EXHIBIT 7.10 Four-Year, Zero-Coupon Bond

Row Number	Current Time			
	0	1	2	3
1	0.930855099	0.951557892	0.980029882	0.989916622
2		0.915743405	0.940645037	0.980777516
3			0.945073689	0.97890891
4			0.917568724	0.954003494
5				0.977728127
6				0.952852752
7				0.963793904
8				0.942565694

EXHIBIT 7.11 Three-Year, Zero-Coupon Bond

Row Number	Current Time			
	0	1	2	3
1	0.961892238	0.974952939	0.994603742	1
2		0.954609057	0.97329298	1
3			0.979056303	1
4			0.962639708	1
5				1
6				1
7				1
8				1

EXHIBIT 7.12 Zero-Coupon Bond Prices

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.997005786	0.98411015	0.961892238	0.930855099
Time 1 up	1	0.990857831	0.974952939	0.951557892	
Time 1 down	2	0.983273439	0.954609057	0.915743405	
Time 2 up-up	3	0.994603742	0.980029882		
Time 2 up-down	4	0.97329298	0.940645037		
Time 2 down-up	5	0.979056303	0.945073689		
Time 2 down-down	6	0.962639708	0.917568724		
Time 3 up-up-up	7	0.989916622			
Time 3 up-up-down	8	0.980777516			
Time 3 up-down-up	9	0.97890891			
Time 3 up-down-down	10	0.954003494			
Time 3 down-up-up	11	0.977728127			
Time 3 down-up-down	12	0.952852752			
Time 3 down-down-up	13	0.963793904			
Time 3 down-down-down	14	0.942565694			

bond prices), cash flows at time $T = 4$ would be discounted by the one-year spot rates at time $t = 0$, by the one-year spot rate at time $t = 1$ in scenario 2 (down), by the one-year spot rate in scenario 6 (down down) at time $t = 2$, and by the one-year spot rate at time $t = 3$ in scenario 14 (down down down). The discount factor is

$$\begin{aligned} &\text{Discount factor}(0, 4, \text{down down down}) \\ &= 1/(1.003003)(1.017011)(1.038810)(1.060934) \end{aligned}$$

These discount factors are displayed in Exhibit 7.16 for each potential cash flow date.


When taking expected values, we can calculate the probability of each scenario coming about since the probability of an upshift is one-half (Exhibit 7.17).

As in Chapter 6, it is convenient to calculate the probability-weighted discount factors for use in calculating the expected present value of cash flows (Exhibit 7.18).

We now use the HJM bushy trees we have generated to value representative types of securities.

EXHIBIT 7.13 Continuously Compounded Zero-Coupon Yields

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.2999%	0.8009%	1.2951%	1.7913%
Time 1 up	1	0.9184%	1.2683%	1.6552%	
Time 1 down	2	1.6868%	2.3227%	2.9340%	
Time 2 up-up	3	0.5411%	1.0086%		
Time 2 up-down	4	2.7070%	3.0595%		
Time 2 down-up	5	2.1166%	2.8246%		
Time 2 down-down	6	3.8076%	4.3014%		
Time 3 up-up-up	7	1.0135%			
Time 3 up-up-down	8	1.9410%			
Time 3 up-down-up	9	2.1317%			
Time 3 up-down-down	10	4.7088%			
Time 3 down-up-up	11	2.2524%			
Time 3 down-up-down	12	4.8295%			
Time 3 down-down-up	13	3.6878%			
Time 3 down-down-down	14	5.9150%			

**EXHIBIT 7.14** Kamakura Corporation, HJM Zero-Coupon Yield Curve Movements

VALUATION OF A ZERO-COUPON BOND MATURING AT TIME $T = 4$

A riskless zero-coupon bond pays \$1 in each of the 8 nodes of the bushy tree that prevail at time $T = 4$ (Exhibit 7.19).

When we multiply this vector of 1s times the probability-weighted discount factors in the time $T = 4$ column in the previous table and add them, we get a zero-coupon

EXHIBIT 7.15 One-Year Forward Rates Simple Interest

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.3003%	1.3104%	2.3098%	3.3343%
Time 1 up	1	0.9227%	1.6313%	2.4586%	
Time 1 down	2	1.7011%	3.0027%	4.2442%	
Time 2 up-up	3	0.5426%	1.4871%		
Time 2 up-down	4	2.7440%	3.4708%		
Time 2 down-up	5	2.1392%	3.5958%		
Time 2 down-down	6	3.8810%	4.9120%		
Time 3 up-up-up	7	1.0186%			
Time 3 up-up-down	8	1.9599%			
Time 3 up-down-up	9	2.1546%			
Time 3 up-down-down	10	4.8214%			
Time 3 down-up-up	11	2.2779%			
Time 3 down-up-down	12	4.9480%			
Time 3 down-down-up	13	3.7566%			
Time 3 down-down-down	14	6.0934%			

EXHIBIT 7.16 Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1	0.997005786	0.987890991	0.982560076	0.972652552
2		0.997005786	0.987890991	0.982560076	0.963672831
3			0.980329308	0.961507367	0.941228129
4			0.980329308	0.961507367	0.917281387
5				0.959797588	0.938421098
6				0.959797588	0.914545773
7				0.94370392	0.909536085
8				0.94370392	0.88950294

bond price of 0.93085510, which is the value we should get in a no-arbitrage economy, the value observable in the market and used as an input to create the tree.

VALUATION OF A COUPON-BEARING BOND PAYING ANNUAL INTEREST

Next we value a bond with no credit risk that pays \$3 in interest at every scenario at times $T = 1, 2, 3$, and 4 plus principal of 100 at time $T = 4$. The valuation is calculated by multiplying each cash flow by the matching probability weighted discount

EXHIBIT 7.17 Probability of Each State

Row Number	Current Time				
	0	1	2	3	4
1	100.00%	50.00%	25.00%	12.50%	12.50%
2		50.00%	25.00%	12.50%	12.50%
3			25.00%	12.50%	12.50%
4			25.00%	12.50%	12.50%
5				12.50%	12.50%
6				12.50%	12.50%
7				12.50%	12.50%
8				12.50%	12.50%

EXHIBIT 7.18 Probability-Weighted Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1	0.498502893	0.246972748	0.12282001	0.121581569
2		0.498502893	0.246972748	0.12282001	0.120459104
3			0.245082327	0.120188421	0.117653516
4			0.245082327	0.120188421	0.114660173
5				0.119974699	0.117302637
6				0.119974699	0.114318222
7				0.11796299	0.113692011
8				0.11796299	0.111187867

EXHIBIT 7.19 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0	0	0	0	1
2		0	0	0	1
3			0	0	1
4			0	0	1
5				0	1
6				0	1
7				0	1
8				0	1
Risk-Neutral Value =					0.93085510

EXHIBIT 7.20 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0	3	3	3	103
2		3	3	3	103
3			3	3	103
4			3	3	103
5				3	103
6				3	103
7				3	103
8				3	103
Risk-Neutral Value =					104.70709974

factor, to get a value of 104.70709974. It will surprise many that this is the same value that we arrived at in Chapter 6, even though the volatilities used are different. The values are the same because, by construction, our valuations for the zero-coupon bond prices at time zero for maturities at $T = 1, 2, 3$, and 4 continue to match the inputs. Multiplying these zero-coupon bond prices times 3, 3, 3, and 103 also leads to a value of 104.70709974 as it should (Exhibit 7.20).

VALUATION OF A DIGITAL OPTION ON THE ONE-YEAR U.S. TREASURY RATE

Now we value a digital option that pays \$1 at time $T = 3$ if (at that time) the one-year U.S. Treasury rate (for maturity at $T = 4$) is over 4 percent. If we look at the previous table of the term structure of one-year spot and forward rates, this happens in three scenarios: scenarios 10 (up down down), 12 (down up down), and 14 (down down down). The cash flow can be input in the table in Exhibit 7.21 and multiplied by the probability-weighted discount factors to find that this option has a value of 0.35812611.

CONCLUSION

This chapter, an extension of Chapter 6, shows how to simulate zero-coupon bond prices, forward rates, and zero-coupon bond yields in an HJM framework with rate-dependent and maturity-dependent interest rate volatility. Monte Carlo simulation, an alternative to the bushy tree framework, can be done in a fully consistent way. We discuss HJM Monte Carlo simulation in Chapter 10.

In Chapter 8, we enrich the realism of the simulation by recognizing one very important flaw in the current example. In the current example, because all of the sigmas are positive, when the single risk factor shocks forward rates, they will all

EXHIBIT 7.21 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0	0	0	0	0
2		0	0	0	0
3			0	0	0
4			0	1	0
5				0	0
6				1	0
7				0	0
8				1	0
Risk-Neutral Value =					0.35812611

move up together, move down together, or remain unchanged together. This means that rates are 100 percent correlated with each other. If one examines the actual correlation of one-year changes in forward rates, however, correlations fall well below 100 percent (see Exhibit 7.22).

In the next two chapters, we introduce a second and third risk factor in order to match both observed volatilities and correlations.

EXHIBIT 7.22 Correlation of One-Year U.S. Treasury Spot and One-Year U.S. Treasury Forward Rates, 1963–2011

U.S. Treasury 1-Year Spot or 1-Year Forward Rate Maturing in Year										
Maturity	1	2	3	4	5	6	7	8	9	10
1	1									
2	0.898446916	1								
3	0.803276999	0.962340885	1							
4	0.73912125	0.912764405	0.969297407	1						
5	0.632113264	0.842699969	0.888021718	0.959258921	1					
6	0.538514767	0.776475154	0.838137715	0.895129164	0.960934857	1				
7	0.487825492	0.726944387	0.809716383	0.842848126	0.889793518	0.973990896	1			
8	0.504654811	0.733078076	0.820332916	0.857470216	0.880224306	0.934078276	0.970495319	1		
9	0.512063356	0.715579552	0.790776168	0.846345826	0.854053003	0.845015972	0.862326831	0.954766006	1	
10	0.491523904	0.67310992	0.734644592	0.804337673	0.809254708	0.758401192	0.755353037	0.878460628	0.979528635	1

HJM Interest Rate Modeling with Two Risk Factors

In Chapters 6 and 7, we provided worked examples of how to use the yield curve simulation framework of Heath, Jarrow, and Morton (HJM) with two different assumptions about the volatility of forward rates. The first assumption was that volatility was dependent on the maturity of the forward rate and nothing else. The second assumption was that volatility of forward rates was dependent on both the level of rates and the maturity of forward rates being modeled. Both of these models were one-factor models, implying that random rate shifts are either all positive, all negative, or zero. This kind of yield curve movement is not consistent with the yield curve twists that are extremely common in the U.S. Treasury market and most other fixed income markets. In this chapter, we generalize the model to include two risk factors in order to increase the realism of the simulated yield curve.

PROBABILITY OF YIELD CURVE TWISTS IN THE U.S. TREASURY MARKET

In this chapter, we use the same data as inputs to the yield curve simulation process that we used in Chapters 6 and 7. A critical change in this chapter is to enhance our ability to model twists in the yield curve as shown in Exhibit 8.1.

EXHIBIT 8.1 Summary of U.S. Treasury Yield Curve Shifts Percentage Distribution of Daily Rate Shifts January 2, 1962, to August 22, 2011

Type of Shift	U.S. Treasury Input Data	Monthly Zero- Coupon Bond Yields	Monthly Forward Rates
Negative shift	18.4%	12.0%	2.4%
Positive shift	18.6%	12.1%	2.6%
Zero shift	0.7%	0.7%	0.7%
Subtotal	37.7%	24.8%	5.7%
Twist	62.3%	75.2%	94.3%
Grand Total	100.0%	100.0%	100.0%

It shows that, for 12,386 days of movements in U.S. Treasury forward rates, yield curve twists occurred on 94.3 percent of the observations.

In order to incorporate yield curve twists in interest rate simulations under the Heath, Jarrow, and Morton (HJM) framework, we introduce a second risk factor driving interest rates in this chapter. The single-factor yield curve models of Ho and Lee (1986), Vasicek (1977), Hull and White (1990b), Black, Derman, and Toy (1990), and Black and Karasinski(1991) cannot model yield curve twists and, therefore, they do not provide a sufficient basis for interest rate risk management.

OBJECTIVES OF THE EXAMPLE AND KEY INPUT DATA

Following Jarrow (2002), we again make the same modeling assumptions for our worked example as in Chapters 6 and 7:

- Zero-coupon bond prices for the U.S. Treasury curve on March 31, 2011, are the basic inputs.
- Interest rate volatility assumptions are based on the Dickler, Jarrow, and van Deventer papers on daily U.S. Treasury yields and forward rates from 1962 to 2011. In this chapter, we retain the volatility assumptions used in Chapter 7 but expand the number of random risk factors driving interest rates to two factors.
- The modeling period is four equal length periods of one year each.
- The HJM implementation is that of a “bushy tree” that we describe below.

In Chapters 6 and 7, the HJM bushy tree that we constructed consisted solely of upshifts and downshifts because we were modeling as if only one random factor was driving interest rates. In this chapter, with two random factors, we move to a bushy tree that has upshifts, midshifts, and downshifts at each node in the tree, as shown in Exhibit 8.2. Please consider these terms as labels only, since forward rates and zero-coupon bond prices, of course, move in opposite directions.

At each of the points in time on the lattice (time 0, 1, 2, 3, and 4) there are sets of zero-coupon bond prices and forward rates. At time 0, there is one set of data. At time one, there are three sets of data, the up-set, mid-set, and down-set. At time two, there are nine sets of data (up up, up mid, up down, mid up, mid mid, mid down, down up, down mid, and down down), and at time three there are $27 = 3^3$ sets of data.

The table in Exhibit 8.3 shows the actual volatilities, resulting from all risk factors, for the one-year changes in continuously compounded forward rates from 1963 to 2011. This is the same data we used in Chapter 7. We use this table later to divide total volatility between two uncorrelated risk factors.

We will again use the zero-coupon bond prices prevailing on March 31, 2011, as our other inputs.

Number of periods:	4
Length of periods (years)	1
Number of risk factors:	1
Volatility term structure:	Empirical, two factor

Period Start	Period End	Parameter Inputs Zero-Coupon Bond Prices
0	1	0.99700579
1	2	0.98411015
2	3	0.96189224
3	4	0.93085510

INTRODUCING A SECOND RISK FACTOR DRIVING INTEREST RATES

In the first two worked examples of the HJM approach in Chapters 6 and 7, the nature of the single factor that shocks one-year spot and forward rates was not specified. In this chapter, we take a cue from the popular academic models of the term structure of interest rates and postulate that one of the two factors driving changes in forward rates is the change in the short-run rate of interest. In the context of our annual model, the short rate is the one-year spot U.S. Treasury yield. For each of the one-year U.S. Treasury forward rates $f_k(t)$, we run the regression

$$\begin{aligned}
 [\text{Change in } f_k(i)] &= a \\
 &+ b[\text{Change in one-year U.S. Treasury spot rate } (i)] \\
 &+ e(i)
 \end{aligned}$$

where the change in continuously compounded yields is measured over annual intervals from 1963 to 2011. The coefficients of the regressions for the one-year forward rates maturing in years 2, 3, . . . , 10 are as shown in Exhibit 8.4.

Graphically, it is easy to see that the historical response of forward rates to the spot rate of interest is neither constant—as assumed by Ho and Lee (1986)—nor declining, as assumed by Vasicek (1977) and Hull and White (1990b). Indeed, the response of the one-year forwards maturing in years 8, 9, and 10 years to changes in the one-year spot rate is larger than the response of the one-year forward maturing in year 7 as shown in Exhibit 8.5).

We use these regression coefficients to separate the total volatility (shown in the previous table) of each forward rate between two risk factors. The first risk factor is changes in the one-year spot U.S. Treasury. The second factor represents all other sources of shocks to forward rates.

Because of the nature of the linear regression of changes in one-year forward rates on the first risk factor, changes in the one-year spot U.S. Treasury rate, we know that risk factor 1 and risk factor 2 are uncorrelated. Therefore, total volatility for the forward rate maturing in $T - t = k$ years can be divided as follows between the two risk factors:

$$\sigma_{k, total}^2 = \sigma_{k, 1}^2 + \sigma_{k, 2}^2$$




EXHIBIT 8.2 Example of Bushy Tree with Two Risk Factors for HJM Modeling of No-Arbitrage, Zero-Coupon Bond Price Movements


EXHIBIT 8.3 Standard Deviation of Annual Changes in One-Year Rate

Data Group	1-Year Rate in Data Group		Number of Observations	1-Year	2-Years	3-Years	4-Years
	Minimum Spot Rate	Maximum Spot Rate					
1	0.0000%	0.0900%	Assumed	0.01000000%	0.02000000%	0.10000000%	0.20000000%
2	0.0900%	0.2500%	Assumed	0.04000000%	0.10000000%	0.20000000%	0.30000000%
3	0.2500%	0.5000%	346	0.06363856%	0.38419072%	0.67017474%	0.86385093%
4	0.5000%	0.7500%	86	0.08298899%	0.46375394%	0.61813347%	0.80963400%
5	0.7500%	1.0000%	20	0.94371413%	1.08296308%	0.96789760%	0.74672896%
6	1.0000%	2.0000%	580	1.01658509%	0.84549720%	0.63127387%	0.66276260%
7	2.0000%	3.0000%	614	1.30986113%	1.28856305%	1.07279248%	0.80122998%
8	3.0000%	4.0000%	1492	1.22296401%	1.11359303%	1.02810902%	0.90449832%
9	4.0000%	5.0000%	1672	1.53223304%	1.21409679%	1.00407389%	0.83492621%
10	5.0000%	6.0000%	2411	1.22653783%	1.09067681%	0.97034198%	0.91785966%
11	6.0000%	7.0000%	1482	1.67496041%	1.40755286%	1.20035444%	1.04528208%
12	7.0000%	8.0000%	1218	1.70521091%	1.34675823%	1.30359062%	1.29177866%
13	8.0000%	9.0000%	765	2.16438588%	1.77497758%	1.72454046%	1.69404480%
14	9.0000%	10.0000%	500	2.28886301%	2.14441843%	2.02862693%	2.02690844%
15	10.0000%	16.6000%	959	2.86758369%	2.64068702%	2.55664943%	2.32042833%
Grand Total			12145	1.73030695%	1.49119219%	1.37014684%	1.26364341%

EXHIBIT 8.4 Regression Coefficients

Dependent Variable	Coefficient b on 1-Year Change in 1-Year U.S. Treasury Rate	Standard Error of Regression
Forward 2	0.774288643	0.006547828
Forward 3	0.636076417	0.008161008
Forward 4	0.539780355	0.008511904
Forward 5	0.439882997	0.009330728
Forward 6	0.353383944	0.009567974
Forward 7	0.309579768	0.009585926
Forward 8	0.309049192	0.009148432
Forward 9	0.322513079	0.009361197
Forward 10	0.327367707	0.010036507

Sources: Kamakura Corporation; Federal Reserve.

**EXHIBIT 8.5** Percent Response by U.S. Treasury One-Year Forward Rates to 1 Percent Shift in One-Year U.S. Treasury Spot Rate, 1962–2011

We also know that the risk contribution of the first risk factor, the change in the spot one-year U.S. Treasury rate, is proportional to the regression coefficient α_k of changes in forward rate maturing in year k on changes in the one-year spot rate. We denote the total volatility of the one-year U.S. Treasury spot rate by the subscript $_{1, total}$:

$$\sigma_{k,1}^2 = \alpha_k^2 \sigma_{1,total}^2$$

This allows us to solve for the volatility of risk factor 2 using this equation:

$$\sigma_{k,2} = [\sigma_{k,total}^2 - \alpha_k^2 \sigma_{1,total}^2]^{1/2}$$

Because the total volatility for each forward rate varies by the level of the one-year U.S. Treasury spot rate, so will the values of $\sigma_{k,1}$ and $\sigma_{k,2}$. In one case (data group 6 for the forward rate maturing at time $T = 3$), we set the sigma for the second risk factor to zero and ascribed all of the total volatility to risk factor 1 because the calculated contribution of risk factor 1 was greater than the total volatility. In the worked example, we select the volatilities from lookup tables for the appropriate risk factor volatility. Using the previous equations, the look-up table for risk factor 1 (changes in the one-year spot rate) volatility is given in Exhibit 8.6.

The lookup table for the general “all other” risk factor 2 is as shown in Exhibit 8.7.

KEY IMPLICATIONS AND NOTATION OF THE HJM APPROACH

We now reintroduce a slightly modified version of the notation used in Chapters 6 and 7:

Δ	= length of time period, which is 1 in this example
$r(t, s_t)$	= the simple one-period, risk-free interest rate as of current time t in state s_t
$R(t, s_t)$	= $1 + r(t, s_t)$, the total return, the value of \$1 dollar invested at the risk-free rate for 1 period
$f(t, T, s_t)$	= the simple one-period forward rate maturing at time T in state s_t as of time t
$F(t, T, s_t)$	= $1 + f(t, T, s_t)$, the return on \$1 invested at the forward rate for one period
$\sigma_1(t, T, s_t)$	= forward rate volatility due to risk factor 1 at time t for the one-period forward rate maturing at T in state s_t (a sequence of ups and downs). Because σ_1 is dependent on the spot U.S. Treasury rate, the state s_t is very important in determining its level.
$\sigma_2(t, T, s_t)$	= forward rate volatility due to risk factor 2 at time t for the one-period forward rate maturing at T in state s_t (a sequence of ups and downs). Because σ_2 is dependent on the spot U.S. Treasury rate, the state s_t is very important in determining its level.
$P(t, T, s_t)$	= zero-coupon bond price at time t maturing at time T in state s_t (i.e., up or down)
Index(1)	An indicator of the state for risk factor 1, defined below
Index(2)	An indicator of the state for risk factor 2, defined below
$K(i, t, T, s_t)$	= the sum of the forward volatilities for risk factor i (either 1 or 2) as of time t for the 1-period forward rates from $t + \Delta$ to T , as shown here:

$$K(i, t, T, s_t) = \sum_{j=t+\Delta}^T \sigma_i(t, j, s_t) \sqrt{\Delta}$$

EXHIBIT 8.6 Look-Up Table for Risk Factor 1

Data Group	1-Year Rate in Data Group		Number of Observations	Standard Deviation of Risk Factor 1, Sigma 1			
	Minimum Spot Rate	Maximum Spot Rate		1-Year	2-Years	3-Years	4-Years
1	0.00%	0.09%	Assumed	0.01000000%	0.00774289%	0.00492507%	0.00265845%
2	0.09%	0.25%	Assumed	0.04000000%	0.03097155%	0.01970027%	0.01063382%
3	0.25%	0.50%	346	0.06363856%	0.04927462%	0.03134242%	0.01691802%
4	0.50%	0.75%	86	0.08298899%	0.06425743%	0.04087264%	0.02206225%
5	0.75%	1.00%	20	0.94371413%	0.73070713%	0.46478558%	0.25088212%
6	1.00%	2.00%	580	1.01658509%	0.78713029%	0.63127387%	0.34074924%
7	2.00%	3.00%	614	1.30986113%	1.01421060%	0.64511544%	0.34822064%
8	3.00%	4.00%	1492	1.22296401%	0.94692714%	0.60231802%	0.32511944%
9	4.00%	5.00%	1672	1.53223304%	1.18639064%	0.75463511%	0.40733721%
10	5.00%	6.00%	2411	1.22653783%	0.94969431%	0.60407816%	0.32606952%
11	6.00%	7.00%	1482	1.67496041%	1.29690282%	0.82492930%	0.44528063%
12	7.00%	8.00%	1218	1.70521091%	1.32032544%	0.83982788%	0.45332259%
13	8.00%	9.00%	765	2.16438588%	1.67585941%	1.06597465%	0.57539217%
14	9.00%	10.00%	500	2.28886301%	1.77224063%	1.12728047%	0.60848385%
15	10.00%	16.60%	959	2.86758369%	2.22033748%	1.41230431%	0.76233412%
	Grand Total		12145	1.73030695%	1.33975702%	0.85218784%	0.45999426%

EXHIBIT 8.7 Look-Up Table for General Risk Factor 2

Data Group	1-Year Rate in Data Group			Standard Deviation of Risk Factor 2, Sigma 2			
	Minimum Spot Rate	Maximum Spot Rate	Number of Observations	1-Year	2-Years	3-Years	4-Years
1	0.00%	0.09%	Assumed	0.00000000%	0.01844038%	0.09987864%	0.19998233%
2	0.09%	0.25%	Assumed	0.00000000%	0.09508293%	0.19902738%	0.29981148%
3	0.25%	0.50%	346	0.00000000%	0.38101775%	0.66944144%	0.86368524%
4	0.50%	0.75%	86	0.00000000%	0.45928064%	0.61678068%	0.80933335%
5	0.75%	1.00%	20	0.00000000%	0.79929726%	0.84899949%	0.70332233%
6	1.00%	2.00%	580	0.00000000%	0.30869310%	0.00000000%	0.56845776%
7	2.00%	3.00%	614	0.00000000%	0.79484061%	0.85715213%	0.72160368%
8	3.00%	4.00%	1492	0.00000000%	0.58601914%	0.83319934%	0.84404654%
9	4.00%	5.00%	1672	0.00000000%	0.25789196%	0.66233695%	0.72881971%
10	5.00%	6.00%	2411	0.00000000%	0.53633630%	0.75937680%	0.85798894%
11	6.00%	7.00%	1482	0.00000000%	0.54703576%	0.87197617%	0.94569540%
12	7.00%	8.00%	1218	0.00000000%	0.26551545%	0.99701437%	1.20962421%
13	8.00%	9.00%	765	0.00000000%	0.58484243%	1.35563190%	1.59333350%
14	9.00%	10.00%	500	0.00000000%	1.20734981%	1.68658411%	1.93341801%
15	10.00%	16.60%	959	0.00000000%	1.42945073%	2.13116232%	2.19162824%
	Grand Total		12145	0.00000000%	0.65475588%	1.07288315%	1.17694518%

Note that this is a slightly different definition of K than we used in Chapters 6 and 7.

We again use the trigonometric function $\cosh(x)$:

$$\text{Cosh}(x) = \frac{1}{2}[e^x + e^{-x}]$$

We again arrange the tree to look like a table by stretching the bushy tree as in Exhibit 8.8.

A completely populated zero-coupon bond price tree would then be summarized like this; prices shown are for the zero-coupon bond price maturing at time $T = 4$ at times 0, 1, 2, and 3 as shown in Exhibit 8.9.

The mapping of the sequence of up and down states is shown in Exhibit 8.10, consistent with the stretched tree in Exhibit 8.8.

In order to populate the trees with zero-coupon bond prices and forward rates, we again need to select the relevant pseudo-probabilities.




EXHIBIT 8.8 Example of Bushy Tree with Two Risk Factors for HJM Modeling of No-Arbitrage, Zero-Coupon Bond Price Movements

EXHIBIT 8.9 Four-Year, Zero-Coupon Bond Price

Row Number	Current Time			
	0	1	2	3
1	0.93085510	0.96002942	0.98271534	0.99567640
2		0.90943520	0.95325638	0.98282580
3		0.93256899	0.96584087	0.98795974
4			0.94311642	0.98343465
5			0.93491776	0.97488551
6			0.91274456	0.96385706
7			0.96285825	0.99123966
8			0.95448798	0.98262266
9			0.93185064	0.97150669
10				0.98272135
11				0.96087480
12				0.95222434
13				0.98272135
14				0.96087480
15				0.95222434
16				0.96472003
17				0.95770868
18				0.93866906
19				0.98545473
20				0.97688802
21				0.96583693
22				0.98545473
23				0.97688802
24				0.96583693
25				0.97845156
26				0.96236727
27				0.95217151

PSEUDO-PROBABILITIES

Consistent with Chapter 7 of Jarrow (2002), and without loss of generality, we set the probability of an upshift to one-fourth, the probability of a midshift to one-fourth, and the probability of a downshift to one-half. The no-arbitrage restrictions that stem from this set of pseudo-probabilities are given next.

The Formula for Zero-Coupon Bond Price Shifts with Two Risk Factors

In Chapters 6 and 7, we used equation 15.17 in Jarrow (2002, 286) to calculate the no-arbitrage shifts in zero-coupon bond prices. Alternatively, when there is one risk factor, we could have used equation 15.19 in Jarrow (2002, 287) to shift forward rates and derive zero-coupon bond prices from the forward rates. Now that we have

EXHIBIT 8.10 Map of Sequence of States

Row Number	Current Time			
	0	1	2	3
1	Time 0	Up	Up Up	Up Up Up
2		Mid	Up Mid	Up Up Mid
3		Down	Up Down	Up Up Down
4			Mid Up	Up Mid Up
5			Mid Mid	Up Mid Mid
6			Mid Down	Up Mid Down
7			Down Up	Up Down Up
8			Down Mid	Up Down Mid
9			Down Down	Up Down Down
10				Mid Up Up
11				Mid Up Mid
12				Mid Up Down
13				Mid Mid Up
14				Mid Mid Mid
15				Mid Mid Down
16				Mid Down Up
17				Mid Down Mid
18				Mid Down Down
19				Down Up Up
20				Down Up Mid
21				Down Up Down
22				Down Mid Up
23				Down Mid Mid
24				Down Mid Down
25				Down Down Up
26				Down Down Mid
27				Down Down Down

two risk factors, it is convenient to calculate the forward rates first. We do this using equations 15.32, 15.39a, and 15.39b in Jarrow (2002, 293–296). We use this equation for the shift in forward rates:

$$F(t + \Delta, T, s_{t+\Delta}) = F(t, T, s_t) e^{[\mu(t, T, s_t)\Delta] \Delta} e^{[\text{Index}(1)\sigma_1(t, T, s_t)\sqrt{\Delta} + \text{Index}(2)\sqrt{2}\sigma_2(t, T, s_t)\sqrt{\Delta}]\Delta} \quad (8.1)$$

Where when $T = t + \Delta$, the expression μ becomes

$$[\mu(t, T)\Delta] \Delta = \log \left[\frac{1}{2} e^{[K(1, t, t+\Delta)]\Delta} \cosh \left(\Delta K(2, t, t + \Delta) \sqrt{2} \right) + \frac{1}{2} e^{-[K(1, t, t+\Delta)]\Delta} \right]$$

We use this expression to evaluate equation (8.1) when $T > t + \Delta$:

$$[\mu(t, T)\Delta]\Delta = \log \left[\frac{1}{2} e^{[K(1, t, T)]\Delta} \cosh \left(\Delta K(2, t, T) \sqrt{2} \right) + \frac{1}{2} e^{-[K(1, t, T)]\Delta} \right] - \sum_{j=t+\Delta}^{T-\Delta} [\mu(t, j)\Delta]\Delta$$

The values for the pseudo-probabilities and Index (1) and Index (2) are set as follows:

Type of Shift	Pseudo-Probability	Index (1)	Index (2)
Upshift	1/4	-1	-1
Midshift	1/4	-1	1
Downshift	1/2	1	0

Building the Bushy Tree for Zero-Coupon Bonds Maturing at Time $T = 2$

We now populate the bushy tree for the two-year, zero-coupon bond. We calculate each element of equation (8.1). When $t = 0$ and $T = 2$, we know $\Delta = 1$ and

$$P(0, 2, s_t) = 0.98411015$$

The one-period risk-free return (1 plus the risk-free rate) is again

$$R(0, s_t) = 1/P(0, 1, s_t) = 1/0.997005786 = 1.003003206$$

The one-period spot rate for U.S. Treasuries is $r(0, s_t) = R(0, s_t) - 1 = 0.3003206$ percent. At this level of the spot rate for U.S. Treasuries, volatilities for risk factor 1 are selected from data group 3 in the look-up table above. The volatilities for risk factor 1 for the one-year forward rates maturing in years 2, 3, and 4 are 0.000492746, 0.000313424, and 0.00016918. The volatilities for risk factor 2 for the one-year forward rates maturing in years 2, 3, and 4 are 0.003810177, 0.006694414, and 0.008636852.

The scaled sum of sigmas $K(1, t, T, s_t)$ for $t = 0$ and $T = 2$ becomes

$$K(1, t, T, s_t) = \sum_{j=t+\Delta}^{T-\Delta} \sigma_1(t, j, s_t) \sqrt{\Delta} = \sum_{j=1}^1 \sigma_1(0, j, s_t) \sqrt{\Delta}$$

and, therefore,

$$K(1, 0, T, s_t) = (\sqrt{1})(0.000492746) = 0.000492746$$

Similarly,

$$K(2, 0, T, s_t) = 0.00381077$$

We also can calculate that

$$[\mu(t, t + \Delta)\Delta] = 0.00000738$$

Using formula (8.1) with these inputs and the fact that the variable Index (1) = -1 and Index (2) = -1 for an upshift gives us the forward returns for an upshift, midshift, and downshift as follows: 1.007170561, 1.018083343, and 1.013610665. From these we calculate the zero-coupon bond prices:

$$P(1, 2, s_t = \text{up}) = 0.992880 = 1/F(1, 2, s_t = \text{up})$$

For a midshift, we set Index (1) = -1 and Index (2) = +1 and calculate

$$P(1, 2, s_t = \text{mid}) = 0.992238 = 1/F(1, 2, s_t = \text{mid})$$

For a downshift, we set Index (1) = 1 and Index (2) = 0 and recalculate formula 1 to get

$$P(1, 2, s_t = \text{down}) = 0.986572 = 1/F(1, 2, s_t = \text{down})$$

We have fully populated the bushy tree for the zero-coupon bond maturing at $T = 2$ (note values have been rounded for display only), since all the up-, mid- and downstates at time $t = 2$ result in a riskless payoff of the zero-coupon bond at its face value of 1 as shown in Exhibit 8.11.

Building the Bushy Tree for Zero-Coupon Bonds Maturing at Time $T = 3$

For the zero-coupon bonds and one-period forward returns ($= 1 + \text{forward rate}$) maturing at time $T = 3$, we use the same volatilities listed above for risk factors 1 and 2 to calculate

$$K(1, 0, 3, s_t) = 0.00080617$$

$$K(2, 0, 3, s_t) = 0.010504592$$

$$[\mu(t, T)\Delta] = 0.00004816$$

The resulting forward returns for an upshift, midshift, and downshift are 1.013189027, 1.032556197, and 1.023468132. Zero-coupon bond prices are calculated from the one-period forward returns, so

$$P(1, 3, s_t = \text{up}) = 1/[F(1, 2, s_t = \text{up})F(1, 3, s_t = \text{up})]$$

The zero-coupon bond prices for an upshift, midshift, and downshift are 0.979956, 0.951268, and 0.963950. To eight decimal places, we have populated the

EXHIBIT 8.11 Two-Year, Zero-Coupon Bond Price

Row Number	Current Time			
	0	1	2	3
1	0.98411015	0.99288049	1	
2		0.98223786	1	
3		0.98657210	1	
4			1	
5			1	
6			1	
7			1	
8			1	
9			1	

second column of the zero-coupon bond price table for the zero-coupon bond maturing at $T = 3$ as shown in Exhibit 8.12.

Building the Bushy Tree for Zero-Coupon Bonds Maturing at Time $T = 4$

We now populate the bushy tree for the zero-coupon bond maturing at $T = 4$. Using the same volatilities as before for both risk factors 1 and 2, we find that

$$K(1, 0, 4, s_t) = 0.000975351$$

$$K(2, 0, 4, s_t) = 0.019141444$$

$$[\mu(t, T)\Delta]\Delta = 0.00012830$$

Using formula (8.1) with the correct values for Index (1) and Index (2) leads to the following forward returns for an upshift, midshift, and downshift: 1.020756042, 1.045998862, and 1.033650059. The zero-coupon bond price is calculated as follows:

$$P(1, 4, s_t = \text{up}) = 1/[F(1, 2, s_t = \text{up})F(1, 3, s_t = \text{up})F(1, 4, s_t = \text{up})]$$

This gives us the three zero-coupon bond prices of the column labeled 1 in this table for up-, mid-, and downshifts: 0.96002942, 0.90943520, and 0.93256899 as shown in Exhibit 8.13.

Now, we move to the third column, which displays the outcome of the $T = 4$ zero-coupon bond price after nine scenarios: up up, up mid, up down, mid up, mid mid, mid down, down up, down mid, and down down. We calculate $P(2, 4, s_t = \text{up})$, $P(2, 4, s_t = \text{mid})$, and $P(2, 4, s_t = \text{down})$ after the initial down state as follows. When $t = 1$, $T = 4$, and $\Delta = 1$ then the volatilities for the two remaining one-period forward rates that are relevant are taken from the look-up table for data group 6 for risk factor 1: 0.007871303 and 0.006312739. For risk factor 2, the volatilities

EXHIBIT 8.12 Three-Year, Zero-Coupon Bond Price

Row Number	Current Time			
	0	1	2	3
1	0.96189224	0.97995583	0.99404203	1
2		0.95126818	0.98121252	1
3		0.96394999	0.98633804	1
4			0.98035908	1
5			0.97183666	1
6			0.96084271	1
7			0.98906435	1
8			0.98046626	1
9			0.96937468	1
10				1
11				1
12				1
13				1
14				1
15				1
16				1
17				1
18				1
19				1
20				1
21				1
22				1
23				1
24				1
25				1
26				1
27				1

for the two remaining one-period forward rates are also chosen from data group 6: 0.003086931 and 0. The zero value was described previously; the implied volatility for risk factor 1 was greater than measured total volatility for the forward rate maturing at $T = 3$ in data group 6, so the risk factor 1 volatility was set to total volatility and risk factor 2 volatility was set to zero.

At time 1 in the downstate, the zero-coupon bond prices for maturities at $T = 2$, 3, and 4 are 0.986572, 0.963950, and 0.932569. We make the intermediate calculations as previously for the zero-coupon bond maturing at $T = 4$:

$$K(1, 1, 4, s_t) = 0.014184042$$

$$K(2, 1, 4, s_t) = 0.003086931$$

$$[\mu(t, T)\Delta]\Delta = 0.00006964$$

EXHIBIT 8.13 Four-Year, Zero-Coupon Bond Price

Row Number	Current Time			
	0	1	2	3
1	0.93085510	0.96002942	0.98271534	0.99567640
2		0.90943520	0.95325638	0.98282580
3		0.93256899	0.96584087	0.98795974
4			0.94311642	0.98343465
5			0.93491776	0.97488551
6			0.91274456	0.96385706
7			0.96285825	0.99123966
8			0.95448798	0.98262266
9			0.93185064	0.97150669
10				0.98272135
11				0.96087480
12				0.95222434
13				0.98272135
14				0.96087480
15				0.95222434
16				0.96472003
17				0.95770868
18				0.93866906
19				0.98545473
20				0.97688802
21				0.96583693
22				0.98545473
23				0.97688802
24				0.96583693
25				0.97845156
26				0.96236727
27				0.95217151

We can calculate the values of the one-period forward return maturing at time $T = 4$ in the upstate, midstate, and downstate as follows: 1.027216983, 1.027216983, and 1.040268304. Similarly, using the appropriate intermediate calculations, we can calculate the forward returns for maturity at $T = 3$: 1.011056564, 1.01992291, and 1.031592858. Since

$$P(2, 4, s_t = \text{down up}) = 1/[F(1, 3, s_t = \text{down up})F(1, 4, s_t = \text{down up})]$$

the zero-coupon bond prices for maturity at $T = 4$ in the down up, down mid, and down down states are as follows:

0.962858

0.954488

0.931851

We have correctly populated the seventh, eighth, and ninth rows of column 3 ($t = 2$) of the previous bushy tree for the zero-coupon bond maturing at $T = 4$ (note values have been rounded to six decimal places for display only). The remaining calculations are left to the reader.

If we combine all of these tables, we can create a table of the term structure of zero-coupon bond prices in each scenario as in examples one and two. The shading highlights two nodes of the bushy tree where values are identical because of the occurrence of $\sigma_2 = 0$ at one point on the bushy tree as shown in Exhibit 8.14.

At any point in time t , the continuously compounded yield to maturity at time T can be calculated as $y(T - t) = -\ln[P(t, T)]/(T - t)$. Note that we have no negative rates on this bushy tree and that yield curve shifts are much more complex than in the two prior examples using one risk factor as shown in Exhibit 8.15.

We can graph yield curve movements as shown in Exhibit 8.16 at time $t = 1$. We plot yield curves for the up-, mid-, and downshifts. These shifts are relative to the one-period forward rates prevailing at time zero for maturity at time $T = 2$ and $T = 3$. Because these one-period forward rates were much higher than yields as of time $t = 0$, all three shifts produce yields higher than time-zero yields.

When we add nine yield curves prevailing at time $t = 3$ and 27 single-point yield curves prevailing at time $t = 4$, two things are very clear. First, yield curve movements in a two-factor model are much more complex and much more realistic than what we saw in the two one-factor examples. Second, in the low-yield environment prevailing as of March 31, 2011, no arbitrage yield curve simulation shows “there is nowhere to go but up” from a yield curve perspective as shown in Exhibit 8.17.

Finally, we can display the one-year U.S. Treasury spot rates and the associated term structure of one-year forward rates in each scenario as shown in Exhibit 8.18.

VALUATION IN THE HJM FRAMEWORK

Jarrow, as quoted in Chapter 6, described valuation as the expected value of cash flows using the risk-neutral probabilities. Note that column 1 denotes the riskless one-period interest rate in each scenario. For the scenario number 39 (three consecutive downshifts in zero-coupon bond prices), cash flows at time $T = 4$ would be discounted by the one-year spot rates at time $t = 0$, by the one-year spot rate at time $t = 1$ in scenario 3 (down), by the one-year spot rate in scenario 12 (down down) at time $t = 2$, and by the one-year spot rate at time $t = 3$ in scenario 39 (down down down). The discount factor is

$$\text{Discount factor}(0.4, \text{down down down}) = 1/(1.003003)(1.013611)(1.031593) \\ (1.050231)$$

These discount factors are displayed in Exhibit 8.19 for each potential cash flow date.

EXHIBIT 8.14 Zero-Coupon Bond Prices

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.99700579	0.98411015	0.96189224	0.93085510
Time 1 Up	1	0.99288049	0.97995583	0.96002942	
Time 1 Mid	2	0.98223786	0.95126818	0.90943520	
Time 1 Down	3	0.98657210	0.96394999	0.93256899	
Time 2 Up Up	4	0.99404203	0.98271534		
Time 2 Up Mid	5	0.98121252	0.95325638		
Time 2 Up Down	6	0.98633804	0.96584087		
Time 2 Mid Up	7	0.98035908	0.94311642		
Time 2 Mid Mid	8	0.97183666	0.93491776		
Time 2 Mid Down	9	0.96084271	0.91274456		
Time 2 Down Up	10	0.98906435	0.96285825		
Time 2 Down Mid	11	0.98046626	0.95448798		
Time 2 Down Down	12	0.96937468	0.93185064		
Time 3 Up Up Up	13	0.99567640			
Time 3 Up Up Mid	14	0.98282580			
Time 3 Up Up Down	15	0.98795974			
Time 3 Up Mid Up	16	0.98343465			
Time 3 Up Mid Mid	17	0.97488551			
Time 3 Up Mid Down	18	0.96385706			
Time 3 Up Down Up	19	0.99123966			
Time 3 Up Down Mid	20	0.98262266			
Time 3 Up Down Down	21	0.97150669			
Time 3 Mid Up Up	22	0.98272135			
Time 3 Mid Up Mid	23	0.96087480			
Time 3 Mid Up Down	24	0.95222434			
Time 3 Mid Mid Up	25	0.98272135			
Time 3 Mid Mid Mid	26	0.96087480			
Time 3 Mid Mid Down	27	0.95222434			
Time 3 Mid Down Up	28	0.96472003			
Time 3 Mid Down Mid	29	0.95770868			
Time 3 Mid Down Down	30	0.93866906			
Time 3 Down Up Up	31	0.98545473			
Time 3 Down Up Mid	32	0.97688802			
Time 3 Down Up Down	33	0.96583693			
Time 3 Down Mid Up	34	0.98545473			
Time 3 Down Mid Mid	35	0.97688802			
Time 3 Down Mid Down	36	0.96583693			
Time 3 Down Down Up	37	0.97845156			
Time 3 Down Down Mid	38	0.96236727			
Time 3 Down Down Down	39	0.95217151			

EXHIBIT 8.15 Continuously Compounded Zero-Coupon Yields

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.2999%	0.8009%	1.2951%	1.7913%
Time 1 Up	1	0.7145%	1.0124%	1.3597%	
Time 1 Mid	2	1.7922%	2.4980%	3.1644%	
Time 1 Down	3	1.3519%	1.8358%	2.3271%	
Time 2 Up Up	4	0.5976%	0.8718%		
Time 2 Up Mid	5	1.8966%	2.3936%		
Time 2 Up Down	6	1.3756%	1.7378%		
Time 2 Mid Up	7	1.9836%	2.9283%		
Time 2 Mid Mid	8	2.8568%	3.3648%		
Time 2 Mid Down	9	3.9945%	4.5650%		
Time 2 Down Up	10	1.0996%	1.8925%		
Time 2 Down Mid	11	1.9727%	2.3290%		
Time 2 Down Down	12	3.1104%	3.5291%		
Time 3 Up Up Up	13	0.4333%			
Time 3 Up Up Mid	14	1.7323%			
Time 3 Up Up Down	15	1.2113%			
Time 3 Up Mid Up	16	1.6704%			
Time 3 Up Mid Mid	17	2.5435%			
Time 3 Up Mid Down	18	3.6812%			
Time 3 Up Down Up	19	0.8799%			
Time 3 Up Down Mid	20	1.7530%			
Time 3 Up Down Down	21	2.8907%			
Time 3 Mid Up Up	22	1.7430%			
Time 3 Mid Up Mid	23	3.9911%			
Time 3 Mid Up Down	24	4.8955%			
Time 3 Mid Mid Up	25	1.7430%			
Time 3 Mid Mid Mid	26	3.9911%			
Time 3 Mid Mid Down	27	4.8955%			
Time 3 Mid Down Up	28	3.5917%			
Time 3 Mid Down Mid	29	4.3212%			
Time 3 Mid Down Down	30	6.3292%			
Time 3 Down Up Up	31	1.4652%			
Time 3 Down Up Mid	32	2.3383%			
Time 3 Down Up Down	33	3.4760%			
Time 3 Down Mid Up	34	1.4652%			
Time 3 Down Mid Mid	35	2.3383%			
Time 3 Down Mid Down	36	3.4760%			
Time 3 Down Down Up	37	2.1784%			
Time 3 Down Down Mid	38	3.8359%			
Time 3 Down Down Down	39	4.9010%			




EXHIBIT 8.16 Kamakura Corporation HJM Zero-Coupon Yield Curve Movements Two-Factor Empirical Volatilities, 1962–2011




EXHIBIT 8.17 Kamakura Corporation HJM Zero-Coupon Yield Curve Movements Two-Factor Empirical Volatility, 1962–2011

When taking expected values, we can calculate the probability of each scenario coming about since the probabilities of an upshift, midshift, and downshift are $1/4$, $1/4$, and $1/2$ as shown in Exhibit 8.20.

It is again convenient to calculate the probability-weighted discount factors for use in calculating the expected present value of cash flows as shown in Exhibit 8.21.

We now use the HJM bushy trees we have generated to value representative securities.

EXHIBIT 8.18 One-Year Forward Rates, Simple Interest

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.3003%	1.3104%	2.3098%	3.3343%
Time 1 Up	1	0.7171%	1.3189%	2.0756%	
Time 1 Mid	2	1.8083%	3.2556%	4.5999%	
Time 1 Down	3	1.3611%	2.3468%	3.3650%	
Time 2 Up Up	4	0.5994%	1.1526%		
Time 2 Up Mid	5	1.9147%	2.9327%		
Time 2 Up Down	6	1.3851%	2.1222%		
Time 2 Mid Up	7	2.0034%	3.9489%		
Time 2 Mid Mid	8	2.8979%	3.9489%		
Time 2 Mid Down	9	4.0753%	5.2696%		
Time 2 Down Up	10	1.1057%	2.7217%		
Time 2 Down Mid	11	1.9923%	2.7217%		
Time 2 Down Down	12	3.1593%	4.0268%		
Time 3 Up Up Up	13	0.4342%			
Time 3 Up Up Mid	14	1.7474%			
Time 3 Up Up Down	15	1.2187%			
Time 3 Up Mid Up	16	1.6844%			
Time 3 Up Mid Mid	17	2.5761%			
Time 3 Up Mid Down	18	3.7498%			
Time 3 Up Down Up	19	0.8838%			
Time 3 Up Down Mid	20	1.7685%			
Time 3 Up Down Down	21	2.9329%			
Time 3 Mid Up Up	22	1.7582%			
Time 3 Mid Up Mid	23	4.0718%			
Time 3 Mid Up Down	24	5.0173%			
Time 3 Mid Mid Up	25	1.7582%			
Time 3 Mid Mid Mid	26	4.0718%			
Time 3 Mid Mid Down	27	5.0173%			
Time 3 Mid Down Up	28	3.6570%			
Time 3 Mid Down Mid	29	4.4159%			
Time 3 Mid Down Down	30	6.5338%			
Time 3 Down Up Up	31	1.4760%			
Time 3 Down Up Mid	32	2.3659%			
Time 3 Down Up Down	33	3.5371%			
Time 3 Down Mid Up	34	1.4760%			
Time 3 Down Mid Mid	35	2.3659%			
Time 3 Down Mid Down	36	3.5371%			
Time 3 Down Down Up	37	2.2023%			
Time 3 Down Down Mid	38	3.9104%			
Time 3 Down Down Down	39	5.0231%			

EXHIBIT 8.19 Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1	0.997005786	0.989907593	0.984009757	0.979755294
2		0.997005786	0.989907593	0.984009757	0.967110174
3		0.997005786	0.989907593	0.984009757	0.972162025
4			0.979296826	0.971309727	0.955219645
5			0.979296826	0.971309727	0.946915774
6			0.979296826	0.971309727	0.93620374
7			0.98361809	0.976383515	0.967830059
8			0.98361809	0.976383515	0.959416563
9			0.98361809	0.976383515	0.948563113
10				0.960062532	0.943473945
11				0.960062532	0.922499895
12				0.960062532	0.914194906
13				0.951716561	0.935272181
14				0.951716561	0.914480461
15				0.951716561	0.906247669
16				0.940950217	0.907753518
17				0.940950217	0.901156195
18				0.940950217	0.883240853
19				0.972861584	0.958711053
20				0.972861584	0.950376831
21				0.972861584	0.939625643
22				0.964404349	0.950376831
23				0.964404349	0.942115059
24				0.964404349	0.931457333
25				0.953494474	0.932948155
26				0.953494474	0.917611875
27				0.953494474	0.907890273

VALUATION OF A ZERO-COUPON BOND MATURING AT TIME $T = 4$

A riskless zero-coupon bond pays \$1 in each of the 27 nodes of the bushy tree that prevail at time $T = 4$ as shown in Exhibit 8.22.

When we multiply this vector of 1s times the probability-weighted discount factors in the time $T = 4$ column in the previous table in Exhibit 8.22 and add them, we get a zero-coupon bond price of 0.93085510, which is the value we should get in a no-arbitrage economy, the value observable in the market and used as an input to create the tree.

EXHIBIT 8.20 Probability of Each State

Row Number	Current Time				
	0	1	2	3	4
1	100.0000%	25.0000%	6.2500%	1.5625%	1.5625%
2		25.0000%	6.2500%	1.5625%	1.5625%
3		50.0000%	12.5000%	3.1250%	3.1250%
4			6.2500%	1.5625%	1.5625%
5			6.2500%	1.5625%	1.5625%
6			12.5000%	3.1250%	3.1250%
7			12.5000%	3.1250%	3.1250%
8			12.5000%	3.1250%	3.1250%
9			25.0000%	6.2500%	6.2500%
10				1.5625%	1.5625%
11				1.5625%	1.5625%
12				3.1250%	3.1250%
13				1.5625%	1.5625%
14				1.5625%	1.5625%
15				3.1250%	3.1250%
16				3.1250%	3.1250%
17				3.1250%	3.1250%
18				6.2500%	6.2500%
19				3.1250%	3.1250%
20				3.1250%	3.1250%
21				6.2500%	6.2500%
22				3.1250%	3.1250%
23				3.1250%	3.1250%
24				6.2500%	6.2500%
25				6.2500%	6.2500%
26				6.2500%	6.2500%
27				12.5000%	12.5000%
Total	100.0000%	100.0000%	100.0000%	100.0000%	100.0000%

VALUATION OF A COUPON-BEARING BOND PAYING ANNUAL INTEREST

Next we value a bond with no credit risk that pays \$3 in interest at every scenario at times $T = 1, 2, 3$, and 4 plus principal of 100 at time $T = 4$. The valuation is calculated by multiplying each cash flow by the matching probability-weighted discount factor, to get a value of 104.70709974. It will again surprise many that this is the same value that we arrived at in Chapters 6 and 7, even though the volatilities used and the number of risk factors used are different. The values are the same because, by construction, our valuations for the zero-coupon bond prices at time zero for maturities at $T = 1, 2, 3$, and 4 continue to match the inputs. Multiplying these zero-coupon bond prices times 3, 3, 3, and 103 also leads to a value of 104.70709974 as it should as shown in Exhibit 8.23.

EXHIBIT 8.21 Probability-Weighted Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1	0.249251447	0.061869225	0.015375152	0.015308676
2		0.249251447	0.061869225	0.015375152	0.015111096
3		0.498502893	0.123738449	0.030750305	0.030380063
4			0.061206052	0.015176714	0.014925307
5			0.061206052	0.015176714	0.014795559
6			0.122412103	0.030353429	0.029256367
7			0.122952261	0.030511985	0.030244689
8			0.122952261	0.030511985	0.029981768
9			0.245904522	0.06102397	0.059285195
10				0.015000977	0.01474178
11				0.015000977	0.014414061
12				0.030001954	0.028568591
13				0.014870571	0.014613628
14				0.014870571	0.014288757
15				0.029741143	0.02832024
16				0.029404694	0.028367297
17				0.029404694	0.028161131
18				0.058809389	0.055202553
19				0.030401924	0.02995972
20				0.030401924	0.029699276
21				0.060803849	0.058726603
22				0.030137636	0.029699276
23				0.030137636	0.029441096
24				0.060275272	0.058216083
25				0.059593405	0.05830926
26				0.059593405	0.057350742
27				0.119186809	0.113486284

VALUATION OF A DIGITAL OPTION ON THE ONE-YEAR U.S. TREASURY RATE

Now, we value a digital option that pays \$1 at time $T = 3$ if (at that time) the one-year U.S. Treasury rate (for maturity at $T = 4$) is over 4 percent. If we look at the table of the term structure of one-year spot rates over time, this happens at the seven shaded scenarios out of 27 possibilities at time $t = 3$ as shown in Exhibit 8.24.

The evolution of the spot rate can be displayed graphically as in Exhibit 8.25.

The cash flow payoffs in the seven relevant scenarios can be input in Exhibit 8.26 and multiplied by the probability weighted discount factors to find that this option has a value of 0.29701554 as shown in the exhibit.

EXHIBIT 8.22 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.00	0.00	0.00	0.00	1.00
2		0.00	0.00	0.00	1.00
3		0.00	0.00	0.00	1.00
4			0.00	0.00	1.00
5			0.00	0.00	1.00
6			0.00	0.00	1.00
7			0.00	0.00	1.00
8			0.00	0.00	1.00
9			0.00	0.00	1.00
10				0.00	1.00
11				0.00	1.00
12				0.00	1.00
13				0.00	1.00
14				0.00	1.00
15				0.00	1.00
16				0.00	1.00
17				0.00	1.00
18				0.00	1.00
19				0.00	1.00
20				0.00	1.00
21				0.00	1.00
22				0.00	1.00
23				0.00	1.00
24				0.00	1.00
25				0.00	1.00
26				0.00	1.00
27				0.00	1.00
Risk-Neutral Value =					0.93085510

REPLICATION OF HJM EXAMPLE 3 IN COMMON SPREADSHEET SOFTWARE

The readers of this book are almost certain to be proficient in the use of common spreadsheet software. In replicating the examples in this and other chapters, any discrepancy in values within the first 10 decimal places is most likely an error by the analyst or, heaven forbid, the authors. In the construction of the examples in this book, every discrepancy in the first 10 decimal places has been an error, not the normal computer science rounding error that the authors would typically blame first as the cause.

EXHIBIT 8.23 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.00	3.00	3.00	3.00	103.00
2		3.00	3.00	3.00	103.00
3		3.00	3.00	3.00	103.00
4			3.00	3.00	103.00
5			3.00	3.00	103.00
6			3.00	3.00	103.00
7			3.00	3.00	103.00
8			3.00	3.00	103.00
9			3.00	3.00	103.00
10				3.00	103.00
11				3.00	103.00
12				3.00	103.00
13				3.00	103.00
14				3.00	103.00
15				3.00	103.00
16				3.00	103.00
17				3.00	103.00
18				3.00	103.00
19				3.00	103.00
20				3.00	103.00
21				3.00	103.00
22				3.00	103.00
23				3.00	103.00
24				3.00	103.00
25				3.00	103.00
26				3.00	103.00
27				3.00	103.00
Risk-Neutral Value =					104.70709974

EXHIBIT 8.24 Spot Rate Process

State	Row Number	0	1	2	3
Up Up Up	1	0.3003%	0.7145%	0.5994%	0.4342%
Up Up Mid	2		1.7922%	1.9147%	1.7474%
Up Up Down	3		1.3519%	1.3851%	1.2187%
Up Mid Up	4			2.0034%	1.6844%
Up Mid Mid	5			2.8979%	2.5761%
Up Mid Down	6			4.0753%	3.7498%

(Continued)

EXHIBIT 8.24 (Continued)

State	Row Number	0	1	2	3
Up Down Up	7			1.1057%	0.8838%
Up Down Mid	8			1.9923%	1.7685%
Up Down Down	9			3.1593%	2.9329%
Mid Up Up	10				1.7582%
Mid Up Mid	11				4.0718%
Mid Up Down	12				5.0173%
Mid Mid Up	13				1.7582%
Mid Mid Mid	14				4.0718%
Mid Mid Down	15				5.0173%
Mid Down Up	16				3.6570%
Mid Down Mid	17				4.4159%
Mid Down Down	18				6.5338%
Down Up Up	19				1.4760%
Down Up Mid	20				2.3659%
Down Up Down	21				3.5371%
Down Mid Up	22				1.4760%
Down Mid Mid	23				2.3659%
Down Mid Down	24				3.5371%
Down Down Up	25				2.2023%
Down Down Mid	26				3.9104%
Down Down Down	27				5.0231%


**EXHIBIT 8.25** Evolution of One-Year U.S. Treasury Spot Rate Two-Factor HJM Model with Rate and Maturity Dependent Volatility

EXHIBIT 8.26 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.00	0.00	0.00	0.00	0.00
2		0.00	0.00	0.00	0.00
3		0.00	0.00	0.00	0.00
4			0.00	0.00	0.00
5			0.00	0.00	0.00
6			0.00	0.00	0.00
7			0.00	0.00	0.00
8			0.00	0.00	0.00
9			0.00	0.00	0.00
10				0.00	0.00
11				1.00	0.00
12				1.00	0.00
13				0.00	0.00
14				1.00	0.00
15				1.00	0.00
16				0.00	0.00
17				1.00	0.00
18				1.00	0.00
19				0.00	0.00
20				0.00	0.00
21				0.00	0.00
22				0.00	0.00
23				0.00	0.00
24				0.00	0.00
25				0.00	0.00
26				0.00	0.00
27				1.00	0.00
Risk-Neutral Value =					0.29701554

CONCLUSION

This chapter, continuing the examples of Chapters 6 and 7, shows how to simulate zero-coupon bond prices, forward rates, and zero-coupon bond yields in an HJM framework with two risk factors and rate-dependent and maturity-dependent interest rate volatility. The results show a rich twist in simulated yield curves and a pull of rates upward from a very low rate environment. Monte Carlo simulation, an alternative to the bushy tree framework, can be done in a fully consistent way. We discuss that in detail in Chapter 10.

In the next chapter, we introduce a third risk factor to further advance the realism of the model.

HJM Interest Rate Modeling with Three Risk Factors¹

In Chapters 6 through 8, we provided worked examples of how to use the yield curve simulation framework of Heath, Jarrow, and Morton (HJM) using two different assumptions about the volatility of forward rates and one or two independent risk factors. The first volatility assumption was that volatility was dependent on the maturity of the forward rate and nothing else. The second volatility assumption was that the volatility of forward rates was dependent on both the level of rates and the maturity of the forward rates being modeled. The first two models were one-factor models, implying that random rate shifts are either all positive, all negative, or zero. Our third example postulated that the change in one-year spot rates, along with an unspecified second risk factor, was in fact driving interest rate movements. In this chapter, we generalize the model to include three risk factors in order to increase further the realism of the simulated yield curve. Consistent with the research of Dickler, Jarrow, and van Deventer cited in Chapters 6, 7, and 8, we explain why even three risk factors understate the number of risk factors driving the U.S. Treasury yield curve.

PROBABILITY OF YIELD CURVE TWISTS IN THE U.S. TREASURY MARKET

In Chapter 8, we explained that one-factor models imply no twists in the yield curve. Random yield shocks drive yields either all up or all down together. Yield curve twists, where some rates rise and others fall, are assumed never to occur. Yet during 12,386 days of movements in U.S. Treasury forward rates, yield curve twists occurred on 94.3 percent of the observations (see Chapter 8). We now ask a simple question. To what extent did the two-factor model of Chapter 8 succeed in explaining the correlations between changes in forward rates? We answer that question by analyzing the correlation between the residuals of the regression we used to determine the impact of our first risk factor, the annual change in the one-year spot U.S. Treasury rate, on changes in forward rates at various maturities:

$$\begin{aligned}
 [\text{Change in } f_k(i)] & \\
 &= a \\
 &+ b[\text{Change in one-year U.S. Treasury spot rate}(i)] \\
 &+ e_k(i)
 \end{aligned}$$

If the two-factor model is correct, all forward rate movements would be explained by our two factors: the one-year change in the one-year spot rate and the unnamed second factor. What would the correlations remaining after use of the first risk factor show if our hypothesis was correct? The correlation of the residuals at any two maturities would show a 100 percent correlation, with all residuals either moving up together or down together after eliminating the common impact of changes in the one-year spot rate. Unfortunately, the correlation of the residuals shows that the two-factor model is not rich enough to achieve that objective (Exhibit 9.1).

The correlations are positive and far from 100 percent, which leads to an immediate conclusion. We need at least one more risk factor to realistically mimic real yield curve movements. We incorporate a third factor in this chapter for that reason.

OBJECTIVES OF THE EXAMPLE AND KEY INPUT DATA

Following Jarrow (2002), we make the same modeling assumptions for our worked example as in Chapters 6 through 8:

- Zero-coupon bond prices for the U.S. Treasury curve on March 31, 2011, are the basic inputs.
- Interest rate volatility assumptions are based on the Dickler, Jarrow, and van Deventer papers on daily U.S. Treasury yields and forward rates from 1962 to 2011. In this chapter, we again retain the volatility assumptions used in Chapter 8 but expand the number of random risk factors driving interest rates to three factors.
- The modeling period is four equal length periods of one year each.
- The HJM implementation is that of a “bushy tree,” which we describe next.

EXHIBIT 9.1 Correlation of Residuals from Regression 1 on U.S. Treasury Spot and Forward Rates, 1962 to 2011

Maturity in Year	U.S. Treasury Forward Rate Maturing in Year							
	2	3	4	5	6	7	8	9
2	1							
3	0.920156	1						
4	0.840913	0.936176	1					
5	0.807619	0.823928	0.942743	1				
6	0.790989	0.808105	0.875853	0.950392	1			
7	0.753103	0.803682	0.820229	0.859572	0.967012	1		
8	0.737788	0.806994	0.833124	0.838943	0.910462	0.961099	1	
9	0.677499	0.741696	0.808673	0.796858	0.786531	0.816912	0.93905	1

With two factors, the bushy tree has upshifts, midshifts, and downshifts at each node in the tree. With three risk factors, there are four branches at each node of the tree. For simplicity and consistency with an n -factor HJM implementation, we abandon the terms up and down and instead label the shifts Shift 1, Shift 2, Shift 3, and Shift 4. The bushy tree starts with one branch at time zero and moves to 4 at time 1, 16 at time 2, and 64 at time 3. We split the tree into its upper and lower halves for better visibility (Exhibit 9.2 and Exhibit 9.3).

At each of the points in time on the lattice (time 0, 1, 2, 3, and 4) there are sets of zero-coupon bond prices and forward rates. At time 0, there is one set of data. At time one, there are four sets of data: the Shift 1 set, Shift 2 set, Shift 3 set, and Shift 4 set. At time two, there are 16 sets of data, and at time three there are $64 = 4^3$ sets of data.

As shown in Chapters 6 through 8, volatilities of the U.S. Treasury one-year spot rate and one-year forward rates with maturities in years 2, 3, . . . , 10 depend dramatically on the starting level of the one-year U.S. Treasury spot rate. As before, we use the same volatility table as in Chapter 8 to measure total volatility, measured as the standard deviation of one-year changes in continuously compounded spot and forward rates from 1963 to 2011. We use this table later to divide total volatility between three uncorrelated risk factors.

We will again use the zero-coupon bond prices prevailing on March 31, 2011, as our other inputs, as we have done beginning in Chapter 6. We now partition total interest rate volatility among the three risk factors.

RISK FACTOR 1: ANNUAL CHANGES IN THE ONE-YEAR U.S. TREASURY SPOT RATE

In Chapters 6 and 7, the nature of the single-factor shocking one-year spot and forward rates was not specified. In Chapter 8, we postulate that the first of the three factors driving changes in forward rates is the change in the one-year spot rate of interest. For each of the one-year U.S. Treasury forward rates $f_k(t)$, we run the regression

$$\begin{aligned} [\text{Change in } f_k(i)] \\ &= a \\ &+ b[\text{Change in one-year U.S. Treasury spot rate } (i)] \\ &+ e_k(i) \end{aligned}$$

where the change in continuously compounded yields is measured over annual intervals from 1963 to 2011. We add an additional assumption that the second risk factor is the one-year change in the one-year forward rate maturing in year 10. To be consistent with the HJM assumptions that the individual risk factors are uncorrelated, we estimate the incremental explanatory power of this factor by adding the one-year change in the one-year forward rate maturing in year 10 as an explanatory variable in a regression where the dependent variable is the appropriate series of residuals from the first set of regressions, which we label the Regression 1 set.




EXHIBIT 9.2 Example of Bushy Tree with Three Risk Factors for HJM Modeling of No-Arbitrage Forward Rate Movements




EXHIBIT 9.3 Example of Bushy Tree with Three Risk Factors for HJM Modeling of No-Arbitrage Forward Rate Movements

$$e_k(i) = a + b[\text{change in forward maturing in year } 10 \text{ } (i)] + e_{2k}(i)$$

Because of the nature of the linear regression of changes in one-year forward rates on the first risk factor (changes in the one-year spot U.S. Treasury rate), we know that risk factor 1 and the residuals $e_k(i)$ are uncorrelated. The regression above will pick up the residual influence of the one-year changes in the one-year forward rate maturing in year 10. There are many econometric approaches to this process that are more elegant, but we use this simple approach for clarity of exposition.

Our hypothesis that the one-year change in the one-year forward rate maturing in year 10 was important was right on the mark. The coefficients, shown in Exhibit 9.4, are very statistically significant with t -statistics ranging from 68 to 175. This graph shows that the impact of the change of the one-year forward rate maturing in year 10 generally increases for rates maturing closest to year 10, as one would expect.

The third risk factor is “everything else” not explained by the changes in the one-year spot rate or the changes in the one-year forward rate maturing in year 10. Because of the sequential nature of the Regression 1 set and the Regression 2 set (which means correlation among the risk factors is zero), total volatility for the forward rate maturing in $T - t = k$ years can be divided as follows between the three risk factors:

$$\sigma_{k,\text{total}}^2 = \sigma_{k,1}^2 + \sigma_{k,2}^2 + \sigma_{k,3}^2$$

We also know that the risk contribution of the first risk factor, the change in the spot one-year U.S. treasury rate, is proportional to the regression coefficient α_k of




EXHIBIT 9.4 Percent Response by U.S. Treasury One-Year Forward Rates to 1 Percent Shift in One-Year U.S. Treasury Forward Rate Maturing in Year 10, 1962–2011

changes in forward rate maturing in year k on changes in the one-year spot rate. Because Regression 2 set is done on the residuals of Regression 1 set, the risk contribution of risk factor 2 (the change in the one-year forward rate maturing in year 10) is also proportional to the relevant regression coefficient. Note again that, by construction, risk factor 2 is not correlated with risk factor 1. We denote the total volatility of the one-year U.S. Treasury spot rate by the subscript $_{1,\text{total}}$ and the total volatility of the one-year forward rate maturing in year 10 by $_{10,\text{total}}$. Then the contribution of risk factors 1 and 2 to the volatility of the one-year forward rate maturing in year k is

$$\begin{aligned}\sigma_{k,1}^2 &= \alpha_{1k}^2 \sigma_{1,\text{total}}^2 \\ \sigma_{k,2}^2 &= \alpha_{10k}^2 \sigma_{10,\text{total}}^2\end{aligned}$$

This allows us to solve for the volatility of risk factor 3 using this equation:

$$\sigma_{k,3} = [\sigma_{k,\text{total}}^2 - \alpha_{1k}^2 \sigma_{1,\text{total}}^2 - \alpha_{10k}^2 \sigma_{10,\text{total}}^2]^{1/2}$$

If ever the quantity in brackets is negative, $\sigma_{k,3}$ is set to zero and $\sigma_{k,2}$ is set to the value, which makes the quantity in brackets exactly zero. Because the total volatility for each forward rate varies by the level of the one-year spot U.S. Treasury spot rate, so will the values of $\sigma_{k,1}$, $\sigma_{k,2}$, and $\sigma_{k,3}$. The regression coefficients used in this process are summarized in Exhibit 9.5.

Using the equations above, the look-up table for risk factor 1 (changes in the one-year spot rate) volatility is given in Exhibit 9.6.

The look-up table for risk factor 2 (changes in the one-year forward rate maturing in year 10) is given in Exhibit 9.7.

The look-up table for the general “all other” risk factor 3 is shown in Exhibit 9.8.

EXHIBIT 9.5 Regression Coefficients

Dependent Variable	Coefficient b_1 on	Standard Error of Regression 1	Coefficient b_2 of	Standard Error of Regression 2
	Change in U.S. Treasury Rate		Residuals in Regression 1 on Change in 10-Year U.S. Treasury Rate	
Forward 2	0.77428864	0.654783%	0.29955296	0.556380%
Forward 3	0.63607642	0.816101%	0.40401267	0.670241%
Forward 4	0.53978035	0.851190%	0.48360419	0.643346%
Forward 5	0.43988300	0.933073%	0.52091272	0.714295%
Forward 6	0.35338394	0.956797%	0.48643906	0.775357%
Forward 7	0.30957977	0.958593%	0.49125712	0.773537%
Forward 8	0.30904919	0.914843%	0.57964967	0.625036%
Forward 9	0.32251308	0.936120%	0.68828249	0.497098%
Forward 10	0.32736771	1.003651%		

EXHIBIT 9.6 Standard Deviation of Risk Factor 1, Sigma 1

Data Group	Rate in Data Group		Number of Observations	Standard Deviation of Risk Factor 1, Sigma 1			
	Minimum Spot Rate	Maximum Spot Rate		1-Year	2-Years	3-Years	4-Years
1	0.00%	0.09%	Assumed	0.01000000%	0.00774289%	0.00492507%	0.00265845%
2	0.09%	0.25%	Assumed	0.04000000%	0.03097155%	0.01970027%	0.01063382%
3	0.25%	0.50%	346	0.06363856%	0.04927462%	0.03134242%	0.01691802%
4	0.50%	0.75%	86	0.08298899%	0.06425743%	0.04087264%	0.02206225%
5	0.75%	1.00%	20	0.94371413%	0.73070713%	0.46478558%	0.25088212%
6	1.00%	2.00%	580	1.01658509%	0.78713029%	0.63127387%	0.34074924%
7	2.00%	3.00%	614	1.30986113%	1.01421060%	0.64511544%	0.34822064%
8	3.00%	4.00%	1492	1.22296401%	0.94692714%	0.60231802%	0.32511944%
9	4.00%	5.00%	1672	1.53223304%	1.18639064%	0.75463511%	0.40733721%
10	5.00%	6.00%	2411	1.22653783%	0.94969431%	0.60407816%	0.32606952%
11	6.00%	7.00%	1482	1.67496041%	1.29690282%	0.82492930%	0.44528063%
12	7.00%	8.00%	1218	1.70521091%	1.32032544%	0.83982788%	0.45332259%
13	8.00%	9.00%	765	2.16438588%	1.67585941%	1.06597465%	0.57539217%
14	9.00%	10.00%	500	2.28886301%	1.77224063%	1.12728047%	0.60848385%
15	10.00%	16.60%	959	2.86758369%	2.22033748%	1.41230431%	0.76233412%

EXHIBIT 9.7 Standard Deviation of Risk Factor 2, Sigma 2

Data Group	Rate in Data Group		Number of Observations	Standard Deviation of Risk Factor 2, Sigma 2		
	Minimum Spot Rate	Maximum Spot Rate		2-Years	3-Years	4-Years
1	0.00%	0.09%	Assumed	0.01844038%	0.09987864%	0.14508126%
2	0.09%	0.25%	Assumed	0.09508293%	0.14140444%	0.16926147%
3	0.25%	0.50%	346	0.11460958%	0.15457608%	0.18502795%
4	0.50%	0.75%	86	0.21277322%	0.28697122%	0.34350527%
5	0.75%	1.00%	20	0.30003119%	0.40465767%	0.48437625%
6	1.00%	2.00%	580	0.28241222%	0.00000000%	0.45593183%
7	2.00%	3.00%	614	0.13927096%	0.18783734%	0.22484177%
8	3.00%	4.00%	1492	0.22509440%	0.30358902%	0.36339682%
9	4.00%	5.00%	1672	0.22361237%	0.30159018%	0.36100420%
10	5.00%	6.00%	2411	0.28294579%	0.38161427%	0.45679324%
11	6.00%	7.00%	1482	0.29059766%	0.39193450%	0.46914658%
12	7.00%	8.00%	1218	0.26551545%	0.47400436%	0.56738441%
13	8.00%	9.00%	765	0.46585233%	0.62830374%	0.75208116%
14	9.00%	10.00%	500	0.58305515%	0.78637737%	0.94129569%
15	10.00%	16.60%	959	0.58088433%	0.78344955%	0.93779108%

EXHIBIT 9.8 Standard Deviation of Risk Factor 3, Sigma 3

Data Group	Rate in Data Group		Number of Observations	Standard Deviation of Risk Factor 3, Sigma 3			
	Minimum Spot Rate	Maximum Spot Rate		1-Year	2-Years	3-Years	4-Years
1	0.00%	0.09%	Assumed	0.00000000%	0.00000000%	0.00000000%	0.13763852%
2	0.09%	0.25%	Assumed	0.00000000%	0.00000000%	0.14005958%	0.24746207%
3	0.25%	0.50%	346	0.00000000%	0.36337194%	0.65135096%	0.84363313%
4	0.50%	0.75%	86	0.00000000%	0.40702120%	0.54595415%	0.73281962%
5	0.75%	1.00%	20	0.00000000%	0.74084910%	0.74635937%	0.50994308%
6	1.00%	2.00%	580	0.00000000%	0.12463857%	0.00000000%	0.33951493%
7	2.00%	3.00%	614	0.00000000%	0.78254406%	0.83631746%	0.68568072%
8	3.00%	4.00%	1492	0.00000000%	0.54106464%	0.77592194%	0.76181187%
9	4.00%	5.00%	1672	0.00000000%	0.12847478%	0.58968940%	0.63313043%
10	5.00%	6.00%	2411	0.00000000%	0.45562957%	0.65652394%	0.72628160%
11	6.00%	7.00%	1482	0.00000000%	0.46346642%	0.77892862%	0.82112196%
12	7.00%	8.00%	1218	0.00000000%	0.00000000%	0.87713028%	1.06830036%
13	8.00%	9.00%	765	0.00000000%	0.35358489%	1.20123780%	1.40466564%
14	9.00%	10.00%	500	0.00000000%	1.05723235%	1.49203773%	1.68880657%
15	10.00%	16.60%	959	0.00000000%	1.30610213%	1.98193330%	1.98085391%

When we graph the relative contributions of each risk factor to the volatility of the forward rate maturing in year 2, we see an interesting pattern as the level of the one-year spot Treasury rate rises (Exhibit 9.9).

The contribution of risk factor 1 (the second-highest line) rises steadily as the level of interest rates rises. The impact of risk factor 2 (the smoother line near the bottom of the graph) rises more gradually as the level interest rates rises. The “all other risk factor,” risk factor 3, is volatile and spikes a number of times near the bottom of the graph, due in part to the spike in total volatility when the one-year spot U.S. Treasury rate is between 75 and 100 basis points. We see in the next section that this has an impact on our simulations and that this is a value for total volatility that one may want to override as a matter of professional judgment.

ALTERNATIVE SPECIFICATIONS OF THE INTEREST RATE VOLATILITY SURFACE

In Chapters 6 through 9, we have used a look-up table to specify the relevant levels of interest rate volatility. A common practice, sometimes used wisely and sometimes not, is to use the multivariate equivalent of Chapter 5’s yield curve smoothing techniques to create a smooth “interest rate volatility surface” for each of the three risk factors. Risk factor values would be chosen from this smoothed surface instead of the look-up table. This is a modest extension of the approach taken in these examples and we leave such an extension to the reader.



EXHIBIT 9.9 Total Volatility and Volatility for Each Risk Factor, One-Year Forward Maturing in Year 2

KEY IMPLICATIONS AND NOTATION OF THE HJM APPROACH

As in prior chapters, we confirm in this example that the zero-coupon bond valuations are 100 percent consistent with the zero-coupon bond price inputs, ensuring that there are no arbitrage opportunities. We now introduce the same notation first used in Chapter 6 with slight modifications:

Δ	= length of time period, which is 1 in this example
$r(t, s_t)$	= the simple one-period risk-free interest rate as of current time t in state s_t
$R(t, s_t)$	= $1 + r(t, s_t)$, the total return, the value of \$1 dollar invested at the risk-free rate for 1 period
$f(t, T, s_t)$	= the simple one-period forward rate maturing at time T in state s_t as of time t
$F(t, T, s_t)$	= $1 + f(t, T, s_t)$, the return on \$1 invested at the forward rate for one period
$\sigma_1(t, T, s_t)$	= forward rate volatility due to risk factor 1, the change in the spot one-year U.S. Treasury rate, at time t for the one-period forward rate maturing at T in state s_t (a sequence of ups and downs). Because σ_1 is dependent on the spot one-year U.S. Treasury rate, the state s_t is very important in determining its level.
$\sigma_2(t, T, s_t)$	= forward rate volatility due to risk factor 2, the change in the one-year forward rate maturing in 10 years, at time t for the one-period forward rate maturing at T in state s_t (a sequence of ups and downs). Because σ_2 is dependent on the spot one-year U.S. Treasury rate and its volatility, the state s_t is very important in determining its level.
$\sigma_3(t, T, s_t)$	= forward rate volatility due to risk factor 3, all volatility remaining after the impacts of risk factors 1 and 2. Because σ_3 is dependent on the spot one-year U.S. Treasury rate and its volatility, the state s_t is very important in determining its level.
$P(t, T, s_t)$	= zero-coupon bond price at time t maturing at time T in state s_t (i.e., up or down).
Index(1)	An indicator of the state for risk factor 1, defined later.
Index(2)	An indicator of the state for risk factor 2, defined later.
Index(3)	An indicator of the state for risk factor 3, defined later.
$K(i, t, T, s_t)$	= the weighted sum of the forward volatilities for risk factor i (either 1, 2 or 3) as of time t for the one-period forward rates from $t + \Delta$ to T , as shown here for $i = 1$:

$$K(1, t, T, s_t) = \Delta \sum_{j=t+\Delta}^{T-t} \sigma_1(t, j, s_t) \sqrt{\Delta}$$

For $i = 2$ or 3 , this expression has a negative sign:

$$K(i, t, T, s_t) = -\Delta \sum_{j=t+\Delta}^{T-t} \sigma_i(t, j, s_t) \sqrt{\Delta}$$

Note that this is a slightly different definition of K than we used in Chapters 6 through 8. Note also that in Chapter 8, for $i = 2$, we could have added a negative sign without changing the results of our analysis. For more than three factors, the positive sign on K for the first factor and negative signs for subsequent factors can be confirmed using the formulas that we discuss in Chapter 10, where we use the HJM framework for Monte Carlo simulation. As in Chapters 6 through 8, we again use the $\cosh(x)$ expression:

$$\text{Cosh}(x) = \frac{1}{2}[e^x + e^{-x}]$$

In the interests of saving space, we again arrange the bushy tree to look like a table by stretching the bushy tree as follows. Note that S-1 means shift 1, S-2 means shift 2, and so on. The “stretched” bushy tree is rearranged from the previous graph as shown in Exhibit 9.10.

A completely populated zero-coupon bond price tree would then be summarized like this; prices shown are for the zero-coupon bond price maturing at time $T = 4$ at times 0, 1, 2, and 3 (Exhibit 9.11).

EXHIBIT 9.10 Map of Sequence of States

Row Number	Current Time			
	0	1	2	3
1	Time 0	S-1	S-1, S-1	S-1, S-1, S-1
2		S-2	S-1, S-2	S-1, S-1, S-2
3		S-3	S-1, S-3	S-1, S-1, S-3
4		S-4	S-1, S-4	S-1, S-1, S-4
5			S-2, S-1	S-1, S-2, S-1
6			S-2, S-2	S-1, S-2, S-2
7			S-2, S-3	S-1, S-2, S-3
8			S-2, S-4	S-1, S-2, S-4
9			S-3, S-1	S-1, S-3, S-1
10			S-3, S-2	S-1, S-3, S-2
11			S-3, S-3	S-1, S-3, S-3
12			S-3, S-4	S-1, S-3, S-4
13			S-4, S-1	S-1, S-4, S-1
14			S-4, S-2	S-1, S-4, S-2
15			S-4, S-3	S-1, S-4, S-3
16			S-4, S-4	S-1, S-4, S-4
17				S-2, S-1, S-1
18				S-2, S-1, S-2
19				S-2, S-1, S-3
20				S-2, S-1, S-4
21				S-2, S-2, S-1
22				S-2, S-2, S-2

EXHIBIT 9.10 (Continued)

Row Number	Current Time			
	0	1	2	3
23				S-2, S-2, S-3
24				S-2, S-2, S-4
25				S-2, S-3, S-1
26				S-2, S-3, S-2
27				S-2, S-3, S-3
28				S-2, S-3, S-4
29				S-2, S-4, S-1
30				S-2, S-4, S-2
31				S-2, S-4, S-3
32				S-2, S-4, S-4
33				S-3, S-1, S-1
34				S-3, S-1, S-2
35				S-3, S-1, S-3
36				S-3, S-1, S-4
37				S-3, S-2, S-1
38				S-3, S-2, S-2
39				S-3, S-2, S-3
40				S-3, S-2, S-4
41				S-3, S-3, S-1
42				S-3, S-3, S-2
43				S-3, S-3, S-3
44				S-3, S-3, S-4
45				S-3, S-4, S-1
46				S-3, S-4, S-2
47				S-3, S-4, S-3
48				S-3, S-4, S-4
49				S-4, S-1, S-1
50				S-4, S-1, S-2
51				S-4, S-1, S-3
52				S-4, S-1, S-4
53				S-4, S-2, S-1
54				S-4, S-2, S-2
55				S-4, S-2, S-3
56				S-4, S-2, S-4
57				S-4, S-3, S-1
58				S-4, S-3, S-2
59				S-4, S-3, S-3
60				S-4, S-3, S-4
61				S-4, S-4, S-1
62				S-4, S-4, S-2
63				S-4, S-4, S-3
64				S-4, S-4, S-4

EXHIBIT 9.11 Four-Year, Zero-Coupon Bond Price

Row Number	Current Time			
	0	1	2	3
1	0.9308550992	0.8945353929	0.8951939726	0.9424127368
2		0.9635638078	0.9550797882	0.9472682532
3		0.9404125772	0.9332467863	0.9508321387
4		0.9325702085	0.8986172056	0.9256048053
5			0.9461410097	0.9698937713
6			0.9829031898	0.9747412918
7			0.9780746840	0.9801122898
8			0.9691466829	0.9609579282
9			0.9579764308	0.9503806226
10			0.9627643885	0.9805996987
11			0.9680693916	0.9691821954
12			0.9372422391	0.9478524225
13			0.9524667898	0.9485398381
14			0.9572272104	0.9534269227
15			0.9625017026	0.9570139788
16			0.9318518473	0.9316226298
17				0.9601375282
18				0.9906668429
19				0.9791321239
20				0.9575833729
21				0.9784392385
22				0.9944994413
23				0.9923910646
24				0.9881386595
25				0.9744393509
26				1.0037479721
27				0.9974134983
28				0.9787812762
29				0.9825458738
30				0.9874566295
31				0.9928976912
32				0.9734934642
33				0.9777040612
34				0.9825906175
35				0.9880048666
36				0.9686962602
37				0.9777040612
38				0.9825906175
39				0.9880048666
40				0.9686962602
41				0.9649211970
42				0.9939435367

EXHIBIT 9.11 (Continued)

Row Number	Current Time			
	0	1	2	3
43				0.9876709370
44				0.9692207112
45				0.9568733730
46				0.9872988979
47				0.9758033932
48				0.9543279009
49				0.9661716333
50				0.9968928133
51				0.9852856029
52				0.9636014261
53				0.9748191606
54				0.9796912981
55				0.9850895715
56				0.9658379388
57				0.9748191606
58				0.9796912981
59				0.9850895715
60				0.9658379388
61				0.9568821144
62				0.9778172454
63				0.9734710878
64				0.9521725485

In order to populate the trees with zero-coupon bond prices and forward rates, we again need to select the pseudo-probabilities.

PSEUDO-PROBABILITIES

We set the pseudo-probabilities in a manner consistent with Jarrow (2002), Chapter 7. Without loss of generality, we set the probability of shift 1 to one-eighth, the probability of shift 2 to one-eighth, the probability of shift 3 to one-fourth, and the probability of shift 4 to one-half. The no-arbitrage restrictions that stem from this set of pseudo-probabilities are given in a later section. We now demonstrate how to construct the bushy tree and use it for risk-neutral valuation.

The Formula for Zero-Coupon Bond Price Shifts with Three Risk Factors

Like the two-risk factor case, with three risk factors it is convenient to calculate the forward rates first and then derive zero-coupon bond prices from them. We do this

using slightly modified versions (developed with Jarrow) of equations 15.40 and 15.42 in Jarrow (2002, 297). We use this equation for the shift in forward rates:

$$F(t + \Delta, T, s_{t+\Delta}) = F(t, T, s_t) e^{[\mu(t, T, s_t)] \Delta^2} \\ \times e^{\Delta [\text{Index}(1) \sigma_1(t, T, s_t) \sqrt{\Delta} + \text{Index}(2) \sqrt{2} \sigma_2(t, T, s_t) \sqrt{\Delta} + \text{Index}(3) 2 \sigma_3(t, T, s_t) \sqrt{\Delta}]} \quad (9.1)$$

where when $T = t + \Delta$, the expression μ becomes

$$\mu(t, T) = \frac{1}{\Delta^2} \log \left[\frac{1}{2} e^{[K(1, t, t+\Delta)]} \left\langle \frac{1}{2} e^{[K(2, t, t+\Delta) \sqrt{2}]} \cosh(2K(3, t, t+\Delta)) \right. \right. \\ \left. \left. + \frac{1}{2} e^{-[K(2, t, t+\Delta) \sqrt{2}]} \right\rangle + \frac{1}{2} e^{-[K(1, t, t+\Delta)]} \right]$$

We use this expression to evaluate equation (9.1) when $T > t + \Delta$:

$$\mu(t, T) = \frac{1}{\Delta^2} \log \left[\frac{1}{2} e^{[K(1, t, t+\Delta)]} \left\langle \frac{1}{2} e^{[K(2, t, t+\Delta) \sqrt{2}]} \cosh(2K(3, t, t+\Delta)) \right. \right. \\ \left. \left. + \frac{1}{2} e^{-[K(2, t, t+\Delta) \sqrt{2}]} \right\rangle + \frac{1}{2} e^{-[K(1, t, t+\Delta)]} \right] - \sum_{j=t+\Delta}^{T-\Delta} \mu(t, j)$$

The values for the pseudo-probabilities plus Index(1), Index(2), and Index(3) are set as follows:

Type of Shift	Pseudo-Probability	Index(1)	Index(2)	Index(3)
Shift 1	12.50%	-1	1	1
Shift 2	12.50%	-1	1	-1
Shift 3	23.00%	-1	-1	0
Shift 4	50.00%	1	0	0

Building the Bushy Tree for Zero-Coupon Bonds Maturing at Time $T = 2$

We now populate the bushy tree for the two-year, zero-coupon bond. We calculate each element of equation (9.1). Note that, even as the bushy tree grows more complex, accuracy of valuation to 10 decimal places in common spreadsheet software is maintained. When $t = 0$ and $T = 2$, we know $\Delta = 1$ and

$$P(0, 2, s_t) = 0.9841101497$$

The one-period risk-free rate is again

$$R(0, s_t) = 1/P(0, 1, s_t) = 1/0.997005786 = 1.003003206$$

The one-period spot rate for U.S. Treasuries is $r(0, s_t) = R(0, s_t) - 1 = 0.3003206$ percent. At this level of the spot rate for U.S. Treasuries, volatilities for risk factors 1, 2, and 3 are selected from data group 3 in the previous look-up table. The volatilities for risk factor 1 for the one-year forward rates maturing in years 2, 3, and 4 are 0.000492746, 0.000313424, and 0.00016918. The volatilities for risk factor 2 for the one-year forward rates maturing in years 2, 3, and 4 are 0.001146096, 0.001545761, and 0.00185028. Finally, the volatilities for risk factor 3 are also taken from the appropriate risk factor table for data group 3: 0.003633719, 0.00651351, and 0.008436331.

The scaled sum of sigmas $K(1, t, T, s_t)$ for $t = 0$ and $T = 2$ becomes

$$K(1, t, T, s_t) = \Delta \sum_{j=t+\Delta}^{T-\Delta} \sigma_1(t, j, s_t) \sqrt{\Delta} = \Delta \sum_{j=1}^1 \sigma_1(0, j, s_t) \sqrt{\Delta}$$

and, therefore,

$$K(1, 0, T, s_t) = (\sqrt{1})(0.000492746) = 0.000492746$$

Similarly, remembering the minus sign,

$$K(2, 0, T, s_t) = -0.001146096$$

and

$$K(3, 0, T, s_t) = -0.003633719$$

We also can calculate that

$$\mu(t, t + \Delta) = 0.0000073730$$

Using equation (9.1) with these inputs and the fact that the variable $\text{Index}(1) = -1$, $\text{Index}(2) = 1$, and $\text{Index}(3) = 1$ for shift 1 gives us the forward returns for shift 1 of 1.021652722. For shifts 2, 3, and 4, the values of the three index variables change. Using these new index values, the forward returns in shifts 2, 3, and 4 are 1.006910523, 1.010972304, and 1.013610654. Note that the forward rates (as opposed to forward returns) are equal to these values minus 1. From the forward returns, we calculate the zero-coupon bond prices:

$$P(1, 2, s_t = \text{shift } 1) = 0.9788061814 = 1/F(1, 2, s_t = \text{shift } 1)$$

$$P(1, 2, s_t = \text{shift } 2) = 0.9931369048 = 1/F(1, 2, s_t = \text{shift } 2)$$

$$P(1, 2, s_t = \text{shift } 3) = 0.9891467808 = 1/F(1, 2, s_t = \text{shift } 3)$$

$$P(1, 2, s_t = \text{shift } 4) = 0.9865721079 = 1/F(1, 2, s_t = \text{shift } 4)$$

We have fully populated the bushy tree for the zero-coupon bond maturing at $T = 2$), since all of the four shifts at time $t = 2$ result in a riskless payoff of the zero-coupon bond at its face value of 1.

Row Number	Two-Year, Zero-Coupon Bond Price		
	0	1	2
1	0.9841101497	0.9788061814	1
2		0.9931369048	1
3		0.9891467808	1
4		0.9865721079	1

Building the Bushy Tree for Zero-Coupon Bonds Maturing at Time $T = 3$

For the zero-coupon bonds and one-period forward returns (which equal 1 plus the forward rate) maturing at time $T = 3$, we use the same volatilities listed above for risk factors 1, 2, and 3 to calculate:

$$K(1, 0, 3, s_t) = 0.000806170$$

Remembering the minus sign for risk factors 2 and 3 gives us

$$K(2, 0, 3, s_t) = -0.002691857$$

$$K(3, 0, 3, s_t) = -0.01047229$$

$$\mu(t, T) = 0.0000479070$$

The resulting forward returns for shifts 1, 2, 3, and 4 are 1.038505795, 1.011797959, 1.020593023, and 1.023467874. Zero-coupon bond prices are calculated from the one-period forward returns, so

$$P(1, 3, s_t = \text{shift 1 shift 1}) = 1/[F(1, 2, s_t = \text{shift 1})F(1, 3, s_t = \text{shift 1})]$$

The zero-coupon bond prices for maturity in time $T = 3$ for the four shifts are 0.9425139335, 0.9815565410, 0.9691882649, and 0.9639502449. We have now populated the second column of the zero-coupon bond price table for the zero-coupon bond maturing at $T = 3$ (Exhibit 9.12).

Building the Bushy Tree for Zero-Coupon Bonds Maturing at Time $T = 4$

We now populate the bushy tree for the zero-coupon bond maturing at $T = 4$. Using the same volatilities as before, we find that

$$K(1, 0, 4, s_t) = 0.000975351$$

EXHIBIT 9.12 Three-Year, Zero-Coupon Bond Price

Row Number	Current Time			
	0	1	2	3
1	0.9618922376	0.9425139335	0.9556681147	1
2		0.9815565410	0.9860553163	1
3		0.9691882649	0.9745742911	1
4		0.9639502449	0.9531258489	1
5			0.9779990715	1
6			0.9940520493	1
7			0.9919446211	1
8			0.9876941290	1
9			0.9811450109	1
10			0.9860487650	1
11			0.9914820691	1
12			0.9721055077	1
13			0.9783890415	1
14			0.9832790213	1
15			0.9886970636	1
16			0.9693749296	1

Again, remembering the minus sign, we calculate

$$K(2, 0, 4, s_t) = -0.004542136$$

$$K(3, 0, 4, s_t) = -0.018583560$$

$$\mu(t, T) = 0.0001272605$$

Using equation (9.1) with the correct values for Index(1), Index(2), and Index(3) leads to the following forward returns for shifts 1, 2, 3, and 4: 1.0536351507, 1.0186731103, 1.0305990034, and 1.0336489801. The zero-coupon bond price for shift 1 is calculated as follows:

$$P(1, 4, s_t = \text{shift } 1) = 1/[F(1, 2, s_t = \text{shift } 1)F(1, 3, s_t = \text{shift } 1)F(1, 4, s_t = \text{shift } 1)]$$

A similar calculation gives us the zero-coupon bond prices for all four shifts 1, 2, 3, and 4: 0.8945353929, 0.9635638078, 0.9404125772, and 0.9325702085.

We have now populated the column labeled 1 for $T = 1$ for the zero-coupon bond price maturing at time $T = 4$ (Exhibit 9.13).

Now we move to the third column, which displays the outcome of the $T = 4$ zero-coupon bond price after 16 scenarios: shift 1 followed by shifts 1, 2, 3, or 4; shift 2 followed by shifts 1, 2, 3, or 4 and so on. We calculate $P(2, 4, s_t = \text{shift } 1)$, $P(2, 4, s_t = \text{shift } 2)$, $P(2, 4, s_t = \text{shift } 3)$, and $P(2, 4, s_t = \text{shift } 4)$ after the initial state of shift 2

EXHIBIT 9.13 Four-Year, Zero-Coupon Bond Price

Row Number	Current Time			
	0	1	2	3
1	0.9308550992	0.8945353929	0.8951939726	0.9424127368
2		0.9635638078	0.9550797882	0.9472682532
3		0.9404125772	0.9332467863	0.9508321387
4		0.9325702085	0.8986172056	0.9256048053
5			0.9461410097	0.9698937713
6			0.9829031898	0.9747412918
7			0.9780746840	0.9801122898
8			0.9691466829	0.9609579282
9			0.9579764308	0.9503806226
10			0.9627643885	0.9805996987
11			0.9680693916	0.9691821954
12			0.9372422391	0.9478524225
13			0.9524667898	0.9485398381
14			0.9572272104	0.9534269227
15			0.9625017026	0.9570139788
16			0.9318518473	0.9316226298
17				0.9601375282
18				0.9906668429
19				0.9791321239
20				0.9575833729
21				0.9784392385
22				0.9944994413
23				0.9923910646
24				0.9881386595
25				0.9744393509
26				1.0037479721
27				0.9974134983
28				0.9787812762
29				0.9825458738
30				0.9874566295
31				0.9928976912
32				0.9734934642
33				0.9777040612
34				0.9825906175
35				0.9880048666
36				0.9686962602
37				0.9777040612
38				0.9825906175
39				0.9880048666
40				0.9686962602
41				0.9649211970
42				0.9939435367
43				0.9876709370

EXHIBIT 9.13 (Continued)

Row Number	Current Time			
	0	1	2	3
44				0.9692207112
45				0.9568733730
46				0.9872988979
47				0.9758033932
48				0.9543279009
49				0.9661716333
50				0.9968928133
51				0.9852856029
52				0.9636014261
53				0.9748191606
54				0.9796912981
55				0.9850895715
56				0.9658379388
57				0.9748191606
58				0.9796912981
59				0.9850895715
60				0.9658379388
61				0.9568821144
62				0.9778172454
63				0.9734710878
64				0.9521725485

as follows. After a shift = shift 2, the prevailing one-period, zero-coupon bond price at time $t = 1$ is 0.9931369048, which implies a one-period (one-year) U.S. Treasury spot rate of 0.0069105228 (i.e., 69 basis points). This is data group 4. When $t = 1$, $T = 4$, and $\Delta = 1$ then the volatilities for the two remaining one-period forward rates that are relevant are taken from the look-up table for data group 4 for risk factor 1: 0.0006425743 and 0.0004087264. For risk factor 2, the volatilities for the two remaining one-period forward rates are also chosen from data group 4: 0.0021277322, and 0.0028697122. Finally, the volatilities for risk factor 3 are also taken from data group 4: 0.0040702120 and 0.0054595415.

At time 1 in the shift 2 state, the zero-coupon bond prices for maturities at $T = 2$, 3, and 4 are 0.9931369048, 0.9815565410, and 0.9635638078. We make the intermediate calculations as above for the zero-coupon bond maturing at $T = 4$:

$$K(1, 1, 4, s_t) = 0.0010573007$$

As before, remembering the minus sign,

$$K(2, 1, 4, s_t) = -0.0049974444$$

$$K(3, 1, 4, s_t) = -0.0095297535$$

$$\mu(t, T) = 0.0000474532$$

We can calculate the values of the one-period forward return maturing at time $T = 4$ in shifts 1, 2, 3, and 4 as follows: 1.0336715790, 1.0113427849, 1.0141808568, and 1.0191379141. Similarly, using the appropriate intermediate calculations, we can calculate the forward returns for maturity at $T = 3$: 1.0224958583, 1.0059835405, 1.0081207950, and 1.0124591923. Since

$$P(2, 4, s_t = \text{shift 2 shift 1}) \\ = 1/[F(1, 3, s_t = \text{shift 2 shift 1})F(1, 4, s_t = \text{shift 2 shift 1})]$$

the zero-coupon bond prices for maturity at $T = 4$ after an initial shift 2 for subsequent shifts 1, 2, 3, and 4 are as follows:

0.9461410097
0.9829031898
0.9780746840
0.9691466829

We have correctly populated the fifth, sixth, seventh, and eighth rows of column 3 ($t = 2$) of the bushy tree above for the zero-coupon bond maturing at $T = 4$ (note values have been rounded for display only). The remaining calculations are left to the reader.

If we combine all of these tables, we can create a consolidated table of the term structure of zero-coupon bond prices in each scenario as in examples one, two, and three (Exhibit 9.14).

At any point in time t , the continuously compounded yield to maturity at time T can be calculated as $y(T - t) = -\ln[P(t, T)]/(T - t)$. Note that we have one incidence of negative rates (in state 46, shaded) on this bushy tree. This can be corrected by overriding the spike in volatility we discussed above in the data group for the one-year spot U.S. Treasury rate levels between 0.75 percent and 1.00 percent. Note that yield curve shifts are much more complex than in the three prior examples using one and two risk factors (Exhibit 9.15).

We can graph yield curve movements as shown next at time $t = 1$. We plot the original spot Treasury yield curve and then shifted yield curves for shifts 1, 2, 3, and 4. These shifts are relative to the one-period forward rates prevailing at time zero for maturity at times $T = 2, 3$, and 4. Because these one-period forward rates were much higher than yields as of time $t = 0$, all four shifts produce yields higher than time zero yields with the exception of one yield for maturity at $T = 4$ (Exhibit 9.16).

When we add 16 yield curves prevailing at time $t = 3$ and 64 single-point yield curves prevailing at time $t = 4$, two things are very clear. First, yield curve movements in a three-factor model are much more complex and much more realistic than what we saw in the one-factor and two-factor examples. Second, in the low-yield environment prevailing as of March 31, 2011, no arbitrage yield curve simulation shows “there is nowhere to go but up” from a yield curve perspective, with just a few exceptions (Exhibit 9.17).

Finally, we can display the one-year U.S. Treasury spot rates and the associated term structure of one-year forward rates in each scenario (Exhibit 9.18).

EXHIBIT 9.14 Zero-Coupon Bond Prices

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.9970057865	0.9841101497	0.9618922376	0.9308550992
Time 1 S-1	1	0.9788061814	0.9425139335	0.8945353929	
Time 1 S-2	2	0.9931369048	0.9815565410	0.9635638078	
Time 1 S-3	3	0.9891467808	0.9691882649	0.9404125772	
Time 1 S-4	4	0.9865721079	0.9639502449	0.9325702085	
Time 2 S-1, S-1	5	0.9556681147	0.8951939726		
Time 2 S-1, S-2	6	0.9860553163	0.9550797882		
Time 2 S-1, S-3	7	0.9745742911	0.9332467863		
Time 2 S-1, S-4	8	0.9531258489	0.8986172056		
Time 2 S-2, S-1	9	0.9779990715	0.9461410097		
Time 2 S-2, S-2	10	0.9940520493	0.9829031898		
Time 2 S-2, S-3	11	0.9919446211	0.9780746840		
Time 2 S-2, S-4	12	0.9876941290	0.9691466829		
Time 2 S-3, S-1	13	0.9811450109	0.9579764308		
Time 2 S-3, S-2	14	0.9860487650	0.9627643885		
Time 2 S-3, S-3	15	0.9914820691	0.9680693916		
Time 2 S-3, S-4	16	0.9721055077	0.9372422391		
Time 2 S-4, S-1	17	0.9783890415	0.9524667898		
Time 2 S-4, S-2	18	0.9832790213	0.9572272104		
Time 2 S-4, S-3	19	0.9886970636	0.9625017026		
Time 2 S-4, S-4	20	0.9693749296	0.9318518473		
Time 3 S-1, S-1, S-1	21	0.9424127368			
Time 3 S-1, S-1, S-2	22	0.9472682532			
Time 3 S-1, S-1, S-3	23	0.9508321387			
Time 3 S-1, S-1, S-4	24	0.9256048053			
Time 3 S-1, S-2, S-1	25	0.9698937713			
Time 3 S-1, S-2, S-2	26	0.9747412918			
Time 3 S-1, S-2, S-3	27	0.9801122898			
Time 3 S-1, S-2, S-4	28	0.9609579282			
Time 3 S-1, S-3, S-1	29	0.9503806226			
Time 3 S-1, S-3, S-2	30	0.9805996987			
Time 3 S-1, S-3, S-3	31	0.9691821954			
Time 3 S-1, S-3, S-4	32	0.9478524225			
Time 3 S-1, S-4, S-1	33	0.9485398381			
Time 3 S-1, S-4, S-2	34	0.9534269227			
Time 3 S-1, S-4, S-3	35	0.9570139788			
Time 3 S-1, S-4, S-4	36	0.9316226298			
Time 3 S-2, S-1, S-1	37	0.9601375282			
Time 3 S-2, S-1, S-2	38	0.9906668429			
Time 3 S-2, S-1, S-3	39	0.9791321239			
Time 3 S-2, S-1, S-4	40	0.9575833729			
Time 3 S-2, S-2, S-1	41	0.9784392385			

(Continued)

EXHIBIT 9.14 *(Continued)*

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 3 S-2, S-2, S-2	42	0.9944994413			
Time 3 S-2, S-2, S-3	43	0.9923910646			
Time 3 S-2, S-2, S-4	44	0.9881386595			
Time 3 S-2, S-3, S-1	45	0.9744393509			
Time 3 S-2, S-3, S-2	46	1.0037479721			
Time 3 S-2, S-3, S-3	47	0.9974134983			
Time 3 S-2, S-3, S-4	48	0.9787812762			
Time 3 S-2, S-4, S-1	49	0.9825458738			
Time 3 S-2, S-4, S-2	50	0.9874566295			
Time 3 S-2, S-4, S-3	51	0.9928976912			
Time 3 S-2, S-4, S-4	52	0.9734934642			
Time 3 S-3, S-1, S-1	53	0.9777040612			
Time 3 S-3, S-1, S-2	54	0.9825906175			
Time 3 S-3, S-1, S-3	55	0.9880048666			
Time 3 S-3, S-1, S-4	56	0.9686962602			
Time 3 S-3, S-2, S-1	57	0.9777040612			
Time 3 S-3, S-2, S-2	58	0.9825906175			
Time 3 S-3, S-2, S-3	59	0.9880048666			
Time 3 S-3, S-2, S-4	60	0.9686962602			
Time 3 S-3, S-3, S-1	61	0.9649211970			
Time 3 S-3, S-3, S-2	62	0.9939435367			
Time 3 S-3, S-3, S-3	63	0.9876709370			
Time 3 S-3, S-3, S-4	64	0.9692207112			
Time 3 S-3, S-4, S-1	65	0.9568733730			
Time 3 S-3, S-4, S-2	66	0.9872988979			
Time 3 S-3, S-4, S-3	67	0.9758033932			
Time 3 S-3, S-4, S-4	68	0.9543279009			
Time 3 S-4, S-1, S-1	69	0.9661716333			
Time 3 S-4, S-1, S-2	70	0.9968928133			
Time 3 S-4, S-1, S-3	71	0.9852856029			
Time 3 S-4, S-1, S-4	72	0.9636014261			
Time 3 S-4, S-2, S-1	73	0.9748191606			
Time 3 S-4, S-2, S-2	74	0.9796912981			
Time 3 S-4, S-2, S-3	75	0.9850895715			
Time 3 S-4, S-2, S-4	76	0.9658379388			
Time 3 S-4, S-3, S-1	77	0.9748191606			
Time 3 S-4, S-3, S-2	78	0.9796912981			
Time 3 S-4, S-3, S-3	79	0.9850895715			
Time 3 S-4, S-3, S-4	80	0.9658379388			
Time 3 S-4, S-4, S-1	81	0.9568821144			
Time 3 S-4, S-4, S-2	82	0.9778172454			
Time 3 S-4, S-4, S-3	83	0.9734710878			
Time 3 S-4, S-4, S-4	84	0.9521725485			

EXHIBIT 9.15 Continuously Compounded Zero Yields

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.2999%	0.8009%	1.2951%	1.7913%
Time 1 S-1	1	2.1422%	2.9602%	3.7150%	
Time 1 S-2	2	0.6887%	0.9308%	1.2372%	
Time 1 S-3	3	1.0913%	1.5648%	2.0479%	
Time 1 S-4	4	1.3519%	1.8358%	2.3270%	
Time 2 S-1, S-1	5	4.5345%	5.5357%		
Time 2 S-1, S-2	6	1.4043%	2.2980%		
Time 2 S-1, S-3	7	2.5755%	3.4543%		
Time 2 S-1, S-4	8	4.8008%	5.3449%		
Time 2 S-2, S-1	9	2.2247%	2.7682%		
Time 2 S-2, S-2	10	0.5966%	0.8622%		
Time 2 S-2, S-3	11	0.8088%	1.1085%		
Time 2 S-2, S-4	12	1.2382%	1.5670%		
Time 2 S-3, S-1	13	1.9035%	2.1466%		
Time 2 S-3, S-2	14	1.4049%	1.8973%		
Time 2 S-3, S-3	15	0.8554%	1.6226%		
Time 2 S-3, S-4	16	2.8291%	3.2407%		
Time 2 S-4, S-1	17	2.1848%	2.4350%		
Time 2 S-4, S-2	18	1.6862%	2.1857%		
Time 2 S-4, S-3	19	1.1367%	1.9110%		
Time 2 S-4, S-4	20	3.1104%	3.5291%		
Time 3 S-1, S-1, S-1	21	5.9312%			
Time 3 S-1, S-1, S-2	22	5.4173%			
Time 3 S-1, S-1, S-3	23	5.0418%			
Time 3 S-1, S-1, S-4	24	7.7308%			
Time 3 S-1, S-2, S-1	25	3.0569%			
Time 3 S-1, S-2, S-2	26	2.5583%			
Time 3 S-1, S-2, S-3	27	2.0088%			
Time 3 S-1, S-2, S-4	28	3.9825%			
Time 3 S-1, S-3, S-1	29	5.0893%			
Time 3 S-1, S-3, S-2	30	1.9591%			
Time 3 S-1, S-3, S-3	31	3.1303%			
Time 3 S-1, S-3, S-4	32	5.3556%			
Time 3 S-1, S-4, S-1	33	5.2831%			
Time 3 S-1, S-4, S-2	34	4.7692%			
Time 3 S-1, S-4, S-3	35	4.3937%			
Time 3 S-1, S-4, S-4	36	7.0827%			
Time 3 S-2, S-1, S-1	37	4.0679%			
Time 3 S-2, S-1, S-2	38	0.9377%			
Time 3 S-2, S-1, S-3	39	2.1089%			
Time 3 S-2, S-1, S-4	40	4.3342%			

(Continued)

EXHIBIT 9.15 (Continued)

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 3 S-2, S-2, S-1	41	2.1797%			
Time 3 S-2, S-2, S-2	42	0.5516%			
Time 3 S-2, S-2, S-3	43	0.7638%			
Time 3 S-2, S-2, S-4	44	1.1932%			
Time 3 S-2, S-3, S-1	45	2.5893%			
Time 3 S-2, S-3, S-2	46	−0.3741%			
Time 3 S-2, S-3, S-3	47	0.2590%			
Time 3 S-2, S-3, S-4	48	2.1447%			
Time 3 S-2, S-4, S-1	49	1.7608%			
Time 3 S-2, S-4, S-2	50	1.2623%			
Time 3 S-2, S-4, S-3	51	0.7128%			
Time 3 S-2, S-4, S-4	52	2.6864%			
Time 3 S-3, S-1, S-1	53	2.2548%			
Time 3 S-3, S-1, S-2	54	1.7563%			
Time 3 S-3, S-1, S-3	55	1.2068%			
Time 3 S-3, S-1, S-4	56	3.1804%			
Time 3 S-3, S-2, S-1	57	2.2548%			
Time 3 S-3, S-2, S-2	58	1.7563%			
Time 3 S-3, S-2, S-3	59	1.2068%			
Time 3 S-3, S-2, S-4	60	3.1804%			
Time 3 S-3, S-3, S-1	61	3.5709%			
Time 3 S-3, S-3, S-2	62	0.6075%			
Time 3 S-3, S-3, S-3	63	1.2406%			
Time 3 S-3, S-3, S-4	64	3.1263%			
Time 3 S-3, S-4, S-1	65	4.4084%			
Time 3 S-3, S-4, S-2	66	1.2782%			
Time 3 S-3, S-4, S-3	67	2.4494%			
Time 3 S-3, S-4, S-4	68	4.6748%			
Time 3 S-4, S-1, S-1	69	3.4414%			
Time 3 S-4, S-1, S-2	70	0.3112%			
Time 3 S-4, S-1, S-3	71	1.4824%			
Time 3 S-4, S-1, S-4	72	3.7078%			
Time 3 S-4, S-2, S-1	73	2.5503%			
Time 3 S-4, S-2, S-2	74	2.0518%			
Time 3 S-4, S-2, S-3	75	1.5023%			
Time 3 S-4, S-2, S-4	76	3.4759%			
Time 3 S-4, S-3, S-1	77	2.5503%			
Time 3 S-4, S-3, S-2	78	2.0518%			
Time 3 S-4, S-3, S-3	79	1.5023%			
Time 3 S-4, S-3, S-4	80	3.4759%			
Time 3 S-4, S-4, S-1	81	4.4075%			
Time 3 S-4, S-4, S-2	82	2.2432%			
Time 3 S-4, S-4, S-3	83	2.6887%			
Time 3 S-4, S-4, S-4	84	4.9009%			




EXHIBIT 9.16 Kamakura Corporation, HJM Zero-Coupon Yield Curve Movements, Three-Factor Empirical Volatilities, 1962–2011




EXHIBIT 9.17 Kamakura Corporation, HJM Zero-Coupon Yield Curve Movements, Three Factors, Empirical Volatility, 1962–2011

VALUATION IN THE HJM FRAMEWORK

Jarrow, as quoted in Chapter 6, described valuation as the expected value of cash flows using the risk-neutral probabilities. Note that column 1 denotes the riskless 1 period interest rate in each scenario. For the state number 84 (three consecutive shift 4s), cash flows at time $T = 4$ would be discounted by the one-year spot rates at time $t = 0$, by the one-year spot rate at time $t = 1$ in state 4 (i.e., shift 4), by the one-year

EXHIBIT 9.18 Spot Rate Process

State	Row Number	0	1	2	3
S-1, S-1, S-1	1	0.3003%	2.1653%	4.6388%	6.1106%
S-1, S-1, S-2	2		0.6911%	1.4142%	5.5667%
S-1, S-1, S-3	3		1.0972%	2.6089%	5.1710%
S-1, S-1, S-4	4		1.3611%	4.9179%	8.0375%
S-1, S-2, S-1	5			2.2496%	3.1041%
S-1, S-2, S-2	6			0.5984%	2.5913%
S-1, S-2, S-3	7			0.8121%	2.0291%
S-1, S-2, S-4	8			1.2459%	4.0628%
S-1, S-3, S-1	9			1.9217%	5.2210%
S-1, S-3, S-2	10			1.4149%	1.9784%
S-1, S-3, S-3	11			0.8591%	3.1798%
S-1, S-3, S-4	12			2.8695%	5.5017%
S-1, S-4, S-1	13			2.2088%	5.4252%
S-1, S-4, S-2	14			1.7005%	4.8848%
S-1, S-4, S-3	15			1.1432%	4.4917%
S-1, S-4, S-4	16			3.1593%	7.3396%
S-2, S-1, S-1	17				4.1517%
S-2, S-1, S-2	18				0.9421%
S-2, S-1, S-3	19				2.1313%
S-2, S-1, S-4	20				4.4295%
S-2, S-2, S-1	21				2.2036%
S-2, S-2, S-2	22				0.5531%
S-2, S-2, S-3	23				0.7667%
S-2, S-2, S-4	24				1.2004%
S-2, S-3, S-1	25				2.6231%
S-2, S-3, S-2	26				−0.3734%
S-2, S-3, S-3	27				0.2593%
S-2, S-3, S-4	28				2.1679%
S-2, S-4, S-1	29				1.7764%
S-2, S-4, S-2	30				1.2703%
S-2, S-4, S-3	31				0.7153%
S-2, S-4, S-4	32				2.7228%
S-3, S-1, S-1	33				2.2804%
S-3, S-1, S-2	34				1.7718%
S-3, S-1, S-3	35				1.2141%
S-3, S-1, S-4	36				3.2315%
S-3, S-2, S-1	37				2.2804%
S-3, S-2, S-2	38				1.7718%
S-3, S-2, S-3	39				1.2141%
S-3, S-2, S-4	40				3.2315%
S-3, S-3, S-1	41				3.6354%
S-3, S-3, S-2	42				0.6093%
S-3, S-3, S-3	43				1.2483%
S-3, S-3, S-4	44				3.1757%

EXHIBIT 9.18 (Continued)

State	Row Number	0	1	2	3
S-3, S-4, S-1	45				4.5070%
S-3, S-4, S-2	46				1.2864%
S-3, S-4, S-3	47				2.4797%
S-3, S-4, S-4	48				4.7858%
S-4, S-1, S-1	49				3.5013%
S-4, S-1, S-2	50				0.3117%
S-4, S-1, S-3	51				1.4934%
S-4, S-1, S-4	52				3.7773%
S-4, S-2, S-1	53				2.5831%
S-4, S-2, S-2	54				2.0730%
S-4, S-2, S-3	55				1.5136%
S-4, S-2, S-4	56				3.5370%
S-4, S-3, S-1	57				2.5831%
S-4, S-3, S-2	58				2.0730%
S-4, S-3, S-3	59				1.5136%
S-4, S-3, S-4	60				3.5370%
S-4, S-4, S-1	61				4.5061%
S-4, S-4, S-2	62				2.2686%
S-4, S-4, S-3	63				2.7252%
S-4, S-4, S-4	64				5.0230%

spot rate in state 20 (shift 4 shift 4) at time $t = 2$, and by the one-year spot rate at time $t = 3$ in state 84 (shift 4 shift 4 shift 4). The discount factor is

$$\begin{aligned} &\text{Discount factor}(0, 4, \text{shift 4-shift 4-shift 4}) \\ &= 1/(1.003003)(1.013519)(1.031104)(1.049009) \end{aligned}$$

These discount factors are displayed here for each potential cash flow date (Exhibit 9.19).

When taking expected values, we can calculate the probability of each scenario coming about since the probabilities of shifts 1, 2, 3, and 4 are one-eighth, one-eighth, one-fourth, and one-half (Exhibit 9.20).

It is again convenient to calculate the probability-weighted discount factors for use in calculating the expected present value of cash flows. The values shown are simply the probability of occurrence of that scenario times the corresponding discount factor (Exhibit 9.21).

We now use the HJM bushy trees we have generated to value representative securities.

VALUATION OF A ZERO-COUPON BOND MATURING AT TIME $T = 4$

A riskless zero-coupon bond pays \$1 in each of the 64 nodes of the bushy tree that prevail at time $T = 4$. Its present value is simply the sum of the probability-weighted

EXHIBIT 9.19 Discount Factors Process

Row Number	Current Time				
	0	1	2	3	4
1	1	0.9970057865	0.9758754267	0.9326130292	0.8789063973
2		0.9970057865	0.9758754267	0.9326130292	0.8834347151
3		0.9970057865	0.9758754267	0.9326130292	0.8867584412
4		0.9970057865	0.9758754267	0.9326130292	0.8632311013
5			0.9901632408	0.9622671525	0.9332969176
6			0.9901632408	0.9622671525	0.9379615274
7			0.9901632408	0.9622671525	0.9431298622
8			0.9901632408	0.9622671525	0.9246982493
9			0.9861850641	0.9510631021	0.9038719431
10			0.9861850641	0.9510631021	0.9326121914
11			0.9861850641	0.9510631021	0.9217534253
12			0.9861850641	0.9510631021	0.9014674653
13			0.9836181004	0.9301320945	0.8822673463
14			0.9836181004	0.9301320945	0.8868129805
15			0.9836181004	0.9301320945	0.8901494165
16			0.9836181004	0.9301320945	0.8665321079
17				0.9683787302	0.9297767604
18				0.9683787302	0.9593406994
19				0.9683787302	0.9481707228
20				0.9683787302	0.9273033706
21				0.9842737987	0.9630521061
22				0.9842737987	0.9788597428
23				0.9842737987	0.9767845229
24				0.9842737987	0.9725989919
25				0.9821871007	0.9570817609
26				0.9821871007	0.9858683106
27				0.9821871007	0.9796466721
28				0.9821871007	0.9613463439
29				0.9779784197	0.9609086610
30				0.9779784197	0.9657112740
31				0.9779784197	0.9710325149
32				0.9779784197	0.9520555997
33				0.9675905555	0.9460172157
34				0.9675905555	0.9507454014
35				0.9675905555	0.9559841777
36				0.9675905555	0.9373013525
37				0.9724265646	0.9507454014
38				0.9724265646	0.9554972185
39				0.9724265646	0.9607621782
40				0.9724265646	0.9419859764
41				0.9777848079	0.9434852873
42				0.9777848079	0.9718628902

EXHIBIT 9.19 (Continued)

Row Number	Current Time				
	0	1	2	3	4
43				0.9777848079	0.9657296374
44				0.9777848079	0.9476892869
45				0.9586759325	0.9173314731
46				0.9586759325	0.9464996916
47				0.9586759325	0.9354792279
48				0.9586759325	0.9148911903
49				0.9623611704	0.9298060639
50				0.9623611704	0.9593709346
51				0.9623611704	0.9482006060
52				0.9623611704	0.9273325962
53				0.9671710430	0.9428168643
54				0.9671710430	0.9475290547
55				0.9671710430	0.9527501083
56				0.9671710430	0.9341304867
57				0.9725003275	0.9480119529
58				0.9725003275	0.9527501083
59				0.9725003275	0.9579999309
60				0.9725003275	0.9392777118
61				0.9534947268	0.9123820502
62				0.9534947268	0.9323435872
63				0.9534947268	0.9281995489
64				0.9534947268	0.9078915040

EXHIBIT 9.20 Probability of Each State

Row Number	Current Time				
	0	1	2	3	4
1	100.0000%	12.5000%	1.562500%	0.1953%	0.1953%
2		12.5000%	1.562500%	0.1953%	0.1953%
3		25.0000%	3.125000%	0.3906%	0.3906%
4		50.0000%	6.250000%	0.7813%	0.7813%
5			1.562500%	0.1953%	0.1953%
6			1.562500%	0.1953%	0.1953%
7			3.125000%	0.3906%	0.3906%
8			6.250000%	0.7813%	0.7813%
9			3.125000%	0.3906%	0.3906%
10			3.125000%	0.3906%	0.3906%
11			6.250000%	0.7813%	0.7813%

(Continued)

EXHIBIT 9.20 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
12			12.500000%	1.5625%	1.5625%
13			6.250000%	0.7813%	0.7813%
14			6.250000%	0.7813%	0.7813%
15			12.500000%	1.5625%	1.5625%
16			25.000000%	3.1250%	3.1250%
17				0.1953%	0.1953%
18				0.1953%	0.1953%
19				0.3906%	0.3906%
20				0.7813%	0.7813%
21				0.1953%	0.1953%
22				0.1953%	0.1953%
23				0.3906%	0.3906%
24				0.7813%	0.7813%
25				0.3906%	0.3906%
26				0.3906%	0.3906%
27				0.7813%	0.7813%
28				1.5625%	1.5625%
29				0.7813%	0.7813%
30				0.7813%	0.7813%
31				1.5625%	1.5625%
32				3.1250%	3.1250%
33				0.3906%	0.3906%
34				0.3906%	0.3906%
35				0.7813%	0.7813%
36				1.5625%	1.5625%
37				0.3906%	0.3906%
38				0.3906%	0.3906%
39				0.7813%	0.7813%
40				1.5625%	1.5625%
41				0.7813%	0.7813%
42				0.7813%	0.7813%
43				1.5625%	1.5625%
44				3.1250%	3.1250%
45				1.5625%	1.5625%
46				1.5625%	1.5625%
47				3.1250%	3.1250%
48				6.2500%	6.2500%
49				0.7813%	0.7813%
50				0.7813%	0.7813%
51				1.5625%	1.5625%
52				3.1250%	3.1250%
53				0.7813%	0.7813%
54				0.7813%	0.7813%

EXHIBIT 9.20 (Continued)

Row Number	Current Time				
	0	1	2	3	4
55				1.5625%	1.5625%
56				3.1250%	3.1250%
57				1.5625%	1.5625%
58				1.5625%	1.5625%
59				3.1250%	3.1250%
60				6.2500%	6.2500%
61				3.1250%	3.1250%
62				3.1250%	3.1250%
63				6.2500%	6.2500%
64				12.5000%	12.5000%
Total	100.0000%	100.0000%	100.0000%	100.0000%	100.0000%

EXHIBIT 9.21 Probability-Weighted Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1	0.1246257233	0.0152480535	0.0018215098	0.0017166141
2		0.1246257233	0.0152480535	0.0018215098	0.0017254584
3		0.2492514466	0.0304961071	0.0036430196	0.0034639002
4		0.4985028932	0.0609922142	0.0072860393	0.0067439930
5			0.0154713006	0.0018794280	0.0018228455
6			0.0154713006	0.0018794280	0.0018319561
7			0.0309426013	0.0037588561	0.0036841010
8			0.0618852026	0.0075177121	0.0072242051
9			0.0308182833	0.0037150902	0.0035307498
10			0.0308182833	0.0037150902	0.0036430164
11			0.0616365665	0.0074301805	0.0072011986
12			0.1232731330	0.0148603610	0.0140854291
13			0.0614761313	0.0072666570	0.0068927136
14			0.0614761313	0.0072666570	0.0069282264
15			0.1229522625	0.0145333140	0.0139085846
16			0.2459045251	0.0290666280	0.0270791284
17				0.0018913647	0.0018159702
18				0.0018913647	0.0018737123
19				0.0037827294	0.0037037919
20				0.0075654588	0.0072445576
21				0.0019224098	0.0018809611
22				0.0019224098	0.0019118354
23				0.0038448195	0.0038155645

(Continued)

EXHIBIT 9.21 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
24				0.0076896391	0.0075984296
25				0.0038366684	0.0037386006
26				0.0038366684	0.0038510481
27				0.0076733367	0.0076534896
28				0.0153466734	0.0150210366
29				0.0076404564	0.0075070989
30				0.0076404564	0.0075446193
31				0.0152809128	0.0151723830
32				0.0305618256	0.0297517375
33				0.0037796506	0.0036953797
34				0.0037796506	0.0037138492
35				0.0075593012	0.0074686264
36				0.0151186024	0.0146453336
37				0.0037985413	0.0037138492
38				0.0037985413	0.0037324110
39				0.0075970825	0.0075059545
40				0.0151941651	0.0147185309
41				0.0076389438	0.0073709788
42				0.0076389438	0.0075926788
43				0.0152778876	0.0150895256
44				0.0305557752	0.0296152902
45				0.0149793114	0.0143333043
46				0.0149793114	0.0147890577
47				0.0299586229	0.0292337259
48				0.0599172458	0.0571806994
49				0.0075184466	0.0072641099
50				0.0075184466	0.0074950854
51				0.0150368933	0.0148156345
52				0.0300737866	0.0289791436
53				0.0075560238	0.0073657568
54				0.0075560238	0.0074025707
55				0.0151120475	0.0148867204
56				0.0302240951	0.0291915777
57				0.0151953176	0.0148126868
58				0.0151953176	0.0148867204
59				0.0303906352	0.0299374978
60				0.0607812705	0.0587048570
61				0.0297967102	0.0285119391
62				0.0297967102	0.0291357371
63				0.0595934204	0.0580124718
64				0.1191868409	0.1134864380

discount factors in the last column on the right-hand side. Their sum is 0.9308550992, which is the value we should get in a no-arbitrage economy, the value observable in the market and used as an input to create the tree.

VALUATION OF A COUPON-BEARING BOND PAYING ANNUAL INTEREST

Next, we value a bond with no credit risk that pays \$3 in interest at every scenario at times $T = 1, 2, 3$, and 4 plus principal of 100 at time $T = 4$. The valuation is calculated by multiplying each cash flow by the matching probability-weighted discount factor, to get a value of 104.7070997370. It will surprise a few (primarily those readers who skipped Chapters 6 through 8) that this is the same value that we arrived at in Chapters 6 through 8, even though the volatilities used and number of risk factors used are different. The values are the same because, by construction, our valuations for the zero-coupon bond prices at time zero for maturities at $T = 1, 2, 3$, and 4 continue to match the inputs. Multiplying these zero-coupon bond prices times 3, 3, 3, and 103 also leads to a value of 104.7070997370 as it should (Exhibit 9.22).

EXHIBIT 9.22 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.00	3.00	3.00	3.00	103.00
2		3.00	3.00	3.00	103.00
3		3.00	3.00	3.00	103.00
4		3.00	3.00	3.00	103.00
5			3.00	3.00	103.00
6			3.00	3.00	103.00
7			3.00	3.00	103.00
8			3.00	3.00	103.00
9			3.00	3.00	103.00
10			3.00	3.00	103.00
11			3.00	3.00	103.00
12			3.00	3.00	103.00
13			3.00	3.00	103.00
14			3.00	3.00	103.00
15			3.00	3.00	103.00
16			3.00	3.00	103.00
17				3.00	103.00
18				3.00	103.00
19				3.00	103.00
20				3.00	103.00
21				3.00	103.00

(Continued)

EXHIBIT 9.22 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
22				3.00	103.00
23				3.00	103.00
24				3.00	103.00
25				3.00	103.00
26				3.00	103.00
27				3.00	103.00
28				3.00	103.00
29				3.00	103.00
30				3.00	103.00
31				3.00	103.00
32				3.00	103.00
33				3.00	103.00
34				3.00	103.00
35				3.00	103.00
36				3.00	103.00
37				3.00	103.00
38				3.00	103.00
39				3.00	103.00
40				3.00	103.00
41				3.00	103.00
42				3.00	103.00
43				3.00	103.00
44				3.00	103.00
45				3.00	103.00
46				3.00	103.00
47				3.00	103.00
48				3.00	103.00
49				3.00	103.00
50				3.00	103.00
51				3.00	103.00
52				3.00	103.00
53				3.00	103.00
54				3.00	103.00
55				3.00	103.00
56				3.00	103.00
57				3.00	103.00
58				3.00	103.00
59				3.00	103.00
60				3.00	103.00
61				3.00	103.00
62				3.00	103.00
63				3.00	103.00
64				3.00	103.00
Risk-Neutral Value =					104.7070997370

VALUATION OF A DIGITAL OPTION ON THE ONE-YEAR U.S. TREASURY RATE

Now we value a digital option that pays \$1 at time $T = 3$ if (at that time) the one-year U.S. Treasury rate (for maturity at $T = 4$) is over 8 percent. If we look at the table of the term structure of one-year spot rates over time, this happens in only one scenario, in row 4. Note also that we have a slightly negative spot rate in row 26. This can be easily overridden by changing the relevant interest rate volatility assumptions (Exhibit 9.23):

EXHIBIT 9.23 Spot Rate Process

State	Row Number	0	1	2	3
S-1, S-1, S-1	1	0.3003%	2.1653%	4.6388%	6.1106%
S-1, S-1, S-2	2		0.6911%	1.4142%	5.5667%
S-1, S-1, S-3	3		1.0972%	2.6089%	5.1710%
S-1, S-1, S-4	4		1.3611%	4.9179%	8.0375%
S-1, S-2, S-1	5			2.2496%	3.1041%
S-1, S-2, S-2	6			0.5984%	2.5913%
S-1, S-2, S-3	7			0.8121%	2.0291%
S-1, S-2, S-4	8			1.2459%	4.0628%
S-1, S-3, S-1	9			1.9217%	5.2210%
S-1, S-3, S-2	10			1.4149%	1.9784%
S-1, S-3, S-3	11			0.8591%	3.1798%
S-1, S-3, S-4	12			2.8695%	5.5017%
S-1, S-4, S-1	13			2.2088%	5.4252%
S-1, S-4, S-2	14			1.7005%	4.8848%
S-1, S-4, S-3	15			1.1432%	4.4917%
S-1, S-4, S-4	16			3.1593%	7.3396%
S-2, S-1, S-1	17				4.1517%
S-2, S-1, S-2	18				0.9421%
S-2, S-1, S-3	19				2.1313%
S-2, S-1, S-4	20				4.4295%
S-2, S-2, S-1	21				2.2036%
S-2, S-2, S-2	22				0.5531%
S-2, S-2, S-3	23				0.7667%
S-2, S-2, S-4	24				1.2004%
S-2, S-3, S-1	25				2.6231%
S-2, S-3, S-2	26				-0.3734%
S-2, S-3, S-3	27				0.2593%
S-2, S-3, S-4	28				2.1679%
S-2, S-4, S-1	29				1.7764%
S-2, S-4, S-2	30				1.2703%
S-2, S-4, S-3	31				0.7153%

(Continued)

EXHIBIT 9.23 (Continued)

State	Row Number	0	1	2	3
S-2, S-4, S-4	32				2.7228%
S-3, S-1, S-1	33				2.2804%
S-3, S-1, S-2	34				1.7718%
S-3, S-1, S-3	35				1.2141%
S-3, S-1, S-4	36				3.2315%
S-3, S-2, S-1	37				2.2804%
S-3, S-2, S-2	38				1.7718%
S-3, S-2, S-3	39				1.2141%
S-3, S-2, S-4	40				3.2315%
S-3, S-3, S-1	41				3.6354%
S-3, S-3, S-2	42				0.6093%
S-3, S-3, S-3	43				1.2483%
S-3, S-3, S-4	44				3.1757%
S-3, S-4, S-1	45				4.5070%
S-3, S-4, S-2	46				1.2864%
S-3, S-4, S-3	47				2.4797%
S-3, S-4, S-4	48				4.7858%
S-4, S-1, S-1	49				3.5013%
S-4, S-1, S-2	50				0.3117%
S-4, S-1, S-3	51				1.4934%
S-4, S-1, S-4	52				3.7773%
S-4, S-2, S-1	53				2.5831%
S-4, S-2, S-2	54				2.0730%
S-4, S-2, S-3	55				1.5136%
S-4, S-2, S-4	56				3.5370%
S-4, S-3, S-1	57				2.5831%
S-4, S-3, S-2	58				2.0730%
S-4, S-3, S-3	59				1.5136%
S-4, S-3, S-4	60				3.5370%
S-4, S-4, S-1	61				4.5061%
S-4, S-4, S-2	62				2.2686%
S-4, S-4, S-3	63				2.7252%
S-4, S-4, S-4	64				5.0230%

The evolution of the spot rate can be displayed graphically in Exhibit 9.24.

The cash flow payoff in the one relevant scenario can be input in the table above and multiplied by the appropriate probability-weighted discount factor to find that this option has a value of only \$0.0072860393. That is because the probability of row 4 occurring is only 0.7813% and the payoff is discounted by the simulated spot rates prior to time $T = 3$.




EXHIBIT 9.24 Evolution of One-Year U.S. Treasury Spot Rate with Three-Factor HJM Model with Rate and Maturity-Dependent Volatility

CONCLUSION

This chapter, the fourth example of HJM interest rate simulation, shows how to simulate zero-coupon bond prices, forward rates, and zero-coupon bond yields in an HJM framework with three risk factors and rate-dependent and maturity-dependent interest rate volatility. The results show a rich twist in simulated yield curves and a pull of rates upward from a very low rate environment. Monte Carlo simulation, an alternative to the bushy tree framework, can be done in a fully consistent way. The technique can be generalized to any user-specified number of risk factors, a long-standing feature of the Kamakura Risk Manager enterprise risk management system. We explain how to do this in Chapter 10. Also in Chapter 10, we discuss the implications of this type of analysis for the Basel-Specified Internal Capital Adequacy Assessment Process (ICAAP) for interest rate risk.

NOTE

1. This chapter incorporates some important revisions in previously published work by Robert A. Jarrow (2002, Ch. 15), to whom we give special thanks for his contributions to this revision.

Valuation, Liquidity, and Net Income

Chapters 6 through 9 introduced modern multifactor interest rate models of Heath, Jarrow, and Morton (HJM). In this chapter, we introduce some of the key issues for the practical application of the HJM concepts. We spend the rest of the book addressing these issues in great detail, but we start here with an overview of the most important topics.

HOW MANY RISK FACTORS ARE NECESSARY TO ACCURATELY MODEL MOVEMENTS IN THE RISK-FREE YIELD CURVE?

In Chapters 6 and 7, we implemented a one-factor model of yield curve movements in the HJM framework. We extended that to a two-factor model in Chapter 8 and a three-factor model in Chapter 9. This begs the question, How many risk factors are necessary to accurately model the movements of the forward rates (zero yields, and zero prices) of the risk-free curve? There are two answers. The most important is the perspective of a sophisticated analyst whose primary concern is accurately reporting the risk that his or her institution faces. The second perspective is that of a sophisticated regulator. We address these both in turn.

We start with a perspective of a sophisticated risk manager whose primary concern is accuracy in modeling. There are two conditions for determining the risk number of risk factors driving the yield curve:

1. *Necessary condition:* All securities, whose value is tied to the risk-free yield curve and its volatility, should be priced accurately if a market price is visible
2. *Sufficient condition:* Providing condition one is met, the residual correlation of changes in forward rates 1 through n should be statistically indistinguishable from zero.

Note that the legacy Nelson-Siegel yield curve-fitting technique fails even this necessary condition because it lacks sufficient parameterization to price all input securities accurately, violating the no-arbitrage condition. We can illustrate these conditions for more sophisticated yield curve-fitting techniques using data from Chapters 8 and 9. In Chapter 9, we faced this question: Do we need to add a third factor to the two risk factors which we used in Chapter 8? In Chapter 8, we denoted the change in the one-year U.S. Treasury spot rate as risk factor 1 and the residual risk factor was denoted “everything else.” We inserted the econometric relationships

that sought to explain the movements of the one-period forward rates maturing at times 2 through 10. We took the residuals from those relationships and analyzed their correlation with each other (Exhibit 10.1).

The lowest correlation of the remaining variation in forward rate changes after accounting for risk factor 1 was 67.75 percent. Even though a one-factor model could meet the necessary condition above for no arbitrage (as we saw in Chapter 6), it fails the second condition that the correlation of the remaining variation in forward rate changes be indistinguishable from zero. What if we add a second variable called “everything else”? That variable would meet the sufficient condition only if the correlations remaining after risk factor 1 were all either 100 percent or -100 percent, indicating all remaining random shifts were perfectly correlated with the “all other” risk factor, either positively or negatively. That is not the case here.

For that reason, we decided to add a third factor in Chapter 9 and we formally denoted the second risk factor as the one-period change in the 10-year forward rate. After we take the residuals remaining after risk factor 1 and try to explain them with our new risk factor 2, we have another correlation matrix remaining (Exhibit 10.2).

The use of the second risk factor has reduced the lowest correlation of residual risk from 67.75 percent to 51.14 percent. That represents serious progress, but adding a third factor “everything else” does not eliminate the risk because the remaining correlations after the two designated risk factors are extracted are not all either 100 percent or -100 percent. From this, a sophisticated analyst knows that we must add more factors. How do we know how many are necessary? One method is to keep incrementally adding risk factors until the residual correlations are either all 100 percent or -100 percent (in which case, the last factor will be “everything else”), or we add factors until the remaining residual risk has a correlation that is statistically indistinguishable from zero. Dickler, Jarrow, and van Deventer estimate that 6 to 10 factors are needed to realistically model forward rate movements of the U.S. Treasury curve from 1962 to 2011. Numerous academic authors have investigated the same issue and conclude that at least three factors, and probably more, are needed.¹

What is the regulatory perspective on the problem? Perhaps the most insightful comment comes from the Basel Committee on Banking Supervision (2010). In its *Revisions to the Basel II Market Risk Framework*, the Committee states its requirements clearly:

For material exposures to interest rate movements in the major currencies and markets, banks must model the yield curve using a minimum of six risk factors. (12)

REVISITING THE PHRASE “NO ARBITRAGE”

Condition 1 sets as a minimum standard for the number (and magnitude) of the risk factors the existence of “no arbitrage.” This phrase often has a different meaning in an academic context than it does in practical application, but the true meaning should be the same as stated in Condition 1: All securities whose pricing is fully determined by the risk-free yield curve and its assumed volatility should be priced accurately if market prices are visible. If that is not the case, there is one of three types of errors: (1) the number of risk factors is too few; (2) the magnitude and term

EXHIBIT 10.1 Correlation of Residuals from Regression 1 on U.S. Treasury Spot and Forward Rates, 1962 to 2011

Maturity in Year	1-Year U.S. Treasury Forward Rate Maturing in Year							
	2	3	4	5	6	7	8	9
2	1							
3	0.92016	1						
4	0.84091	0.93618	1					
5	0.80762	0.82393	0.94274	1				
6	0.79099	0.80811	0.87585	0.95039	1			
7	0.7531	0.80368	0.82023	0.85957	0.96701	1		
8	0.73779	0.80699	0.83312	0.83894	0.91046	0.9611	1	
9	0.6775	0.7417	0.80867	0.79686	0.78653	0.81691	0.93905	1

Sources: Kamakura Corporation; Federal Reserve.

EXHIBIT 10.2 Correlation of Residuals from Regression 2 on U.S. Treasury Spot and Forward Rates, 1962 to 2011

Maturity in Year	1-Year U.S. Treasury Forward Rate Maturing in Year							
	2	3	4	5	6	7	8	9
2	1.000000							
3	0.887509	1.000000						
4	0.771819	0.906347	1.000000					
5	0.720069	0.726640	0.901236	1.000000				
6	0.700086	0.711934	0.803607	0.924316	1.000000			
7	0.644186	0.704231	0.710768	0.776320	0.949579	1.000000		
8	0.607691	0.695730	0.687453	0.705748	0.871682	0.960999	1.000000	
9	0.511367	0.592157	0.632456	0.619087	0.674017	0.738484	0.882839	1.000000

Sources: Kamakura Corporation; Federal Reserve.

structure of at least one risk factor is not correct; or (3) there is an error in the software implementation or market data inputs. One common market data input “error” might be simply the bid-offered spread on a given bond, the most benign type of value discrepancy that will not go away. A more refined statement of “priced accurately” is “priced such that the valuation is at or between the bid and offered prices” for the security.

In the HJM context, it was assumed that the market for zero-coupon bonds is complete, perfectly competitive, and that one can go long or short zero-coupon bonds in any amount costlessly. Because of this assumption, it is common but wrong for an analyst to conclude that his HJM implementation for the U.S. Treasury curve is complete if we accurately value the zero-coupon bonds used as input as we saw in Chapters 6 through 9. This conclusion would be wrong if there were other securities tied to the risk-free curve and its volatility that are mispriced after implementation. Consider the 3 percent coupon bond that we valued in Chapters 6 through 9. What if this bond had been callable, and the theoretical pricing were not correct, even though another noncallable bond with the same terms is priced correctly? This means that the number of risk factors and their magnitude remains incorrect. The solution is to add *all* securities with observable prices to assure that this source of error does not come about.

We discuss this in most detail in Chapter 14, estimating the parameters of interest rate models. It is most complex in modeling a yield curve tied to LIBOR and the interest rate swap curve because of many complications: (1) the large number of instruments tied to this curve; (2) the fact that this curve is not risk free; (3) the fact that the credit risk of Eurodollar deposits (for which LIBOR is quoted) and the credit risk of the swap curve are not the same; and (4) the mounting evidence and legal proceedings on the potential manipulation of LIBOR. These issues get special attention in Chapter 17.

VALUATION, LIQUIDITY RISK, AND NET INCOME

A complete analysis of risk of a large portfolio or a large financial institution usually includes a number of calculations, just a few of which are listed here:

1. Valuation of all assets and liabilities, both at time 0 and at user-specified dates in the future
2. Gross and net cash received and cash paid, for liquidity risk assessment, for any set of user-defined time intervals
3. Financial accounting—that is for generally accepted accounting principles (GAAP) or international financial reporting standards (IFRS) basis—net income, interest income, interest expense and (often) noninterest income and expense for any set of user-defined time periods

What extensions of the HJM examples in Chapters 6 through 9 are needed to achieve these objectives? Three major extensions are necessary:

1. A distinction between risk-neutral probabilities and empirical probabilities that a given interest rate scenario will come about

2. Monte Carlo simulation, not just bushy tree simulation, particularly to ensure that a large number of interest rate scenarios can be generated even in near-term time intervals
3. An implementation of GAAP and IFRS accounting accruals
4. Expansion of analysis to include “insurance events,” such as default, prepayment, mortality, property and casualty risks, and a large array of macroeconomic risk factors beyond the risk-free yield curve

The third extension is very important, hard to do well, and best covered in other volumes and in the user documentation for a state-of-the-art enterprise risk management system. We do not address it in this book. The fourth extension is at the heart of the rest of this book. We turn now, however, to the important distinction between risk-neutral and empirical default probabilities.

RISK-NEUTRAL AND EMPIRICAL PROBABILITIES OF INTEREST RATE MOVEMENTS

In Chapters 6 through 9, we followed Jarrow (2002) and showed that risk-neutral valuation allows us to accurately value zero-coupon bonds (and other securities, as we see in the following chapters) using pseudo-probabilities that, without loss of generality, can be selected for the user’s convenience. In Chapters 6 and 7, the probabilities of an up- and downshift were set at one-half. In Chapter 8, with two risk factors, the probabilities of an upshift, midshift, and downshift were set at 1/4, 1/4 and 1/2. In Chapter 9, with three factors, the probabilities for shifts 1, 2, 3, and 4 were set at 1/8, 1/8, 1/4, and 1/2. All of these choices (and many others) result in accurate valuation provided that the formulas for the shifted forward rates and zero-coupon bond prices are adjusted correctly. We need to ensure that the expected value and variance of the changes in forward returns fit these two constraints and that the drift in forward rates μ (adjusted for the choice of pseudo-probabilities and the number of risk factors) is consistent with the formula in Chapter 9:

$$E \left[\frac{\log F(t + \Delta, T)}{\Delta} - \frac{\log F(t, T)}{\Delta} \right] = \mu(t, T) \Delta$$

$$\text{Var}_t \left[\frac{\log F(t + \Delta, T)}{\Delta} - \frac{\log F(t, T)}{\Delta} \right] = \left[\sum_{i=1}^n \sigma_i(t, T)^2 \right] \Delta$$

The reader can confirm that these relationships hold for the three-factor example in Chapter 9, for example, by taking the expected value of the four shifted forward rate values, weighted by the probabilities 1/8, 1/8, 1/4, and 1/2. The variance can be confirmed in the same way. These are the key relationships that one uses to parameterize a Monte Carlo simulation using the HJM approach.

What if an analyst is interested in using the actual probabilities of forward rate shifts (also called the *empirical* or *statistical probabilities*) instead of the

pseudo-probabilities (also called *risk-neutral probabilities*). Jarrow shows that these empirical probabilities are straightforward shifts of the pseudo-probabilities. The amount of the shift is a function of the risk aversion embedded in the term structure of risk-free interest rates.²

When would the empirical probabilities be used instead of (or in addition to) the pseudo-probabilities? Whenever the analyst is doing a simulation of actual cash flows and probabilities, actual net income and actual probabilities, or the actual self-assessment of the institution's probability of default. For a one-factor interest rate model, the previous expected value and variance would be shifted in a simple way using the time-dependent function $\theta(t)$:

$$E \left[\frac{\log F(t + \Delta, T)}{\Delta} - \frac{\log F(t, T)}{\Delta} \right] = [\mu(t, T) - \theta(t)\sigma(t, T)]\Delta$$

$$\text{Var}_t \left[\frac{\log F(t + \Delta, T)}{\Delta} - \frac{\log F(t, T)}{\Delta} \right] = \sigma(t, T)^2\Delta - \theta(t)^2\sigma(t, T)^2\Delta^2$$

Instead of the pseudo-probability of an upshift, which was set at 1/2 in Chapters 6 and 7, the empirical probability of an upshift would be set at³

$$q_t(s_t = up) = \frac{1}{2} + \frac{1}{2}\theta(t, s_t)\sqrt{\Delta}$$

Within a net income simulation, the empirical probabilities would be used to generate an empirical distribution of net income in each user-defined period. Part of the net income calculation—under GAAP accounting, for example—might involve the amortization of a premium or discount of a security purchased at a market value different from par value. This amortization is required under GAAP accounting to be updated as the security's value changes. Valuation of this security would be done using the pseudo probabilities that are associated with the empirical probabilities. This is a feature one can find in a best practice enterprise-wide risk management system.

How is the function $\theta(t)$ determined? It is done using a historical study, such as Dickler, Jarrow, and van Deventer (2011), in which forward rates maturing at time T^* , observed at various points in time t , are compared statistically with the actual spot interest rates, which came about for maturity at time T^* . The graph in Exhibit 10.3, for example, shows the convergence of all one-year U.S. Treasury forward rates (out to nine years forward) to the actual one-year U.S. Treasury spot rate for 24 dates on which the one-year spot Treasury rate was 0.27 percent.

MONTE CARLO SIMULATION USING HJM MODELING

How does one use Monte Carlo simulation in the HJM framework? For reasons outlined next, in most cases, Monte Carlo simulation is preferred to the bushy tree approach we illustrated in Chapters 6 through 9. The process follows the pseudo




EXHIBIT 10.3 One-Year U.S. Treasury Spot and Forward Rates Out to Nine Years Forward for All Spot Rates Equal to 0.27 Percent, 1962–2011

code here for M scenarios and N periods into the future, using the example of the three-factor model in Chapter 9:

```

001 Load-relevant zero-coupon bond prices
002 Load assumptions for risk factors 1, 2, and 3
003 Set ScenarioCounter = 1
004 Set RemainingPeriodCounter = 1
010 Get relevant zero-coupon bond prices, forward rates, and sigmas for new
    current time  $t$ . Remember that in Chapters 7, 8, and 9, the sigmas chosen
    were a function of the spot rate at the current time  $t$ 
011 If RemainingPeriodCounter  $\leq N$ , take one random draw each for risk factors
    1, 2, and 3 else go to 100
012 Shift all forward rates and zero-coupon bond prices based on the new shift
    amounts, which will be a randomly drawn number of sigmas up or down
    for all three risk factors
013 Store results for valuation and other calculations
014 RemainingPeriodCounter = RemainingPeriodCounter + 1
015 Go to 010
100 If ScenarioCounter  $< M$ , ScenarioCounter = ScenarioCounter + 1 and goto
    004, else 500
500 Do valuations
501 Do net income simulations with default "turned on"
502 Do liquidity risk analysis with default "turned on"
503 Do value at risk with default "turned off"
504 Do credit adjusted value at risk with default "turned on"
998 Generate reports
999 End

```

This approach is a very straightforward modification of what we did in Chapters 6 through 9. That being said, many analysts “get it wrong,” even with a good understanding of Chapters 6 through 9 and this Monte Carlo process. In the next section, we discuss some common pitfalls.

COMMON PITFALLS IN INTEREST RATE RISK MANAGEMENT

The most common pitfalls in interest rate risk analysis are very fundamental mistakes, but they are surprisingly common:

1. Using a number of risk factors or volatility assumptions that are unrealistic, violating conditions 1, 2, or both. The most common mistake is the most serious one: using one- or two-factor interest rate models in a world where both realism and regulators demand 6–10 factors.
2. Selecting interest rate paths from a bushy tree or interest rate lattice as a substitute for Monte Carlo, without realizing the huge errors that result from this detour from best practice.

We discuss each of these problems in turn.

Pitfalls in the Use of One-Factor Term Structure Models

One-factor models of the term structure of interest rates were developed to provide insights into the valuation of fixed income options. Our implementation of the one-factor HJM model in Chapters 6 and 7 is a good example of a modern implementation. Since the one-factor model’s early development (from 1977 to 1993), they have been consistently misapplied in asset and liability (interest rate risk) management. Many analysts, choosing one of the common one-factor models, create random interest rate scenarios to evaluate the interest rate risk of the firm. This section uses 50 years of Dickler, Jarrow, and van Deventer’s U.S. Treasury yield curve information to prove that the major implications of one-factor term structure models are inconsistent with historical data. A multifactor approach is necessary to create realistic interest rate simulations, as we have emphasized in Chapters 8 and 9.

In this section, we take the same approach as Jarrow, van Deventer, and Wang (2003) in their paper “A Robust Test of Merton’s Structural Model for Credit Risk.” In that paper, the authors explained how, for any choice of model parameters, the Merton model of risky debt implies that stock prices would rise and credit spreads would fall if the assets of the firm were to rise in value. The opposite movements would occur if the assets of the firm were to fall in value. Jarrow, van Deventer, and Wang used a wide array of bond data and proved that this implication of the Merton model was strongly rejected by the data at extremely high levels of statistical significance.

We now take a similar approach to one-factor term structure models. Speaking broadly, the original Macaulay (1938) duration concept can be thought of as the first one-factor term structure model. Merton (1970) developed a similar term structure model in a continuous time framework. Vasicek (1977) brought a greater

realism to one-factor term structure modeling by introducing the concept of mean reversion to create interest rate cycles. Cox, Ingersoll, and Ross (1985) pioneered the affine term structure model, where the magnitude of random shocks to the short-term interest rate varies with the level of interest rates. Ho and Lee (1986) and Hull and White (1993) developed one-factor term structure models that could be fit exactly to observable yields. Black, Derman, and Toy (1990) and Black and Karasinski (1991) focused on term structure models based on the log of the short rate of interest, insuring that rates would not be negative. We cover the basics of these legacy models briefly in Chapter 13. All are special cases of the general HJM (1990a, 1990b, 1992a) framework for N -factor term structure models that we have illustrated in Chapters 6 through 9.⁴

Among one-factor term structure models, for any set of parameters, there is almost always one characteristic in common. The majority of one-factor models make all interest rates positively correlated (which implies the time-dependent interest rate volatility $\sigma(T)$ is greater than 0). We impose this condition on a one-factor model. The nonrandom, time-dependent drift(T) in rates is small and, if risk premia are consistent and positive (for bonds) across all maturities T , then the drift will be positive as in the HJM (1992) drift condition that we labeled μ in Chapters 6 through 9. With more than 200 business days per year, the daily magnitude of the interest rate drift is almost always between plus one and minus one basis point. Hence, to remove the effect of the drift, we might want to exclude from consideration yield curve changes that are small. However, we ignore the impact of drift in what follows.

Given these restrictions, for a random shift upward in the single-risk factor (usually the short rate of interest in most of the legacy one-factor models), the zero-coupon yields at all maturities will rise. For a random shift downward in the single-factor, all zero-coupon yields will fall. If the short rate is unchanged, there will be no random shock to any of the other zero-coupon yields either.

Following Jarrow, van Deventer, and Wang, we ask this simple question: What percent of the time is it true that all interest rates either rise together, fall together, or stay unchanged? If this implication of one-factor term structure models is consistent with actual yield curve movements, we cannot reject their use. If this implication of one-factor term structure models is inconsistent with actual yield movements, we reject their use in interest rate risk management and asset and liability management.

Results from the U.S. Treasury Market, 1962 to 2011 In this section, we use the same data as Dickler, Jarrow, and van Deventer (2011): U.S. Treasury yields reported by the Federal Reserve in its H15 statistical release from January 2, 1962, to August 22, 2011 (Exhibit 10.4). During this period, the Federal Reserve changed the maturities on which it reported data frequently.

We calculate the daily changes in yields over this period, eliminating the first daily change after a change in data regime, leaving 12,386 data points for analysis. We calculate daily yield changes on three bases:

- Using the yields from the Federal Reserve with no other analysis
- Using monthly forward rates
- Using monthly zero-coupon bond yields, expressed on a continuous compounding basis

EXHIBIT 10.4 Analysis of Federal Reserve H15 Release Data Regimes

Start Date	End Date	Data Regime	Number of Observations
1/2/1962	6/30/1969	1, 3, 5, and 10 years	1,870
7/1/1969	5/28/1976	1, 3, 5, 7, and 10 years	1,722
6/1/1976	2/14/1977	1, 2, 3, 5, 7, and 10 years	178
2/15/1977	12/31/1981	1, 2, 3, 5, 7, 10, and 30 years	1,213
1/4/1982	9/30/1993	3 and 6 months with 1, 2, 3, 5, 7, 10, and 30 years	2,935
10/1/1993	7/30/2001	3 and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	1,960
7/31/2001	2/15/2002	1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	135
2/19/2002	2/8/2006	1, 3, and 6 months with 1, 2, 3, 5, 7, 10, and 20 years	994
2/9/2006	8/22/2011	Second era: 1, 3, and 6 months with 1, 2, 3, 5, 7, 10, 20, and 30 years	1,388
		Total	12,395

Monthly forward rates and zero-coupon yields were calculated using the maximum smoothness forward rate approach that we outlined in Chapter 5. We then asked the question posed previously: What percentage of the time were yield curve movements consistent with the implications of one-factor term structure models? Yield curve movements were considered consistent with one-factor term structure models in three cases:

- Negative shift, in which interest rates at all maturities declined
- Positive shift, in which interest rates at all maturities increased
- Zero shift, in which interest rates at all maturities remained unchanged

By definition, any other yield curve movement consists of a combination of positive, negative, and zero shifts at various maturities on that day. We label this a yield curve “twist,” something that cannot be explained by one-factor term structure models.

The results are striking. Using only the input data from the H15 statistical release, with no yield curve smoothing, 7,716 of the 12,386 days of data showed yield curve twists, 62.3 percent of all observations (Exhibit 10.5). This is a stunningly high level of inconsistency with the assumptions of one-factor term structure models.

For monthly zero-coupon bond yields, yield curve twists occurred on 75.2 percent of the 12,386 days in the sample. When forward rates were examined, yield curve twists prevailed on 94.3 percent of the days in the sample (Exhibit 10.6). One-factor term structure models would have produced results that are consistent with actual forward rates only 5.7 percent of the 1962–2011 period.

Using the approach of Jarrow, van Deventer, and Wang (2003), we can now pose this question: Is the consistency ratio we have derived different to a statistically significant degree from the 100 percent ratio that would prevail if the hypothesis of a one-factor term structure is true?

We use the fact that the measured consistency ratio p is a binomial variable (true or not true, consistent or not consistent) with a standard deviation s as follows:

$$s = \sqrt{\frac{p(1-p)}{n}}$$

EXHIBIT 10.5 Summary of U.S. Treasury Yield Curve Shifts Type of Daily Rate Shifts, January 2, 1962, to August 22, 2011

Type of Shift	U.S. Treasury Input Data	Monthly Zero-Coupon Bond Yields	Monthly Forward Rates
Negative shift	2,276	1,486	297
Positive shift	2,306	1,496	327
Zero shift	88	88	88
Subtotal	4,670	3,070	712
Twist	7,716	9,316	11,674
Grand Total	12,386	12,386	12,386

EXHIBIT 10.6 Summary of U.S. Treasury Yield Curve Shifts Percentage Distribution of Daily Rate Shifts, January 2, 1962, to August 22, 2011

Type of Shift	U.S. Treasury Input Data	Monthly Zero-Coupon Bond Yields	Monthly Forward Rates
Negative shift	18.4%	12.0%	2.4%
Positive shift	18.6%	12.1%	2.6%
Zero shift	0.7%	0.7%	0.7%
Subtotal	37.7%	24.8%	5.7%
Twist	62.3%	75.2%	94.3%
Grand Total	100.0%	100.0%	100.0%

where p is the measured consistency ratio and n is the sample size, 12,386 observations. The chart in Exhibit 10.7 shows the standard deviations for each of the three data classes and the number of standard deviations between the measured consistency ratio p and the ratio (100 percent) that would prevail if the one-factor term structure model hypothesis were true.

We reject the hypothesis that the consistency ratio is 100 percent, the level that would prevail if one-factor term structure models were an accurate description of yield curve movements. The hypothesis that one-factor term structure models are “true” is rejected with an extremely high degree of statistical significance, by more than 100 standard deviations in all three cases.

There are a number of very serious errors that can result from an interest rate risk and asset and liability management process that relies solely on the assumption that one-factor term structure models are an accurate description of potential yield movements:

1. Measured interest rate risk will be incorrect, and the degree of the error will not be known. Using the data above, 62.3 percent of the time actual yield curves will show a twist, but the modeled yield curves will never show a twist.
2. Hedging using the duration or one-factor term structure model approach assumes that interest rate risk of one position (or portfolio) can be completely eliminated with the proper short position in an instrument with a different maturity. The duration/one-factor term structure model approach assumes that if interest rates on the first position rise, interest rates will rise on the second position as well; therefore, “going short” is the right hedging direction. (We discuss this in Chapter 11.) The data above shows that on 62.3 percent of the days from 1962 to 2011, this “same direction” assumption was potentially false—some maturities show the same direction changes and some the opposite—and the hedge could actually *add* to risk, not reduce risk.
3. All estimates of prepayments and interest rate–driven defaults will be measured inaccurately.
4. Economic capital will be measured inaccurately.
5. Liquidity risk will be measured inadequately.
6. Nonmaturity deposit levels will be projected inaccurately.

EXHIBIT 10.7 Consistency of U.S. Treasury Yield Daily Shifts with One-Factor Models of the Term Structure of Interest Rates, January 2, 1962, to August 22, 2011

Type of Shift	U.S. Treasury Input Data	Monthly Zero- Coupon Bond Yields	Monthly Forward Rates
Consistency with 1- Factor Models	37.7%	24.8%	5.7%
Standard Deviation of Consistency %	0.4%	0.4%	0.2%
Number of Standard Deviations from 100%	143.1	193.9	450.6

This is an extremely serious list of deficiencies. The remedy is to move to the recommended 6 to 10 factors. Academic assumptions about the stochastic processes used to model random interest rates have been too simple to be realistic. Jarrow (2009) notes:

Which forward rate curve evolutions [HJM volatility specifications] fit markets best? The literature, for analytic convenience, has favored the affine class but with disappointing results. More general evolutions, but with more complex computational demands, need to be studied. How many factors are needed in the term structure evolution? One or two factors are commonly used, but the evidence suggests three or four are needed to accurately price exotic interest rate derivatives.

This view is shared by the Basel Committee of Banking Supervision, which we quoted previously as requiring six factors for market risk analysis. A move to two risk factors does not begin to address the problem, as we explained earlier in this chapter. We now turn to another common pitfall in implementing interest rate risk models.

Common Pitfalls in Asset and Liability Management⁵

Interpolating Monte Carlo Results, or How to Prove Waialae Golf Course in Honolulu is Not a Golf Course With the advances in software design and computer speeds of recent years, risk managers now have very powerful techniques for Monte Carlo simulation at their disposal. That is the reason why the multifactor interest rate framework of HJM is so useful and important. That said, we find many shortcuts have been taken that can provide very misleading estimates of risk levels. This section explains how the misuse of Monte Carlo interest rate simulation can prove that Waialae Golf Course in Honolulu is not a golf course and that the Grand Canyon does not exist.

A common approach to risk measurement can be described by this example, which was explained in a public forum by one of the largest banks in the world:

- We have three factors driving the yield curve: level, shift, and bend.
- We chose no-arbitrage values for each of these factors and choose seven values for each factor, the mean and plus or minus 1, 2, and 3 standard deviations from the mean.
- We then calculate portfolio values for all $7 \times 7 \times 7 = 343$ combinations of these risk factors.
- We calculate our n th percentile value at risk from the smoothed surface of values that results from these 343 portfolio values, using 1 million draws from this smoothed surface to get the 99.97th percentile of our risk.

What is wrong with this approach? At first glance, it certainly seems reasonable and could potentially save a lot of computing time versus a large Monte Carlo simulation like the one outlined earlier in this chapter. Let's look at it out of the finance context. We pose the problem in this way:

Our golf ball is placed in the location marked by the white square in the lower left-hand corner of this aerial photo of a golf course in Honolulu, Waialae Golf Course (Exhibit 10.8). However, we are blindfolded and cannot see. We are told that the ball is either on a green at Waialae Golf Course or that it is not, and our task is to determine which case is the true one.

We calculate the standard deviation of the length of the first putt made by every golfer on every hole at last year's Masters Golf Tournament. From the location of our golf ball, we then get the latitude and longitude of 49 points: our current golf ball location plus all other locations that fall on a grid defined by 1, 2, or 3 standard deviations (of putting length) North-South or East-West with our golf ball in the middle. The grid looks like the one in Exhibit 10.9.

At each point where the North-South and East-West lines intersect, we are given very precise global positioning satellite estimates of the elevation of the ground at that point. Upon reviewing the 49 values, we smooth the surface connecting them and we find that the 99.97th percentile difference in elevation on this smoothed surface is only two inches. Moreover, we find that the greatest difference in elevation among any of the points on the grid is only slightly more than two inches. We know by the rules of golf that a golf hole must be at least four inches in depth. We reject the hypothesis that we are on a green. *But we are completely wrong!*

This silly example shows the single biggest pitfall of interpolating between Monte Carlo values using some subset of a true Monte Carlo simulation. For the same reasons, using the HJM bushy tree is also not the best practice either when Monte Carlo is available (with the exception of selected securities valuations, which we discuss later in this book). In the Waialae example, we have assumed that the surface connecting the grid points is smooth, but the very nature of the problem contradicts this assumption. The golf green is dramatically "unsmooth" where there is a hole in the green. Our design of the Monte Carlo guaranteed a near certainty that our analysis of "on a green"/"not on a green" would be no better than flipping a coin.

Similarly, we could have been placed blindfolded on a spot somewhere in the United States and told either we are within 100 miles of the Grand Canyon or we are not. We again form a grid of nine points, all of which are plus or minus 100



EXHIBIT 10.8 Waialae Golf Course




EXHIBIT 10.9 Grid

miles North–South or East–West of where we were placed blindfolded, as shown in Exhibit 10.10.

Again, we take the GPS elevation at these nine grid points, smooth the surface, and take the 99.97th percentile elevation difference on the smoothed surface. We find that our elevation difference is only 100 meters at the 99.97th percentile. We reject the hypothesis that we are within 100 miles of the Grand Canyon since we know the average depth of the Grand Canyon is about 1,600 meters. Again, we are grossly wrong in another silly example because our assumption that intervening values between our grid points are connected by a smooth surface is wrong.



EXHIBIT 10.10 Grand Canyon Grid

SUMMARIZING THE PROBLEMS WITH INTERPOLATED MONTE CARLO SIMULATION FOR RISK ANALYSIS

For relatively simple, low-dimensional systems that arise in physics and engineering, where correlations are well controlled even in the tails, a reliable and believable description of the asymptotic behavior of the global probability density function may indeed be possible, using a phrase like “smooth” and multifactor smoothing in an extension of Chapter 5. In high-dimensional cases, such as in the risk management of financial institutions, things are much more complex with respect to the global probability density function of values, cash flows, and net incomes. Nothing illustrates that more clearly than the \$2 billion losses announced by JPMorgan Chase on May 8, 2012. The previous simple examples illustrate two problems that can be extremely serious problems if risk is measured by interpolating between Monte Carlo valuations on a grid or small subset of possible outcomes (like the branches of an HJM bushy tree or a Hull and White interest rate lattice):

- *Smoothness is an assumption, not a fact*, as our simple examples show. Interpolating yield curves using a smoothness criterion is reasonable for good economic reasons, but as we show next, it is highly likely that portfolio values and cash flows in financial institutions have sharp discontinuities.
- *The n th percentile “worst case” derived from the simulation is not relative to the true “worst case,” it is relative to the worst case on the grid* defined by the analyst. In our Grand Canyon example, the “worst case” elevation differential was 100 meters, but the true “worst case” elevation difference was the true depth of the Grand Canyon, 1,600 meters.

Almost the entire balance sheet of financial institutions is filled with caps, floors, prepayment risk, and defaults. In the JPMorgan Chase 2012 case, the balance sheet was filled with huge credit default swap positions taken by the trader called the “London Whale.” One cannot assume that the law of large numbers will make these 0/1 changes in values and cash flows vary in a smooth way as the number of transactions rises. We can show that with a simple example. Assume it is February 2000. Three-month U.S. dollar LIBOR is 6.00 percent. The five-year monthly standard deviation of three-month U.S. dollar LIBOR from March 1995 to February 2000 was 32 basis points. We have only three transactions on our balance sheet:

- The first digital option pays us \$1 if three-month LIBOR is 6.25 percent or higher in 18 months.
- The second digital option pays us \$1 if three-month Libor is 6.05% or less in 18 months
- The third digital option requires us to pay \$100 if three-month LIBOR is below 4.00 percent in 18 months.

We set up our grid of seven values for three-month LIBOR in 18 months centered around the current level of LIBOR (6.00 percent) and plus and minus 1, 2, and 3 standard deviations. We do valuations at these LIBOR levels in 18 months:

Scenario 1: 6.96 percent LIBOR

Scenario 2: 6.64 percent LIBOR

Scenario 3: 6.32 percent LIBOR

Scenario 4: 6.00 percent LIBOR

Scenario 5: 5.68 percent LIBOR

Scenario 6: 5.36 percent LIBOR

Scenario 7: 5.04 percent LIBOR

Our valuations of our portfolio are \$1.00 in every scenario:

- Scenario 1: 6.96 percent LIBOR means we earn \$1 on digital option 1
- Scenario 2: 6.64 percent LIBOR means we earn \$1 on digital option 1
- Scenario 3: 6.32 percent LIBOR means we earn \$1 on digital option 1
- Scenario 4: 6.00 percent LIBOR means we earn \$1 on digital option 2
- Scenario 5: 5.68 percent LIBOR means we earn \$1 on digital option 2
- Scenario 6: 5.36 percent LIBOR means we earn \$1 on digital option 2
- Scenario 7: 5.04 percent LIBOR means we earn \$1 on digital option 2

Typical interpolation of these Monte Carlo results would smooth these calculated values. Ardent interpolation users then draw from this smoothed surface. (We have heard up to 1 million draws from this interpolated surface.) In this case, all values on the grid are 1, the smoothed surface is 1 everywhere, and we announce proudly “With 99.97% confidence, we have no risk because we earn \$1 in every scenario.”

Wrong! We earn nothing in the range between 6.05 percent and 6.25 percent. We completely missed this because we didn't sample in that range and assumed cash flows were smooth, but they weren't! Even worse, the actual value of three-month LIBOR in August 2001 was 3.47 percent, and we had to pay out \$100 on digital option 3!

This is a simple example of a very real problem with valuing from a small number of grid points instead of doing a true Monte Carlo simulation. One of the authors remembers a scary period in the mid-1980s when typical interest rate risk simulations never simulated interest rates high enough to detect the huge balance of adjustable rate mortgages in bank portfolios with interest rate caps at 12.75 percent and 13.00 percent. The young analysts, with no knowledge of history and no examination of the portfolio at a detailed level, did not detect this risk, just as our previous simple example missed medium-sized risk within the grid and huge risk outside the grid.

This graph of the 10-year and 5-year moving average standard deviations of month-end three-month LIBOR in Exhibit 10.11 shows how easy it is to design “grid widths” that end up being horribly wrong.

Take the 32 basis point standard deviation of three-month LIBOR over the five years ended in February 2000. The standard deviation of three-month LIBOR for the five years ended in September 1981 was 450 basis points!

Going back to our Waialae Golf Course example, if we don't know the true width of the hole on a golf course, we at least know that 49 grid points is not enough. We do 1,000 or 10,000 or 100,000 random North–South/East–West coordinates



EXHIBIT 10.11 Five-Year and Ten-Year Simple Standard Deviations of Month-End Values for U.S. Dollar Three-Month LIBOR Rate

Source: Federal Reserve Board, H15 Statistical Release

looking for an elevation change of at least four inches (the depth of a golf hole). The degree of confidence we seek determines the number of random coordinates we draw. If we know by the rules of golf that the whole is exactly 4.25 inches wide, we can be much more refined in our search techniques. The trouble in risk management, however, is that we don't know if we're looking for a golf hole (a few mortgage defaults in 2005) or the Grand Canyon (massive mortgage defaults from the huge drop in U.S. and UK home prices in 2007–2009).

Given the current state of risk systems technology, there is no reason not to do a true Monte Carlo simulation that defines the amount of risk embedded in the portfolio with any reasonable desired level of precision. The shortcut of interpolating between a small subsample of scenarios is extremely dangerous. The full power of N -factor interest rate modeling using the HJM approach deserves a full Monte Carlo simulation.

We now turn to interest rate mismatching and hedging in a modern framework.

NOTES

1. For a video of the daily movements of the forward rate curve, see www.kamakuraco.com or YouTube.com, searching for “Kamakura.”
2. The formulas are given in Jarrow (2002, 280).
3. See Jarrow (2002, 279) for details.
4. For two excellent summaries of the literature on term structure models, see Duffie and Kan (1994) and Jarrow (2009).
5. The authors would like to thank colleagues Dr. Predrag Miocinovic and Dr. Alexandre Telnov for their contributions to this section.

Interest Rate Mismatching and Hedging

In Chapters 1 through 4, we introduced the concept of the interest rate risk safety zone. The “safety zone” is that set of portfolio strategies such that it is impossible for the institution to fail to achieve its mission. In banking, for example, a new bank with \$10 billion in capital might fund \$90 billion in fixed rate loans (with no credit risk) for five years with \$90 billion in five-year fixed rate borrowings. The \$10 billion remaining from the capital raise could be invested in any maturity from 1 day to 30 years without changing the default risk of the organization, provided that the assets concerned have no credit risk. This set of potential zero failure strategies on the 10 billion of capital constitutes the “safety zone” for this new bank. Within this safety zone, shareholders of the bank should be indifferent as to what degree of interest rate risk is taken, because the shareholder can take offsetting positions for the shareholder’s own account to achieve the interest rate the shareholder (not the bank) prefers. The assessment of the safety zone gets more complex when we allow for credit risk, an important topic in this book.

We return to this safety zone concept frequently as we discuss interest rate risk and the implications of interest rate risk for the total risk, including credit risk, of the institution. In this chapter, we focus on the strategy and tactics of changing the interest rate risk of a financial institution from both a political and economic point of view, keeping in mind the lessons of the interest rate simulation in Chapters 6 through 10. It is these more modern techniques that will allow actual implementation of the Merton-Jarrow approach to risk management, using the put option concept as a comprehensive measure of risk for the institution. We compare them to legacy approaches to interest rate risk management in Chapters 12 and 13 after setting the stage in this chapter. Our focus in this chapter is quite simple: When should interest rate risk be hedged? When is it not necessary to hedge interest rate risk?

We start by reviewing the kinds of analytical issues and political issues that frequently arise when making a decision to change the interest rate risk of an institution. By *institution*, we mean any legal entity that has interest rate risk, including (but not limited to) pension funds, mutual funds, hedge funds, banks, insurance firms, university endowments, credit unions, special purpose vehicles and so on.

POLITICAL FACTIONS IN INTEREST RATE RISK MANAGEMENT

In Chapters 2 and 3, we talked about how performance measurement systems differ by type of institution. These institutional differences have a big impact on the politics of measuring and changing the interest rate risk of a financial institution.

Pension Fund Considerations

As we discussed in Chapters 2 and 3, a pension fund is typically divided between the asset side of the organization and the liability side of the organization. The liability side of a pension fund has a heavy actuarial component, and expert actuaries lay out the most likely benefits that must be paid to the beneficiaries of the pension fund. The Board of Directors of the organization then establishes the “best mix” of various asset classes that the fund will invest in. Pension funds that consist solely of “defined contribution” plans instead of “defined benefit” plans have the easiest task, because in the former case the beneficiary of the pension selects the fund’s assets. In the defined benefit case, however, the asset classes selected include:

- *Fixed income securities*, which are designed to provide a steady amount of cash flow toward the pension obligations. The manager of this subportfolio is typically measured against a naïve benchmark like the Barclays (pre-2008, “Lehman Brothers”) Government Bond index.
- *Equity securities*, which are designed to provide (but sometimes don’t, as in the post-bubble Japan case) a higher long-run expected return than fixed income securities while at the same time partially capturing any benefits from an inflationary environment to help offset a rise in pension obligations that inflation would trigger. The manager of such a fund would be measured versus an index like the DAX, Topix, or S&P 500 equity index.
- *Commodities and real estate*, which are almost exclusively designed to capture returns from inflation to offset an increase in pension benefits that occur in an inflationary environment.

The amount of risk that the fund has is analyzed on two levels. The first level is at the macroeconomic level, where the decision is made about what the proper share of the total portfolio is that should go to these three components. The second level is to analyze the level of risk of the subportfolio versus the naïve benchmark that is selected as the performance benchmark for the subportfolio. The total amount of risk taken by the organization is largely determined by the first of these two decisions. The Jarrow-Merton put option framework offers considerable benefits in quantifying the many factors that can cause a mismatch in returns on the asset side and obligations on the liability side (inflation, interest rate movements, changes in the life expectancy of the pension beneficiaries, movements in home prices, building prices, commodity prices, and so on). This is an extremely important issue that most sophisticated pension funds are just beginning to quantify. There will be enormous progress on this issue over the next decade that will involve simultaneous analysis of mark-to-market technology and multiperiod simulation, techniques, which we discuss in the rest of this book. What does the Jarrow-Merton put option measure in the pension fund case? The expected risk-neutral present value of the cost of bailing out

the pension fund when it does not have sufficient funds to meet its promised “defined benefit” obligations.

Things are more manageable at the subportfolio level. Consider the fixed income subportfolio. The manager of the fixed income subportfolio will have “tracking error” versus the Barclays government bond index in two dimensions. One contribution to tracking error comes from the potential for default if the manager is buying fixed income securities with credit risk. The second contribution to tracking error comes from the mismatch in the maturity structure of cash flows in his actual portfolio versus the cash flow associated with the components of the Barclays government bond index.

The politics of what fixed income portfolio tracking error to accept is more straightforward in pension funds than it is in most institutions because of the relatively small number of people involved. Normally, the Board of Directors will specify the maximum tracking error versus the benchmark index and the head of fixed income (or a small committee) will decide where the portfolio should fall within that tracking error benchmark. This will be determined in many cases by advantages in execution rather than by outright interest rate “bets.”

How about the risk-return trade-off at other institutions?

Life Insurance Companies and Property and Casualty Insurance Companies

Changes in the level of interest rate risk in a life insurance company share many of the characteristics of the pension fund decision process. The liability side is analyzed by actuarial staff and the head of the investment department, supervised by the Board of Directors, is given specific performance benchmarks whose performance he must exceed while keeping tracking error within preset limits. Most of the politics of risk taking in an insurance company is focused on the credit risk of the benchmark fixed income portfolio and who bears credit losses—if the actuaries ask for a BBB-rated return but the investment department suffers any related credit losses, the incentive system is flawed.

One other complexity is found in the life insurance business, which pension funds don’t face so directly. Insurance companies must report earnings quarterly or semi-annually, at least to regulators and often to outside shareholders. The level of scrutiny is much higher for life insurance firms than it is for pension funds except in the worst of circumstances, as in the credit crisis of 2006–2011.

In that sense, the insurance company risk-return trade-off involves both a market-to-market orientation (“How does the value of the Jarrow-Merton put option change if we take less risk”) and a net income orientation. We turn to that issue in our discussion of the banking industry in the next section.

Property and casualty insurance companies have still another added complexity—the payoffs on the liability side are much more volatile from year to year because many important classes of property and casualty insurance losses are much more volatile than payouts on life insurance policies. This brings into play many of the liquidity risk considerations that we discuss in Chapter 37.

State-of-the-art risk management at insurance companies is moving rapidly toward an enterprise-wide view of risk, in which both sides of the balance sheet are managed jointly. This is an essential move toward best practice.

Commercial Banks

The politics of interest rate risk-return strategy in commercial banking is the most complex because all of the complications mentioned previously (with the exception of property and casualty insurance losses) are present and the number of players in the decision is very large. A partial list includes the following parties to the interest rate risk and return strategy decision:

- Board of directors of the bank
- Chief executive officer
- Chief financial officer
- Chief risk officer
- Managing committee of the bank, for example asset and liability management committee (ALCO)
- Heads of major business units
- Day-to-day manager of the transfer pricing book, which includes all internal funds transfers among business units
- Manager of the trading floor
- Bank regulators

The politics is particularly complex because (unlike pension funds, for example) many of the decision makers are decision makers not because of their expertise in interest rate risk but because of the size of the businesses they run. Someone who is running a network of 3,000 retail branches is too busy to be an expert in interest rate risk strategy, but they are often still a voter in major interest rate risk decisions.

In Chapter 2, we introduced the impact that a “common practice” transfer pricing system has on bank interest rate risk strategy. We continue that discussion later in this book when we discuss performance measurement in more detail.

- Day-to-day interest rate risk mismatches are moved via the transfer pricing system to a central transfer pricing book, where the manager has interest rate risk limits consistent with his seniority and the limits he would have if he were trading bonds for the bank.
- Huge macro mismatches are the responsibility of the ALCO and must be approved by the Board.
- Basis risk from nonmarket sensitive pricing indexes (such as the prime rate in many countries or the average cost of funds of the 12th district savings and loan associations in the United States, once an extremely important pricing index for home mortgages) is born by the heads of line units.

Within this general structure, a number of political games and tactics exist to advance the career aspirations or personal wealth of one manager instead of the aspirations of the shareholders. These are too numerous (and too humorous) to cover in depth, so we just hit the highlights:

- “*I can predict interest rates, why can’t you?*” This strategy is the oldest in the book. A member of senior management claims he can predict interest rates and urges the transfer pricing book management to take on a huge interest rate risk

position. He does. If rates move as predicted, the senior manager takes the credit. If rates move in the opposite direction, the transfer pricing book manager takes the blame. This is combated with the limits structure on the transfer pricing book and a separate “book” for the ALCO committee where every position is carefully recorded and the advocates noted by name.

- *“My assets are liquid so give me a short-term interest rate.”* This is the second-oldest ploy in the book, usually argued by the head of the trading floor or the head of a portfolio that can be securitized. It was an argument that almost certainly contributed to the failures of Bear Stearns and Lehman Brothers in 2008. The advocates of this position want to avoid paying a matching maturity cost of funds, even on a floating-rate basis, as we discuss in detail in Chapter 37 on liquidity risk. This political strategy is combated by saying “no” and following a more modern approach to transfer pricing, which we discuss in Chapters 36 to 41. On the argument that “my assets are liquid and we can securitize,” we note that a sharp rise in interest rates puts the mark-to-market value of a fixed rate loan portfolio below par value—very few institutions will securitize if it means recognizing this mark-to-market loss. The credit crisis of 2006–2011 made it very clear that securitization markets can completely disappear if the integrity of a market (like the subprime home loan market in the United States) is suddenly called into question.
- *“I want the long-run historical weighted average cost of funds, not the marginal cost of funds.”* If rates are rising, this is a common request by the asset side of the organization. If rates are falling, the liability side of the organization wants the long-run weighted average cost of funds. This is why the marginal cost of funds system has to be very comprehensive or this kind of petty gaming of the transfer pricing system is endless.

If the performance measurement system is designed correctly, politics (for the most part) falls by the wayside and the relevant decision makers can focus on the right risk-return trade-offs.

It is to that task we now turn.

MAKING A DECISION ON INTEREST RATE RISK AND RETURN: THE SAFETY ZONE

There are two environments in which management has to make a decision about whether to change the interest rate risk of the institution.

- Inside the safety zone
- Outside the safety zone

If the institution is still inside the safety zone that we discussed above, then shareholders don’t care (as a group) what the interest rate risk position is. They don’t care because they can form a portfolio of the institution’s common stock and fixed income securities for their own account that exactly offsets any decisions made by the management of the financial institution and gets the combined “pseudo-financial institution” exactly where the individual shareholder would like it to be from an

interest rate risk perspective. If the financial institution is inside the safety zone, we can focus on the pure economics of changing the interest rate risk position.

If the financial institution is outside the safety zone, shareholders are being forced to suffer some of the expected costs of the potential bankruptcy of the institution. The interest rate risk-return estimation should include the reduction of these expected losses from bankruptcy as part of the calculated benefits of reducing risk. As this book was written, the legal costs and advisory costs associated with the bankruptcy of Lehman Brothers are still being tallied, but they are certain to be in the billions of dollars. An advocate of an interest rate risk strategy that makes the likelihood of incurring these costs higher should include these potential bankruptcy costs in the pros and cons of the strategy suggested.

OBVIOUS INTEREST RATE RISK DECISIONS

Some decisions about interest rate risk have obvious “yes or no” decisions, while others are more complex. Some of the obvious decisions are listed here:

1. Is the put option that eliminates the interest rate risk of the institution in the Jarrow-Merton sense from Chapter 1 so expensive that the financial institution can't afford to buy it? If that is the case, the interest rate risk of the institution is way too large and has to be reduced by other means, like a capital issue.
2. If a change in the cash flow structure from an interest rate-related transaction smooths net income but leaves the Jarrow-Merton put option value unchanged, then the transaction is worth doing. This has an information benefit to shareholders who might otherwise misperceive the risk of the institution based on short-run changes in net income.
3. If the change in cash flow structure from an interest rate-related transaction smooths net income but leaves the bank well within the safety zone, then it is probably worth doing.

ASSESSING THE RISK AND RETURN TRADE-OFFS FROM A CHANGE IN INTEREST RATE RISK

For all other changes in the financial institution's interest rate risk and return policy, there are a number of tools for quantifying the benefits of a change in the interest rate risk position of the institution. The two most important components of the benefits of a change in the interest rate risk of an institution consist of the following:

- The change in cost of the Jarrow-Merton put option, which would be necessary to completely eliminate the financial institution's interest rate risk. Most of the rest of this book is devoted to valuing this put option, using the tools of Chapters 6 through 10 and the tools of later chapters on credit and market risk.
- The change in the expected costs of potential bankruptcy of the institution, which we discussed above.

These two issues are critical to deciding when to hedge and when not to hedge. If the financial institution is deep inside the safety zone, we don't need to spend much time on this analysis. If the bank is near the boundaries of the safety zone or outside it, this analysis becomes very important.

In the next chapter, we introduce traditional or "legacy" tools for supplementing the Jarrow-Merton put option analysis and analysis of the expected costs of bankruptcy. In the view of the authors, the worst of these legacy tools have been replaced by the Heath, Jarrow, and Morton (HJM) techniques of Chapters 6 through 10. The best of the legacy tools can be shown to be special cases of the HJM approach, but, in almost every case, they are too simple to be used accurately in practice. We explain why in the next chapter.

Legacy Approaches to Interest Rate Risk Management

In this chapter, we introduce a number of legacy interest rate risk concepts for one purpose: to aid in the reader's communication with his or her colleagues who are still fond of these techniques. We believe that the time has come and gone for the techniques discussed in this chapter: gap analysis, simple simulation methods, duration, and convexity. As we go through this chapter, we will be liberal in adding the heading/prompt "Model Risk Alert" when an analysis involves assumptions that are false and potentially very harmful from a career standpoint and an accuracy point of view. For readers who have no interest in the horse-and-buggy days of interest rate risk management, we recommend that you skip this chapter except for the section that compares a modern hedge using a multifactor Heath, Jarrow, and Morton (HJM) modeling approach with the legacy duration approach.

GAP ANALYSIS AND SIMULATION MODELS

Uyemura and van Deventer (1993) discuss the "80/20 rule" in the context of risk management. Getting 80 percent of the benefits for 20 percent of the work was the right strategy in 1993, when most serious risk analysis on an enterprise-wide basis was still done on mainframe computers because of the severe limitations of the 16-bit operating systems of the personal computers of that era. This rule of thumb is probably still a good one in some senses, but it is simply an excuse for laziness when it involves the computer science aspects of risk management. The state of computer science and risk management software has evolved so far in the last decade that 98 percent of the benefits can now be achieved with 2 percent of the effort required for a perfect answer. This is particularly true with the rapid market penetration of advanced personal computers/servers running a 64-bit operating system, which can address 4.3 billion times (2^{32}) more memory than the 32-bit operating systems they replace. We keep the distinction between the 98/2 rule and the 80/20 rule in mind through this chapter, where we review legacy interest rate risk management tools for only two distinct purposes: communication with senior management (essential) and exposition (not decision making) of risk issues in day-to-day interest rate risk operations.

MEASURING INTEREST RATE RISK: A REVIEW

In earlier chapters, we discussed how financial institutions can measure risk in a number of ways:

- *Jarrow-Merton put option*: the value of a put option that would eliminate all of the interest rate risk (the downside risk only) on the financial institution's balance sheet for a specific time horizon. We expand this concept in later chapters to include all risks, with a special focus on credit risk.
- *Interest rate sensitivity of the market value of the financial institution's equity*: often proxied by the mark to model valuation of the financial institution's assets less the value of its liabilities.
- *Interest rate sensitivity of the financial institution's net income (or cash flow) over a specific time horizon*.

Model Risk Alert

The short time horizons of this analysis truncate visibility of risk, and generally accepted accounting principles drastically distort true cash flow and obscure risk. The authors recommend cash flow simulation over net income simulation, with the use of the latter recommended only if required by regulators or for communication with an unsophisticated audience (from a risk management perspective) such as nonfinancial executives in senior management.

In Chapters 13 through 41, we emphasize a modern approach to each of these risk measures with respect to all risk: interest rate risk, market risk, liquidity risk, and credit risk. Our emphasis in this chapter is solely on tools in use since the 1938 introduction of the Macaulay duration concept. A sophisticated reader should review this chapter as if it is an exercise in archaeology, with the exception of the HJM hedging example.

LEGACY RATE RISK TOOLS: INTEREST RATE SENSITIVITY GAP ANALYSIS

In earlier chapters, we discussed the concept of the interest rate risk management safety zone in which it is impossible for the financial institution to go bankrupt from interest rate risk. Every senior officer of a financial institution needs to know the answers to a number of questions related to the safety zone concept:

- How do we know whether the financial institution is in the safety zone or not?
- If we are out of the safety zone, how far out are we?
- How big should my interest rate risk hedge be to get back in the safety zone?
- How much of a move in interest rates would it take to cause bankruptcy from an interest rate risk move?

Interest rate sensitivity “gap” analysis was the first attempt that most financial institutions made to answer these questions. Today, financial institutions use this

expository technique to communicate with senior management team members who do not have a background in finance or risk management. Gap analysis should not be used for decision making because its implications are too crude to be accurate. Gap analysis is too subject to arbitrary assumptions that can literally change the measured impact of higher interest rates on the financial institution from positive to negative or vice versa. Gap analysis does have a 30,000-foot-level perspective, but there is no reason in the twenty-first century to use a Sopwith Camel, the World War I biplane fighter, to hit risk management targets when risk management cruise missiles are available.

How was an interest rate sensitivity gap supposed to work? Consider a simple example in which \$1,000 is invested in a floating-rate loan with a floating interest rate that resets semiannually. We assume the maturity of the loan is three years. The loan is financed with \$800 in a matching maturity, six-month certificate of deposit and \$200 in equity since we assume the bank was newly formed. Assume that the interest rate on the loan is set at 2 percent over the six-month certificate of deposit (all-in) cost and that this cost changes from 6 percent in the first period to 4 percent to the second period. How would this balance sheet look in a standard interest rate sensitivity gap table as shown in Exhibit 12.1?

We show the \$1,000 in loans in the six-month “time zone,” not its actual maturity because the interest rate resets in approximately six months. The \$800 in certificates of deposit goes in the same time bucket.

What about the effective maturity of equity? While some large financial institutions believe they know what the effective maturity of equity is, independent of the actual balance sheet, most sophisticated practitioners and essentially all (if such a thing is possible) academics believe that the effective maturity of the financial institution’s capital is derived from the net of the impact of assets less the impact of nonequity liabilities. Therefore, the mismatch we show in month six of \$200 is the effective maturity of the common stock of this financial institution.

If the financial institution had just issued \$200 in equity for cash, the financial institution could have made \$800 in loans funded with \$800 in CDs and then invested the \$200 in equity at any maturity. If the proceeds were invested in overnight funds, the interest rate mismatch would look like the one in Exhibit 12.2.

This financial institution is “asset sensitive” because it has more short-term asset exposure than it does liabilities. Without doing any multiperiod simulation, we know from this simple table that net income will rise sharply if interest rates rise.

What if equity instead were invested in long-term securities? Then the interest rate sensitivity gap would look like Exhibit 12.3.

The net income of the financial institution will not change in response to short-run changes in interest rates because the capital of the institution is invested in long-term securities.

THE SAFETY ZONE

All of the interest rate risk positions of the financial institutions discussed so far are in the safety zone. The financial institution can never go bankrupt from interest rate risk alone because all short-term liabilities are invested in matched maturity assets. The capital of the bank can be invested at any maturity. What does this imply for

EXHIBIT 12.1 Standard Interest Rate Sensitivity Gap Analysis

Balance Sheet Category	1 Day	2–7 Days	8–30 Days	2 Months	3 Months	4 Months	5 Months	6 Months	12 Months	2 years	3–5 years	Over 5	Other	Total
Assets														
Commercial Loan								1000						1000
	0	0	0	0	0	0	0	1000	0	0	0	0	0	1000
Liabilities														
Certificates of Deposit								800						800
	0	0	0	0	0	0	0	800	0	0	0	0	0	800
Interest Rate Sensitivity Gap	0	0	0	0	0	0	0	200	0	0	0	0	0	200
Cumulative Rate Sensitivity Gap	0	0	0	0	0	0	0	200	200	200	200	200	200	

EXHIBIT 12.2 Standard Interest Rate Sensitivity Gap Analysis: Equity Invested in Overnight Funds

Balance Sheet Category	1 Day	2-7 Days	8-30 Days	2 Months	3 Months	4 Months	5 Months	6 Months	12 Months	2 years	3-5 years	Over 5	Other	Total
Assets														
Overnight Funds	200													200
Commercial Loan								800						800
	200	0	0	0	0	0	0	800	0	0	0	0	0	1000
Liabilities														
Certificates of Deposit								800						800
	0	0	0	0	0	0	0	800	0	0	0	0	0	800
Interest Rate Sensitivity Gap	200	0	0	0	0	0	0	0	0	0	0	0	0	200
Cumulative Rate Sensitivity Gap	200	200	200	200	200	200	200	200	200	200	200	200	200	

EXHIBIT 12.3 Standard Interest Rate Sensitivity Gap Analysis: Long-Term Bonds

Balance Sheet Category	1 Day	2-7 Days	8-30 Days	2 Months	3 Months	4 Months	5 Months	6 Months	12 Months	2 years	3-5 years	Over 5	Other	Total
Assets														
Long Term Bonds												200		200
Commercial Loan								800						800
	0	0	0	0	0	0	0	800	0	0	0	200		1000
Liabilities														
Certificates of Deposit								800						800
	0	0	0	0	0	0	0	800	0	0	0	0	0	800
Interest Rate Sensitivity Gap	0	0	0	0	0	0	0	0	0	0	0	200	0	200
Cumulative Rate Sensivity Gap	0	0	0	0	0	0	0	0	0	0	0	200	200	

the interest rate sensitivity of the financial institution's common stock? As we discussed in Chapter 4, it moves in three different ways for these three balance sheets:

- Case 1: The common stock trades like \$200 in six-month instruments, since capital is invested in the six-month reset loan, plus the present value of the constant spread on \$800 of the loan.
- Case 2: The common stock trades like \$200 in overnight instruments, since capital is invested in the overnight market, plus the present value of the constant spread on \$800 of the loan.
- Case 3: The common stock trades like \$200 in long-term bonds, since capital is invested in the long-term bond market, plus the present value of the constant spread on \$800 of the loan.

Since all of these investments are “at market” at the time they are done, they neither create nor destroy value. They are all in the safety zone, but the financial institution's common stock will respond differently in the three scenarios to changes in interest rates. In no case, however, can the financial institution go bankrupt from interest rate risk so no hedge is necessary.

What if the management thinks shareholders want stable net income? Then they should invest in long-term bonds instead of the overnight market, but either position is in the safety zone and shareholders should be indifferent—as discussed in Chapters 4 and 5, the shareholders can convert the \$200 from overnight money to long-term bonds or vice versa by doing transactions for their own account in combination with their holdings of the financial institution's common stock.

WHAT'S WRONG WITH GAP ANALYSIS?

What is wrong with gap analysis? As an expositional tool for communicating with nonspecialists, nothing is wrong with it.

As a decision-making tool, though, it has serious flaws that led the Office of the Comptroller of the Currency of the United States to stop requiring gap analysis to be reported on the Quarterly Report of Condition for commercial banks in the United States.

Model Risk Alert

We can summarize some of the problems with gap analysis briefly:

- It ignores the coupons and cash flows on all assets and liabilities.
- Gap analysis typically uses accounting values of assets and liabilities, not market values. Two bonds with accounting values of \$100 but market values of \$120 and \$80 respectively are both put in the gap chart at \$100.
- The time buckets typically used are very focused on short-term maturities and tend to be very crude for the long-term interest rate risk, where the real big risk to the institution often lies.
- Gap analysis ignores options—An 8 percent 30-year mortgage may have an effective interest rate sensitivity of 30 years when market interest rates are 12

percent but an effective interest rate sensitivity of 30 days when interest rates are at 3 percent.

- Gap analysis can't deal with other instruments where payments are random, such as property and casualty insurance policies, savings deposits, demand deposits, charge cards, and so on.
- Some interest rates have a lagged adjustment to changes in open market rates, like the prime rate in the United States and Japan and home mortgage indexes in Japan and Australia.
- Many balance sheet items (charge card balances, deposit balances) have a strong seasonal variation—they flow out and flow back in again. Do you classify them by the first outflow date and ignore the seasonal inflow?
- Assets with varying credit risks are mixed, and visibility on credit risk is completely blurred across the balance sheet

Sophisticated risk managers use detailed transaction level, exact day count interest rate risk analysis like that in the remainder of this book. They use gap charts to communicate their conclusions to senior management, if the senior managers are nonspecialists.

LEGACY RATE RISK TOOLS: MULTIPERIOD SIMULATION

Simulation analysis was the next tool to which major financial institutions turned once progress in computer science allowed financial institutions to take a more detailed look at interest rate risk. In the 1970s, most of the focus was on net income for a one- to two-year horizon. This focus was understandable but fatal for thousands of U.S. financial institutions who had interest rate risk of 30-year fixed rate mortgages (the market standard at the time), which one-year net income simulation ignores after the first year. Fortunately, we have much more sophisticated tools available, which we outline in the rest of this book. Nonetheless, it is important to review the kinds of considerations that played an important role in the simulation exercises of major financial institutions over the past four decades. One of those pitfalls we outlined in Chapter 10 is the predilection of some bankers to select interest rate scenarios from an interest rate lattice or bushy tree like those studied in Chapters 6 through 9 instead of a true Monte Carlo simulation. In this section, we are focused on other areas of model risk that stem from a focus on net income simulation.

Key Assumptions in Simulation

Simulation can be done in many ways:

1. Simulate n discrete movements in interest rates with the current balance sheet held constant and produce market values and net income. This is a variation on the assumption behind traditional value-at-risk (VaR), a risk calculation that critics have assaulted with full force in the wake of the \$2 billion in hedge losses announced by VaR sponsor JPMorgan Chase on May 8, 2012.

Model Risk Alert

Assets and liabilities mature, prepay, and default. They are not constant.

2. Simulate the current balance sheet as if it were constant and as if interest rates remained at current levels.

Model Risk Alert

Assuming interest rates do not change implies a riskless arbitrage versus the current yield curve.

3. Move yield curves through multiple specified changes over the forecasting horizon with specific new assets and liabilities added.

Model Risk Alert

While specific scenario analysis is a good supplement to other risk techniques and a frequent requirement of regulators, it is safe on a stand-alone basis only for expositional purposes.

4. Randomly move interest rates over the forecast horizon and dynamically simulate which assets and liabilities will be added as interest rates changes.

Model Risk Alert

The simulation of interest rates that has typically been done by major financial institutions has not been consistent with no arbitrage. The HJM no-arbitrage constraints of Chapters 6 through 9 need to be imposed on any Monte Carlo simulation of rates.

Most financial institutions employ some combination of these techniques, but they do so in the more modern way that we discuss in later chapters.

Even in the more modern style, however, simulation involves a number of key assumptions:

- *Movements in balance sheet amounts:* How will seasonal variation in deposit balances and commercial loans be modeled? How will cash flow thrown off from coupon payments be invested? Will these decisions be independent of interest rate levels or shall they be set as a function of interest rates?

Model Risk Alert

These projections are no more accurate than any other forecast made by humans. One needs to avoid hedging assets and liabilities that one does not have now or that one does not have a legal obligation to originate later.

- *Maturity of New Assets and Liabilities:* What will be the maturity structure and options characteristics of new assets and liabilities? For example, the mix of home mortgage loans between fixed rate loans and floating-rate loans changes substantially over the interest rate cycle. How should this be modeled?
- *Administered Rate Analysis:* How should interest rates, which respond with a lag to changes in open market rates, be modeled?

Model Risk Alert

Many of these open market rates have been historically tied to the LIBOR in major currencies. We discuss lawsuits filed in 2011 and revised in 2012 alleging manipulation of LIBOR rates by market participants in Chapter 17.

- *New business*: How should “new business” (as opposed to cash flow from existing assets rolling over) be modeled?

Data Aggregation in Simulation Modeling

One of the issues in traditional simulation was the degree to which individual transactions should be aggregated for analysis. Many analysts even today think that aggregation is normally a harmless convenience, but more thoughtful risk managers are horrified at the thought. It is done only as a necessary evil that produces some offsetting benefit, like allowing a larger number of Monte Carlo scenarios or faster processing times.

Model Risk Alert

Except in the most extreme circumstances, like Chinese banks with total assets and liabilities consisting of more than 700 million separate transactions, a systems need for data aggregation is an indication that the relevant risk system is very out of date.

In the twenty-first century with 64-bit operating systems, data aggregation serves no purpose except to boost employment at the financial institution’s information technology department. In an earlier era with limited computer speeds and limited memory, data summarization and how to do it was a major issue. In a more modern era, no one would average a \$100,000, 4 percent coupon mortgage with a \$100,000, 6 percent mortgage and say they were the effective equivalent of \$200,000 in a 5 percent mortgage. The entire discipline of market risk is built around the constant search for more accuracy (and, slightly different, higher precision), but in an earlier era in asset and liability management this kind of approximation was necessary under the 80/20 rule we introduced earlier in this chapter. Now that the 98/2 rule is relevant, data stratification is a euphemism for model risk.

Constraining the Model

Assets and liabilities of a financial institution can’t be accurately modeled independently of each other. This is a truism, but, particularly in the insurance industry, the “Great Wall of China” between the liability side and asset side of insurance firms is often a serious impediment to integrated risk management. If the objective of the modeling exercise is a detailed estimate of statutory (insurance industry) or generally accepted accounting principles net interest income, then of course assets and liabilities plus accounting capital have to be equal. On a cash flow basis, the basis preferred by the authors, total cash in must equal total cash out. This is not a trivial exercise in simulation, as the amount of certificates of deposit that a bank, for example, must issue depends on cash thrown off or consumed by every other asset and liability on the

balance sheet. This is determined by insuring that cash in equals cash out, and cash in's "plug" item is the amount of new certificates of deposit or some other residual funding source. Similar concepts apply in pension funds, endowments, mutual funds, hedge funds, property and casualty insurance, and life insurance.

MODELING THE MATURITY STRUCTURE OF A CLASS OF ASSETS

In the early days of simulation modeling, the maturity structure of a pool of assets (like home mortgages) was not visible to the analyst because computer systems were too antiquated to provide that information on a timely basis to management. The team led by Wm. Mack Terry at Bank of America (see Chapters 2 and 38) in the 1970s, for example, was not able to get the exact maturity dates and payment schedules of each mortgage loan at the bank. They had to estimate a "tractor" type roll-off of old mortgages and their replacement by new mortgages.¹

Fortunately, modern financial systems now should be based on exact day count and exact maturity dates at the individual transaction level. This is actually less work than summarizing the data in a way that does not distort the results.

Periodicity of the Analysis

In an earlier era, the periods of the simulation were often "hard coded" to be of equal length. For example, periods could be monthly but the differences in lengths of the month (from 28 days to 31 days) were often ignored. Now analysts specify the exact start date and end date of each period, which can differ in length from period to period. Modern risk analysis is firmly rooted (for the most part) in exact day count technology. Intraday and near real-time analysis is increasingly standard.

Exceptions to the Exact Day Count Trend

An exception to the trend toward exact day count precision of simulation is the brief risk management detour of "value-at-risk" analysis, which we discussed earlier in this chapter and that we describe fairly harshly in Chapter 36.

Model Risk Alert

This single-period analysis effectively ignores all cash flows (and their exact day count) between the date of the analysis and the value-at-risk horizon.

A similar set of assumptions is commonly applied to get quick (and very dirty) answers on instruments like collateralized debt obligations. These simplistic approaches are being rapidly replaced by more sophisticated simulation like that discussed in Chapter 20 on collateralized debt obligations, which many claim was the "financial instrument of mass destruction" at the heart of the 2007–2011 credit crisis.

Legacy Rate Risk Tools: Duration and Convexity

The use of the duration has gone through a curious evolution. For most of the period since 1938, duration was underutilized, for nearly inexplicable reasons, as

an enterprise risk management tool. For the past 15 years or so, the opposite is true. Duration has become the conventional wisdom and now, for equally inexplicable reasons, it is overutilized as a risk management tool in spite of enormous progress in interest rate risk management like the HJM models of Chapters 6 through 9.

As we discussed in Chapter 1, the use of duration as a tool for total balance sheet risk management met with more than five decades of resistance, in spite of the fact that duration and its variants have been standard techniques on trading floors for most of that time. The duration concepts are really the foundation upon which all further advances in interest rate analytics are based. The main model risk now is clinging to duration long after the progress triggered by duration has resulted in far better techniques. On the other hand, from a total balance sheet management perspective, there are still many management teams at very large financial institutions who sincerely believe that net income simulation is a solid basis for risk management and that interest rate risk analytics used on the trading floor either do not apply or are too hard to understand. While many of these CEOs were fired in the aftermath of the 2006–2011 credit crisis, this will be a persistent risk management issue for decades.

In this chapter, we discuss the duration concepts in view of these two constituencies—one group views duration as “too simple” and has moved on to more advanced techniques, and another group views duration as “too hard” to apply. As we discussed in Chapter 1, a lot of this difference in view is due to gaps between the generations and differences in educational experience. As one of the authors is the “older generation,” he advocates the latter explanation!

Over 30 years ago, Ingersoll, Skelton, and Weil (1978) published an appreciation of Frederick R. Macaulay (1938) and his work on duration. In their paper, these authors pointed out a number of key aspects of duration and potential pitfalls of duration that we want to emphasize in this chapter. In what follows, we introduce the concepts of duration and convexity as they have traditionally been used for “bond math.” We then relate them to Macaulay’s 1938 work. Standard practice in financial markets has gradually drifted away from Macaulay’s original concept of duration. Since Ingersoll, Skelton, and Weil published their article, the standard practice has forked into two streams—the traditional adherents of bond math and the term structure model school of thought, which we outlined in Chapters 6 through 10 and that we review from a historical perspective in Chapter 13. The purpose of this chapter is to reemphasize the points made about duration by Ingersoll, Skelton, and Weil and to highlight the errors that can result from deviating from Macaulay’s original formulation of duration. Most of the vocabulary of financial markets still stems from the use of the traditional fixed income mathematics used in this chapter, but the underlying meanings have evolved as financial instruments and financial mathematics have become more complex.

MACAULAY’S DURATION: THE ORIGINAL FORMULA

As in Chapter 4, we use the basic valuation formulas for calculating the present value of a bond.

We define the price of a zero-coupon bond with maturity of t years as $P(t)$. We let its continuously compounded yield be $y(t)$, and we label the cash flow in period t as $X(t)$. As we discussed in Chapter 4, the present value of a series of n cash flows is

$$\begin{aligned}\text{Present value} &= \sum_{i=1}^n P(t_i)X(t_i) \\ &= \sum_{i=1}^n e^{-y(t_i)t_i} X(t_i)\end{aligned}$$

The last line substitutes the relationship between $y(t)$ and $P(t)$ when yields $y(t)$ are continuously compounded. Macaulay investigated the change in present value as each yield $y(t)$ makes a parallel shift of amount x (lowercase), so the new yield at any maturity $y(t)^* = y(t) + x$. The change in present value that results is

$$\begin{aligned}\frac{\partial \text{Present value}}{\partial x} &= \sum_{i=1}^n -t_i e^{-y(t_i)t_i} X(t_i) \\ &= \sum_{i=1}^n -t_i P(t_i)X(t_i)\end{aligned}$$

In earlier chapters, particularly Chapters 9 and 10, we reviewed the Dickler, Jarrow, and van Deventer (2011) result that 94 percent of forward rate curve movements are inconsistent with the parallel shift assumption in the U.S. Treasury market on a daily basis from 1962 to 2011. For the time being, however, let's assume that a parallel shift assumption is good enough that we need to at least measure our exposure to this kind of risk.

Model Risk Alert

The parallel shift assumption is not good enough for interest rate risk management, and we use it in the remainder of this chapter to explain to a modern reader how duration adherents think about risk.

Macaulay defined duration as the percentage change (expressed as a positive number, which requires changing the sign in the equation above) in present value that results from this parallel shift in rates:

$$\begin{aligned}\text{Duration} &= -\frac{\frac{\partial \text{Present value}}{\partial x}}{\text{Present value}} \\ &= \frac{\sum_{i=1}^n t_i P(t_i)X(t_i)}{\sum_{i=1}^n P(t_i)X(t_i)}\end{aligned}$$

From this formula, one can see the reason that duration is often called the “present value-weighted average time to maturity” of a given security. The time to maturity of

each cash flow t_j is weighted by the share of the present value of the cash flow at that time in total present value.

In the case of a bond with n coupon payments of C dollars (note that C is not the annual percentage interest payment unless payments are made only once per year) and a principal amount of 100, this duration formula can be rewritten as

$$\begin{aligned} \text{Duration} &= - \frac{\frac{\partial \text{Present value}}{\partial x}}{\text{Present value}} \\ &= \frac{C \left[\sum_{i=1}^n t_i P(t_i) \right] + t_n P(t_n) 100}{C \left[\sum_{i=1}^n P(t_i) \right] + P(t_n) 100} \end{aligned}$$

This form of duration, the form its inventor intended, has come to be known as Fisher-Weil duration. All discounting in the present value calculations is done at different continuous yields to maturity $y(t_i)$ for each maturity. We know from Chapter 5 that we can get these continuous yields to maturity from the various yield curve-smoothing techniques in that chapter.

USING DURATION FOR HEDGING

How is the duration concept used for hedging? Let's examine the case where a huge financial institution has its entire balance sheet in one unit of a security with present value B_1 . Let's also assume that the financial institution is uncomfortable with its risk level. One thing the financial institution could do is to try to buy the Jarrow-Merton put option that we discussed in Chapter 1. This put option would provide the financial institution with price insurance against the decline in value of this security. The valuation of this put option should be consistent with the valuation methods used for interest rate options in Chapters 6 through 9 and reviewed in more detail in Chapter 24. Let's assume that the financial institution wants to hedge in a more traditional way. Let's assume that the financial institution wants to form a no-risk portfolio that includes the hedging security (which has present value B_2). What amount of this second security should the financial institution use to hedge? Let the zero risk amount of the hedging security be w . Then the total value of this portfolio W is

$$W = B_1 + wB_2$$

We want the change in the value of this hedged portfolio to be zero for infinitely small parallel shifts in the term structure of interest rates. We can call the amount of this shift x . For the hedge to be successful, we must have

$$\frac{\partial W}{\partial x} = \frac{\partial B_1}{\partial x} + w \frac{\partial B_2}{\partial x} = 0$$

From this equation, we can calculate the right amount of B_2 to hold for a perfect hedge as

$$w = -\frac{\frac{\partial B_1}{\partial x}}{\frac{\partial B_2}{\partial x}}$$

Notice that this formula does not directly include duration. We can incorporate the duration calculation by modifying the original equation for a perfect hedge as follows:

$$\begin{aligned}\frac{\partial W}{\partial x} &= B_1 \frac{\frac{\partial B_1}{\partial x}}{B_1} + w B_2 \frac{\frac{\partial B_2}{\partial x}}{B_2} \\ &= -B_1 \text{Duration}[B_1] - w B_2 \text{Duration}[B_2] = 0\end{aligned}$$

Therefore the hedge amount w can be rewritten

$$w = \frac{-B_1 \text{Duration}[B_1]}{B_2 \text{Duration}[B_2]}$$

Using either formula for w , we can establish the correct hedge for the financial institution's portfolio. While we have talked about hedging as if the financial institution owned only one security, this traditional analysis applies equally well to a financial institution that owns millions of transactions whose risk is captured by the artificial "security" we have described above.

Model Risk Alert

This hedging technique has the following characteristics:

- It is correct only for infinitely small parallel shifts in interest rates.
- It has to be "rebalanced" like a Black-Scholes options delta hedge whenever:
 1. Rates change
 2. Time passes

The hedge will not be correct if there are large parallel jumps in interest rates, nor will it be correct for nonparallel shifts in interest rates. In other words, the basic theory of duration does not result in a "buy-and-hold" hedge. This insight is perhaps the model's greatest contribution to the world of finance. At the same time, it leaves many management teams looking for the insurance policy against risk that they can buy once and do not have to deal with again for many years. That insurance policy is the Jarrow-Merton put option.

COMPARING A DURATION HEDGE WITH HEDGING IN THE HJM FRAMEWORK

How does this traditional hedge compare with hedging in the HJM framework of Chapters 6 and 7 (with one risk factor driving interest rates), Chapter 8 (two risk

factors driving rates), or Chapter 9 (three risk factors driving rates)? We summarize the differences here:

Assumption about Yield Curve Shifts

Duration: Parallel

HJM: No restrictions on nature of yield curve movements, which are implied by the number and nature of the risk factors

Size of Yield Curve Movements

Duration: Assumed to be infinitely small

HJM: Can be implemented in either discrete form (Chapters 6 through 9) or in continuous time form with infinitely small yield shifts

Number of Instruments for “Perfect” Hedge Implied by Model

Duration: 1 hedging instrument, and any maturity can be made into a perfect hedge by holding in the correct amount

HJM One-Factor Model: 1 hedging instrument and “cash”

HJM Two-Factor Model: 2 hedging instruments and “cash”

HJM Three-Factor Model: 3 hedging instruments and “cash”

HJM n-Factor Model: n hedging instruments and “cash”

Size of Time Step over Which Interest Rates Change

Duration: Implicitly assumed to be zero

HJM: Equal to length of period in discrete implementation or infinitely small in the continuous time case

In the HJM discrete implementations, the analysis is simple. Since the HJM examples are built on the assumption of no arbitrage, a riskless hedged position will earn the one-period risk-free rate prevailing at time zero. The return on \$1 invested in “cash,” the riskless one period bond, earns $1/P(0, 1)$, where $P(0, 1)$ is the time 0 value of a zero-coupon bond maturing at time 1. Let’s assume that the instrument being hedged is bond B . We write the value of bond B as $B(t, j)$, where t is the then-current time and j is the relevant scenario that prevails. Let’s assume we are at time 0 in Chapter 9 where three risk factors are driving interest rates. At time 1, there are four nodes on the bushy tree, denoting scenarios 1 (shift 1), 2 (shift 2), 3 (shift 3), and 4 (shift 4). The hedging instruments are a position of w_1 dollars invested in cash (the riskless bond maturing at time 1), w_2 dollars invested in the zero-coupon bond maturing at time 2, w_3 dollars invested in the zero-coupon bond maturing at time 3, and w_4 dollars invested in the zero-coupon bond maturing at time 4. We know that the initial investment in bond B and the hedging instruments is

$$\text{Time zero value } Z(0) = B(0, *) + w_1 + w_2 + w_3 + w_4$$

If the value of a zero-coupon bond at time 1 maturing at time T in scenario j is $P(1, T, j)$, then the time 1 value of our hedged portfolio in scenario j will be

$$Z(1, j) = B(1, j) + w_1[P(1, 1, j)/P(0, 1, *)] + w_2[P(1, 2, j)/P(0, 2, *)] \\ + w_3[P(1, 3, j)/P(0, 3, *)] + w_4[P(1, 4, j)/P(0, 4, *)] = Z(0)/P(0, 1, *)$$

The last equality is the no-arbitrage condition that insures that a hedged portfolio earns the time zero risk-free rate. Since there are four unknowns and we have four scenarios, we can solve for the perfect hedge in the three-factor HJM case in Chapter 9. In the two-factor HJM case, we have three unknowns for the three scenarios. In the one-factor HJM case, we have two unknowns for the two scenarios.

Model Risk Alert

Note that the HJM framework allows us to measure the error in a duration hedge precisely. Since only one instrument is used in a duration hedge, it will be a perfect hedge in (at most) only one of the four shifts in a three-factor HJM model. We can calculate the hedging error from a duration strategy by valuing the hedged portfolio just as we valued simpler instruments in Chapters 6 through 9.

Note also that the HJM hedge needs to be rebalanced with each forward movement in time. In general, the hedge amounts will be different at different nodes on the bushy tree, even at the same current time t . This is a common phenomenon in finance, like the continuous delta hedging that is necessary when using the Black-Scholes options model.

DURATION: THE TRADITIONAL MARKET CONVENTION

For many years, the conventional wisdom has drifted over time from Macaulay's original formulation of the duration concept. With the explosion in interest rate derivatives over the past two decades, the original formulation and its more modern relatives dominate among younger and more analytical interest rate risk experts. The evolved Macaulay duration is passing back out of fashion because it relies on old-fashioned "bond math" rather than the Fisher-Weil continuous compounding and the related concepts of Chapter 4. The reason for the change to bond math from Macaulay's original formulation is a simple one: until relatively recently, it has been difficult for market participants to easily take the $P(t)$ zero-coupon bond values from the current yield curve using the techniques that were outlined in Chapter 5. Now that yield curve smoothing has become easier, this problem has declined in importance, but the market's simpler formulation persists in many areas of major financial institutions. The best way to explain the bond math market convention for duration is to review the formula for yield to maturity.

The Formula for Yield to Maturity

For a settlement date (for either a secondary market bond purchase or a new bond issue) that falls exactly one (Model Risk Alert) equal-length period from the next

coupon date, the formula for yield to maturity can be written as follows for a bond that matures in n periods and pays a fixed dollar coupon amount C m times per year as shown in Chapter 4:

$$\text{Price} + \text{Accrued interest} = \text{Present value} = C \left[\sum_{i=1}^n P(t_i) \right] + P(t_n) * \text{Principal}$$

where

$$P(t_i) = \frac{1}{\left(1 + \frac{y}{m}\right)^i}$$

Yield to maturity is the value of y that makes this equation true. This relationship is the present value equation for a fixed coupon bond, with the addition that the discount factors are all calculated by dividing one dollar by the future value of one dollar invested at y for i equal length periods with interest compounded m times per year. In other words, y is the internal rate of return on this bond. Note also that the present value of the bond, using these discount factors, and the price of the bond are equal since there is no accrued interest.

Yield to Maturity for Long or Short First Coupon Payment Periods

For most outstanding bonds or securities, as we saw in Chapter 4, the number of periods remaining usually contains a short (or occasionally a long) first period. The length of the short first period is equal to the number of days from settlement to the next coupon, divided by the total number of days from the last coupon (or hypothetical last coupon, if the bond is a new issue with a short coupon date) to the next coupon. The exact method of counting days for the purpose of this calculation depends on the interest accrual method associated with the bond (examples are actual/actual, the National Association of Securities Dealers 30/360 method in the United States, or the Euro market implementation of the 30/360 accrual method). We call the fraction of remaining days (as appropriately calculated) to days between coupons (as appropriately calculated) x . Then the yield to maturity can be written

$$\text{Present value} = C \sum_{i=1}^n [P(t_i)] + P(t_n) * \text{Principal}$$

where

$$P(t_i) = \frac{1}{\left(1 + \frac{y}{m}\right)^{i-1+x}}$$

As in the case of even first coupon periods, y is the yield to maturity that makes this relationship true, the internal rate of return on the bond. Note that when x is 1, this reduces to the formula above for even-length first periods. We can solve

for y in simple spreadsheet software in a number of ways, including standard yield-to-maturity/internal rate of return formulas in the spreadsheet, or more generally using the solver function.

Applying the Yield-to-Maturity Formula to Duration

The conventional definition of duration is calculated by applying the implications of the yield-to-maturity formula to the original Macaulay duration formula, which is given below:

$$\begin{aligned} \text{True duration} &= - \frac{\frac{\partial \text{ Present value}}{\partial x}}{\text{Present value}} \\ &= \frac{\sum_{i=1}^n t_i P(t_i) X(t_i)}{\sum_{i=1}^n P(t_i) X(t_i)} \end{aligned}$$

We can rewrite this formula for a conventional bond by substituting directly into this formula on the assumptions that market interest rates are equal for all maturities at the yield-to-maturity y and that interest rates are compounded at discrete, instead of continuous, intervals.

Model Risk Alert

This assumption is false except by coincidence. It assumes that the yield curve is flat.

Model Risk Alert

Instead of allowing for time periods of different length as we did with the Macaulay formulation (say semiannual periods of 181 and 184 days, for example), this formulation assumes all periods except the first or last have equal length. This variation of the Macaulay formula results in the following changes:

- The discount factors $P(t_i)$, taken from the smoothed yield curve by Macaulay, are replaced by a simpler formulation that uses the same interest rate (the yield to maturity) at each maturity:

$$P(t_i) = \frac{1}{\left(1 + \frac{y}{m}\right)^{i-1+x}}$$

- The time t_i is directly calculated as

$$t_i = \frac{i - 1 + x}{m}$$

Using this formulation, the “conventional” definition of duration for a bond with dollar coupon C and principal of 100 as

$$\begin{aligned} \text{Conventional duration} &= \frac{C \left[\sum_{i=1}^n t_i P(t_i) \right] + t_n P(t_n) 100}{\text{Present value}} \\ &= \frac{C \left[\sum_{i=1}^n \frac{i - 1 + x}{m} \left(\frac{1}{\left(1 + \frac{y}{m}\right)^{i-1+x}} \right) \right] + \frac{n - 1 + x}{m} \left(\frac{1}{\left(1 + \frac{y}{m}\right)^{n-1+x}} \right) 100}{\text{Present value}} \end{aligned}$$

This is just a discrete time transformation of Macaulay’s formula with the additional false assumptions of a flat yield curve and equal length periods (other than the first or last periods). As shown below, however, this formula does not measure the percentage change in present value for a small change in the discrete yield to maturity, as shown in the next section.

Modified Duration

In the case of Macaulay duration, the formula was based on the calculation of

$$\begin{aligned} \text{Duration} &= - \frac{\frac{\partial \text{Present value}}{\partial x}}{\text{Present value}} \\ &= \frac{\sum_{i=1}^n t_i P(t_i) X(t_i)}{\sum_{i=1}^n P(t_i) X(t_i)} \end{aligned}$$

The second line of the equation immediately above results from the fact that the yield to maturity of a zero-coupon bond with maturity t_i , written as $y(t_i)$, is continuously compounded. The conventional measure of duration, given in the previous section, is not consistent with Macaulay’s derivation of duration when interest

payments are discrete instead of continuous. What happens to the percentage change in present value when the discrete yield to maturity (instead of the continuous yield to maturity analyzed by Macaulay) shifts from y to $y + z$ for infinitely small changes in z ? Using the yield-to-maturity formula as a starting point and differentiating with respect to z , we get the following formula, which is called modified duration:

$$\text{Modified duration} = \frac{C \left[\sum_{i=1}^n \frac{i-1+x}{m} \left(\frac{1}{\left(1+\frac{y}{m}\right)^{i+x}} \right) \right] + \frac{n-1+x}{m} \left(\frac{1}{\left(1+\frac{y}{m}\right)^{n+x}} \right) 100}{\text{Present value}}$$

$$= \frac{\text{Conventional duration}}{1 + \frac{y}{m}}$$

This modified duration measures the same thing that Macaulay's original formula was intended to measure: the percentage change in present value for an infinitely small change in yields.

Model Risk Alert

Unfortunately, the measurement is inaccurate because the assumptions underlying the model are false. The Macaulay formulation and the conventional duration are different for the following reasons:

- The Macaulay measure (most often called Fisher-Weil duration) is based on continuous compounding and correctly measures the percentage change in present value for small parallel changes in these continuous yields. The conventional duration measure, a literal translation to discrete compounding of the Macaulay formula with the additional implied assumption of a flat yield curve, does *not* measure the percentage change in price for small changes in discrete yield to maturity. The percentage change in present value is measured by conventional modified duration. When using continuous compounding, duration and modified duration are the same.
- The Macaulay measure uses a different yield for each payment date instead of using the yield to maturity as the appropriate discount rate for each payment date.

The impact of these differences in assumptions on hedging is outlined in the next section.

THE PERFECT DURATION HEDGE: THE DIFFERENCE BETWEEN THE ORIGINAL MACAULAY AND CONVENTIONAL DURATIONS

Both the original Macaulay (Fisher-Weil) and conventional formulation of duration are intended to measure the percentage change in “value” for a parallel shift in yields.

Model Risk Alert

We remind the reader that yield curve twists on daily movements in the U.S. Treasury curve occurred 94 percent of the time between 1962 and 2011.

Of the remaining shifts, although rates moved up or down together, a small fraction were truly parallel shifts in equal amounts at each maturity. Using duration, the percentage change is measured as the percentage change in present value; it is important to note that it is *not* the percentage change in *price* (present value—accrued interest) that is measured. The difference between a hedge calculated based on the percentage change in present value (the correct way) versus a hedge based on the percentage change in price (the incorrect way) can be significant if the bond to be hedged and the instrument used as the hedging instrument have different amounts of accrued interest. The impact of this “accrued interest” effect varies with the level of interest rates and the amount of accrued interest on the underlying instrument and the instrument used as a hedge. If payment dates on the two bonds are different and the relative amounts of accrued interest are different, the hedge ratio error can exceed 5 percent.

The second source of error, when using modified duration to calculate hedge ratios, is the error induced by using the yield to maturity as the discount rate at every payment date instead of using a different discount rate for every maturity. The hedging error (relative to Fisher-Weil duration) from this source can easily vary from –2 percent to plus 2 percent as a linear yield curve takes on different slopes. The hedge ratio error is obviously most significant when the yield curve has a significant slope.

The third source of error is the dominant yield curve movement: yield curve twists instead of parallel shifts. If one is hedging a five-year bond with a three-year bond and one yield rises while the other yield falls, one can lose money on both sides of the hedge. The only way to avoid this phenomenon is to *avoid* duration models and to use an HJM framework with an accurate number of underlying risk factors.

CONVEXITY AND ITS USES

The word “convexity” can be heard every day on trading floors around the world, and yet it’s a controversial subject. It is usually ignored by leading financial academics, and it’s an important topic to two diverse groups of financial market participants: adherents to the historical yield-to-maturity bond mathematics that predates the derivatives world, and modern “rocket scientists” who must translate the latest in derivatives mathematics and term-structure modeling to hedges that work. This section is an attempt to bridge the gap between these constituencies.

Convexity: A General Definition

The duration concept, reduced to its most basic nature, involves the ratio of the first derivative of a valuation formula f to the underlying security's value:

$$\text{Duration} = \frac{-f'(y)}{f(y)}$$

The variable y could be any stochastic variable that determines present value according to the function f . It may be yield to maturity, it may be the short rate in a one-factor HJM term structure model, and so forth. It is closely related to the delta hedge concept, which has become a standard hedging tool for options-related securities. The underlying idea is exactly the same as that discussed at the beginning of this chapter: for infinitely small changes in the random variable, the first derivative (whether called “duration” or “delta”) results in a perfect hedge if the model incorporated in the function f is “true.” So far in our discussion, the constituencies mentioned above are all in agreement.

As we showed in the section above, the duration and modified duration concepts are *not* true whenever the yield curve is not flat. The academic community and the derivatives hedgers agree on this point and neither group would be satisfied with a hedge ratio based on the discrete yield-to-maturity formulas presented in this chapter, despite their long and glorious place in the history of fixed income markets. All constituencies agree, however, that when the random variable in a valuation formula “jumps” by more than an infinitely small amount, a hedging challenge results.

Leading academics solve this problem by either assuming it away or by breaking a complex security into more liquid “primitive” securities and assembling a perfect hedge that exactly replicates the value of the complex security. This is exactly what we did in the case of the three-factor HJM model hedge of our bond B above. In the case of a noncallable bond, we know that the perfect hedge (perhaps one of many perfect hedges, depending on the number of risk factors driving rates) is the portfolio of zero-coupon bonds that exactly offsets the cash flows of the underlying bond portfolio. Traditional practitioners of bond math and derivatives experts often are forced to deal with (1) the fact that no simple cost-efficient perfect hedge exists; or (2) the fact that required rebalancing of a hedge can only be done at discrete, rather than continuous, intervals.

The fundamental essence of the convexity concept involves the second derivative of a security's value with respect to a random input variable y . It is usually expressed as a ratio to the value of the security, as calculated using the formula f :

$$\text{Convexity} = \frac{f''(y)}{f(y)}$$

It is very closely linked with the “gamma” calculation (the second derivative of an option's value with respect to the price of the underlying security).

When a security's value makes a discrete jump because of a discrete jump in the input variable y , the new value can be calculated to any desired degree of precision using a Taylor expansion, which gives the value of the security at the new input variable level, say $y + z$, relative to the old value at input variable level y :

$$f(y+z) = f(y) + f'(y)z + f''(y)\frac{z^2}{2!} + \dots + f^{[n]}(y)\frac{z^n}{n!} + \dots$$

Rearranging these terms and dividing by the security's original value $f(y)$ gives us the percentage change in value that results from the shift z :

$$\begin{aligned}\frac{f(y+z) - f(y)}{f(y)} &= \frac{f'(y)}{f(y)}z + \frac{f''(y)}{f(y)}\frac{z^2}{2!} + \dots + \frac{f^{[n]}(y)}{f(y)}\frac{z^n}{n!} + \dots \\ &= (\text{Duration})z + (\text{Convexity})\frac{z^2}{2!} + \text{Error}\end{aligned}$$

The percentage change can be expressed in terms of duration and convexity and an error term. Note that modified duration should replace duration in this expression if the random factor is the discrete yield to maturity. We will see in Chapter 13 that this Taylor expansion is very closely related to Ito's lemma, the heart of the stochastic mathematics we introduce in Chapter 13. These techniques, combined with the yield curve-smoothing technology of Chapter 5, form the basis for an integrated analysis of interest rate risk, market risk, liquidity risk, and credit risk, which is booming in popularity today. The HJM framework of Chapters 6 through 9 is the fundamental foundation for this integrated risk framework.

Convexity for the Present Value Formula

For the case of the present value of a security with cash flows of X_i per period, payments m times per year, and a short first coupon of x periods, the convexity formula can be derived by differentiating the present value formula twice with respect to the yield-to-maturity y and then dividing by present value:

$$\text{Convexity} = \frac{\left[\sum_{i=1}^n X(t_i) \left(\frac{-(i-1+x)}{m} \right) \left(\frac{-(i-1+x)-1}{m} \right) \left(\frac{1}{1+\frac{y}{m}} \right)^{-(i-1+x)-2} \right]}{\text{Present value}}$$

When present value is calculated using the continuous time approach,

$$\text{Present value} = \sum_{i=1}^n e^{-t_i y(t_i)} X(t_i)$$

then the continuous compounding or Fisher-Weil version of convexity is

$$\text{Convexity} = \frac{\left[\sum_{i=1}^n t_i^2 P(t_i) X(t_i) \right]}{\text{Present value}}$$

Hedging Implications of the Convexity Concept

To the extent that rebalancing a hedge is costly or to the extent that the underlying random variable jumps rather than moving in a continuous way, the continuous

rebalancing of a duration (or delta) hedge will be either expensive or impossible. In either case, the best hedge is the hedge that comes closest to a buy-and-hold hedge, a hedge that can be put in place and left there with no need to rebalance until the portfolio being hedged has matured or has been sold. A hedge that matches both the duration (or modified duration in the case of a hedger using a present value based valuation formula) and the convexity of the underlying portfolio will come closer to this objective than most simple duration hedges.

Model Risk Alert

Even with a convexity adjustment, however, the implicit assumption is still that interest rate changes are driven only by one parallel yield curve shift factor. Since this assumption is dramatically false, hedging accuracy will be low except by coincidence.

Consider the case of a financial institution with a portfolio that has a value P and two component securities (say mortgages and auto loans) with values S_1 and S_2 . Denoting duration by D and convexity by C , the best (simple) hedge ratios w_1 and w_2 for a hedge using both securities S_1 and S_2 can be obtained by solving the same kind of equations we used previously in the duration case. We need to solve either

$$\begin{aligned} P' &= w_1 S'_1 + w_2 S'_2 \\ P'' &= w_1 S''_1 + w_2 S''_2 \end{aligned}$$

or

$$\begin{aligned} PD_P &= w_1 S_1 D_1 + w_2 S_2 D_2 \\ PC_P &= w_1 S_1 C_1 + w_2 S_2 C_2 \end{aligned}$$

for the hedge ratios w_1 and w_2 .

CONCLUSION

The concepts we have illustrated in this chapter are generally analyzed using a very simple position of a trader or a financial institution—there was generally only one security in the portfolio and only one hedging instrument. In practical application, traders have 10, 50, 100, or 100,000 positions and financial institutions have 50,000, 1 million, or 700 million positions. The process of measuring the change in the value of the position (some of which are assets and some of which are liabilities) is identical whether there is one position or millions. We illustrate the calculation of these interest rate “deltas” or durations in the next few chapters, but we do it with emphasis on the multifactor HJM approach. The calculation of the perfect hedge is also the same from a total balance sheet perspective with one modest variation. As we saw in Chapters 1 through 5, the interest rate risk of the proceeds of the financial institution’s equity is irrelevant to shareholders as long as all of the other assets are funded in such a way that there is a constant hedged cash flow from them. If

this assumption is correct, then the financial institution cannot go bankrupt from interest rate risk and the financial institution is in the “safety zone.” Hedging of the interest rate risk of an institution that is already in the safety zone, regardless of the hedging technique (HJM or duration) used, is simply wasteful because it creates no shareholder value.

Over the past seven decades, the market’s conventional definition of duration has drifted from Macaulay’s original concept to a measure that appears similar but can be very misleading if used inappropriately. Fortunately, the trends in interest rate risk management are now swinging forcefully back to the original Macaulay concept and its implementation in a multifactor HJM interest rate framework. As Ingersoll, Skelton, and Weil (1978) point out, the Macaulay (1938) formulation offers considerable value as a guide to the correct hedge if one takes a back-to-basics approach: the measure should focus on the percentage change in present value, not price, and the discount rate at each maturity should reflect its actual continuous yield to maturity as obtained using the techniques in Chapter 5 and driven by multiple interest rate factors as we did in Chapters 8 and 9. Assuming the yield to maturity is the same for each cash flow or coupon payment on a bond, even when the payment dates are different, can result in serious hedging errors. The convexity approach provides a useful guide to maximizing hedging efficiency, but the term convexity has outlasted the simple convexity formula commonly used for bonds. The formula itself has been supplanted by the HJM hedging techniques that we illustrated previously.

How do we implement these concepts for a bond option? How do we calculate them for a life insurance policy or a first-to-default swap? How do we calculate them for a mortgage that prepays? How do we calculate them for a savings deposit without an explicit maturity? That is a task to which we will turn after a little more foundation in interest rate risk analytics.

NOTE

1. See Chapter 7 in Uyemura and van Deventer (1993) for more on this approach.

Special Cases of Heath, Jarrow, and Morton Interest Rate Modeling

In this chapter, we continue our review of the history of interest rate analytics, but this time we focus on term structure models. We start with the example of parallel yield curves shifts and duration as a naïve term structure model to illustrate the approach academics have typically taken to interest rate modeling:

- Step 1: Make an assumption about how interest rates vary
- Step 2: Impose no-arbitrage restrictions
- Step 3: Derive what the shape of the yield curve must be

The approach taken by Heath, Jarrow, and Morton (HJM) was the exact opposite:

- Step 1: Observe what the shape of the yield curve actually is
- Step 2: Add assumptions about the volatility of rates
- Step 3: Impose no-arbitrage conditions on yield curve movement from its current shape
- Step 4: Derive the nature of yield curve movements from the existing yield curve shape

The power of the HJM approach was illustrated in Chapters 6 through 9, with four different sets of assumptions about interest rate volatility. In this chapter, we demonstrate that the early academic work on term structure models, in every case, is a special case of HJM. By special case, we mean a specific assumption in Step 2 of the general HJM process above.

In Chapter 12, we were introduced to Frederick Macaulay's duration concept. At the heart of the duration concept is the implicit assumption that rates at all maturities move at the same time in the same direction by the same absolute amount. This assumption about how rates move can be labeled a "term structure model." At the same time, it is known on trading floors all over the world as PVBP, that is, the present value of a basis point, a common measure of risk.

The purpose of this chapter is to introduce the traditional academic approach to term structure models and the procedures for evaluating the reasonableness of a

given term structure model. We show in each case that the resulting model is a special case of the HJM approach. We review historical developments in term structure modeling to show readers and adherents to these traditional models that a more realistic approach to interest rate modeling using 6 to 10 risk factors in the HJM framework is a logical step forward from what has been done in the past. We go on to fully implement the HJM term structure model concept in later chapters so that we can calculate the credit-adjusted value of the Jarrow-Merton put option, the best measure of risk, as outlined in Chapter 1.

WHAT IS AN ACADEMIC TERM STRUCTURE MODEL AND WHY WAS IT DEVELOPED?

When analyzing fixed income portfolios, the value of a balance sheet, or a fixed income option, it is not enough to know where interest rates currently are. We demonstrated this clearly in Chapters 6 through 9. It is not possible to know where rates will be in the future with any certainty. Instead, we can analyze a portfolio or a specific security by making an assumption about the random process by which rates will move in the future. Once we have made an assumption about the random (or stochastic) process by which rates move, we can then derive the probability distribution of interest rates at any point in the future. Without this kind of assumption, we can't analyze interest rate risk effectively. As we noted above, the early academic work on term structure modeling ignored the current shape of the yield curve in this derivation.

THE VOCABULARY OF TERM STRUCTURE MODELS

The mathematics for analyzing something whose value changes continuously but in a random process is heavily used in physics, and these tools from physics have been responsible for most of the derivatives valuation formulas used in finance in the past five decades. At the risk of sending some readers searching for the remote control device, we want to spend a little time on these mathematical techniques because of the power they bring to a clear understanding of risk and what to do about it. The tools we describe in this chapter were employed by HJM to derive the no-arbitrage n factor term structure model framework that we employed in Chapters 6 through 9.

In the stochastic process, mathematics used in physics, the change in a random variable x over the next instant is written as dx . The most common assumption, but not the only assumption, about the random jumps in the size of a variable x is that they are normally distributed. The noise or shock to the random variable is assumed to come from a random number generator called a *Wiener process*. It is a random variable, say z , whose value changes from instant to instant, but which has a normal distribution with mean zero and standard deviation of 1. For example, if we want to assume that the continuously compounded zero-coupon bond yield $y(t)$ for a given maturity t changes continuously with a mean change of zero and a standard deviation of 1, we would write

$$dy = dz$$

In stochastic process mathematics, all interest rates are written as decimals, such as .03 instead of 3, so the assumption that rates jump each instant by 1 (which means 100 percent) is too extreme. We can scale the jumps of interest rates by multiplying the noise generating Wiener process by a scale factor. In the case of interest rates, this scale factor is the volatility of the stochastic process driving interest rate movements. This volatility is related to but different from the volatility of stock prices assumed by Black and Scholes to be lognormally distributed. In the Black-Scholes model, the return on stock prices is assumed to be normally distributed, not stock prices themselves. In our simple interest rate model, the change in interest rates is assumed to be driven by this normally distributed “shock term” z with changes dz . We label this interest rate “volatility σ .” In that case, the stochastic process for the movement in yields is

$$dy = \sigma dz$$

The volatility term allows us to scale the size of the shocks to interest rates up and down. We can set this interest rate volatility scale factor to any level we want, although normally it would be set to levels that most realistically reflect the movements of interest rates. We can calculate this implied interest rate volatility from the prices of observable securities like interest rate caps, interest rate floors, or callable bonds. We will show the relationship between interest rate volatility and the volatility of the Black-Scholes model in Chapter 18 on the random interest rates version of the Merton credit model.

The yield that we have modeled above is a random walk, a stochastic process where the random variable drifts randomly with no trend up or down. As time passes, ultimately the yield may rise to infinity or fall to negative infinity. This isn’t a very realistic assumption about interest rates, which usually move in cycles, depending on the country and the central bank’s approach to monetary policy. How can we introduce interest rate cycles to our model? We need to introduce some form of drift in interest rates over time. One form of drift is to assume that interest rates change by some formula as time passes:

$$dy = \alpha(t)dt + \sigma dz$$

The formula above assumes that on average, the change in interest rates over a given instant will be the change given by the function $\alpha(t)$, with random shocks in the amount of σdz . This formula is closely linked to the term structure models proposed by Ho and Lee (1986), which is the simplest of the single factor models within the HJM framework. The Ho and Lee model, generalized by Heath, Jarrow, and Morton (1992), was the first term structure model to fit actual observable yield curve data exactly. The function $\alpha(t)$ can be chosen to fit a yield curve smoothed using the techniques of Chapter 5 exactly. We go into more detail on the Ho and Lee special case of HJM later in this chapter.

The assumption about yield movements above is *not* satisfactory in the sense that it is still a random walk of interest rates, although there is the drift term $\alpha(t)$ built into the random walk. The best way to build the interest rate cycle into the random movement of interest rates is to assume that the interest rate drifts back to

some long-run level, except for the random shocks from the dz term. One process that does this is

$$dy = \alpha(\gamma - y)dt + \sigma dz$$

This is the Ornstein-Uhlenbeck process that is at the heart of the Vasicek (1977) model, which we introduce below. (As we note there, the Vasicek model is another special case of the HJM framework.) The term α is called the “speed of mean reversion.” It is assumed to be a positive number. The larger it is, the faster y drifts back toward its long-run mean γ and the shorter and more violent interest rate cycles will be. Since α is positive, when y is above γ , it will be pulled down except for the impact of the shocks that emanate from the dz term. When y is less than γ , it will be pulled up.

In almost every early term structure model, the speed of mean reversion and the volatility of interest rates play a key role in determining the value of securities, which either are interest options or which have interest rate options embedded in them. Many of the early term structure models, in spite of their simplicity, were powerful enough and analytically attractive enough that they were used in both the random interest rate version of the Merton credit model (Chapter 18) and the more modern reduced form credit models of Duffie and Singleton (1999) and Jarrow (1999, 2001) (which we discuss in Chapter 16).

For that reason, it is important to understand a little more about how these interest rate models are derived.

ITO'S LEMMA

Once a term structure model has been chosen, we need to be able to draw conclusions about how securities whose values depend on interest rates should move around. For instance, if a given interest rate y is assumed to be the random factor that determines the price of a zero-coupon bond P , it is logical to ask how the zero-coupon bond price moves as time passes and y varies. The formula used to do this is called *Ito's lemma*.¹ Ito's lemma puts the movement in the bond price P in stochastic process terms like this:

$$dP = P_y dy + \frac{1}{2} P_{yy} (dy)^2 + P_t$$

where the subscripts denote partial derivatives. (This formula looks very much like the Taylor series expansion in the discussion on convexity in Chapter 12.) The change in the price of a zero-coupon bond equals its first derivative times the change in y plus one-half of its second derivative times the volatility of y plus the “drift” in prices over time. The terms dy and $(dy)^2$ depend on the stochastic process chosen for y . If the stochastic process is the Ornstein-Uhlenbeck process given above, we can substitute in the values for dy and $(dy)^2$. If we do this, then the movements in P can be rewritten

$$\begin{aligned}
dP &= P_y[\alpha(\gamma - y)dt + \sigma dz] + \frac{1}{2}P_{yy}\sigma^2 dt + P_t dt \\
&= \left[P_y\alpha(\gamma - y) + \frac{1}{2}P_{yy}\sigma^2 + P_t \right] dt + P_y\sigma dz \\
&= g(y, t)dt + h(y, t)dz
\end{aligned}$$

For any stochastic process, the dy term is the stochastic process itself. The term $(dy)^2$ is the instantaneous variance of y . In the case of the Ornstein-Uhlenbeck process, it is the square of the coefficient of dz , σ^2 . The term $g(y, t)$, which depends on the level of rates y and time t , is the drift in the bond price. The term $h(y, t)$ is the bond price volatility in the term structure model sense, not the Black-Scholes sense. We have very neatly divided the expression for movements in the bond's price into two pieces. The first term depends on the level of interest rates and time, and the second term depends on the dz term, which introduces random shocks to the interest rate market. If dz or h were always zero, then interest rates would not be random. They would drift over time, but there would be no surprises. Not surprisingly, we will devote a lot of time to discussing the hedging that eliminates these interest rate shocks.

ITO'S LEMMA FOR MORE THAN ONE RANDOM VARIABLE

What if the zero-coupon bond price depended on two random variables, x and y ? Then we would be dealing with a two-factor term structure model, like the HJM two-factor model in Chapter 8. The formula for the movement in the bond's price for such a two-factor term structure model is given by Ito's lemma as follows:

$$dP = P_x dx + P_y dy + \frac{1}{2}P_{xx}(dx)^2 + \frac{1}{2}P_{yy}(dy)^2 + P_{xy}(dxdy) + P_t dt$$

The instantaneous correlation between the two Wiener processes Z_x and Z_y , which provide the random shocks to x and y is reflected in the formula for the random movement of P by the instantaneous correlation coefficient ρ ; $(dxdy)$ is defined as $\rho\sigma_x\sigma_y$. The sigmas are the instantaneous variances of the two variables driving this term structure model, parallel to the variance of our original single-factor term structure model above. Shimko (1992) has an excellent discussion of how to apply this kind of arithmetic to many problems in finance. Our need is restricted to term structure models, so we now turn to the construction of a single-factor term structure model based on the duration concept.

USING ITO'S LEMMA TO BUILD A TERM STRUCTURE MODEL

In their pioneering work, Heath, Jarrow, and Morton (1992) clearly outline the steps that must be followed in developing a term structure model. There are potentially hundreds of alternative term structure models but their derivation has a common

core. The steps involved in building a term structure model almost always follow this process:

1. Make an assumption about the random (stochastic) process that interest rates follow that is consistent with market behavior in general.
2. Use Ito's lemma to specify how zero-coupon bond prices move.
3. Impose the constraint that there not be riskless arbitrage in the bond market via the following steps:
 - a. Use a bond of one maturity (in the case of a one-factor model) to hedge the risk of a bond, which has another maturity.²
 - b. After eliminating the risk of this portfolio, impose the constraint that this hedged portfolio earn the instantaneous (short-term) risk-free rate of interest.
4. Solve the resulting partial differential equation for zero-coupon bond prices.
5. Examine whether this implies reasonable or unreasonable conditions on the market.
6. If the economic implications are reasonable, proceed to value other securities with the model. If they are unreasonable, reject the assumption about how rates move as a reasonable basis for a term structure model.

We illustrate this process by analyzing the economic implications of the assumption of parallel yield curve shifts underlying the traditional duration analysis.

DURATION AS A TERM STRUCTURE MODEL

In this section, we use the continuous compounding formulas in Chapters 4 and 12 to analyze the yield to maturity y_i on a zero-coupon bond with maturity t_i and its value P .

$$P(t_i) = e^{-y_i t_i}$$

We want to examine what happens when all yields at time zero y_0 shift by a parallel amount x , with x being the same for all maturities. After the shift of x , the price of a zero-coupon bond with maturity $\tau = T - t$ (t is the current time in years and T is the maturity date of the bond in years, for a net maturity of $T - t$) will be

$$P(\tau) = e^{-(y_0+x)\tau}$$

where $y = y_0 + x$. We assume that x is initially zero and that we are given today's yield curve and know the values of y for all maturities.

We now make our assumption about the term structure model for interest rates. We assume that the movements in x follow a random walk, that is, we assume changes in x have no drift term and that the change in x is normally distributed as follows:

$$dx = \sigma dz$$

The scale factor σ controls the volatility of the parallel shifts x . We now proceed to step 2 of the previous section by using Ito's lemma to specify how the zero-coupon bond price for a bond with time to maturity $T - t$ moves. We use the formulas for

the first and second derivatives of the bond price P and our knowledge of what dx and $(dx)^2$ are. Ito's lemma states

$$\begin{aligned}
 dP &= P_x dx + \frac{1}{2} P_{xx} (dx)^2 + P_t dt \\
 &= -\tau P \sigma dz + \frac{1}{2} \tau^2 P \sigma^2 dt + y P dt \\
 &= \left[\frac{1}{2} \tau^2 P \sigma^2 + y P \right] dt - \tau P \sigma dz \\
 &= g(y, t) P dt + h(y, t) P dz
 \end{aligned}$$

because

$$\begin{aligned}
 y &= y_0 + x \\
 \tau &= T - t \\
 P &= \exp^{-y\tau} \\
 P_x &= -\tau P \\
 P_{xx} &= \tau^2 P \\
 P_t &= y P \\
 g(y, t) &= \frac{1}{2} \tau^2 \sigma^2 + y \\
 h(y, t) &= -\tau \sigma
 \end{aligned}$$

We now go to step 3 as outlined in the previous section. We want to form a portfolio of 1 unit of bond 1 with maturity τ_1 and w units of the bond 2 with maturity τ_2 . This is exactly the same kind of hedge construction that we did in Chapter 12. The value of this portfolio is

$$W = P_1 + w P_2$$

We then apply Ito's lemma to changes in the value of the portfolio as the parallel shift amount x moves:

$$\begin{aligned}
 dW &= dP_1 + w dP_2 \\
 &= g_1 P_1 dt + h_1 P_1 dz + w [g_2 P_2 dt + h_2 P_2 dz] \\
 &= [g_1 P_1 + w g_2 P_2] dt + [h_1 P_1 + w h_2 P_2] dz
 \end{aligned}$$

As in step 3b listed previously, we want to eliminate the interest rate risk in this portfolio. Since all interest rate risk comes from the random shock term dz , "eliminating the interest rate risk" means choosing w such that the coefficient of dz is zero. This means that the proper hedge ratio w for zero interest rate risk is

$$w = \frac{-h_1 P_1}{h_2 P_2} = \frac{\tau_1 P_1}{\tau_2 P_2}$$

Substituting this into the previous equation means

$$dW = \left[g_1 P_1 + \frac{\tau_1 P_1}{\tau_2 P_2} g_2 P_2 \right] dt$$

Now we impose the no-arbitrage condition. Since the interest rate risk has been eliminated from this portfolio, the instantaneous return on the portfolio dW should equal the riskless³ short-term rate (the rate y with maturity 0, $y(0)$) times the value of the portfolio W :

$$\begin{aligned} dW &= r(P_1 + wP_2)dt \\ &= [g_1 P_1 + w g_2 P_2]dt \end{aligned}$$

Rewritten, this means that

$$\frac{rP_1 - (\frac{1}{2}\tau_1^2\sigma^2 + y_1)P_1}{\tau_1 P_1} = \frac{rP_2 - (\frac{1}{2}\tau_2^2\sigma^2 + y_2)P_2}{\tau_2 P_2} = k$$

This ratio has to be equal for any two maturities or there will be riskless arbitrage opportunities in this bond market. We define this ratio k as the “market price of risk.” Choosing bond 1 and dropping the subscript 1 means that the yield y must satisfy the following relationship:

$$rP - \frac{1}{2}\tau^2\sigma^2P - yP = k\tau P$$

Rearranging means that the yield y for any maturity must adhere to the following relationship:

$$y(\tau) = r - k\tau - \frac{1}{2}\tau^2\sigma^2$$

This equation comes from step 4 listed previously. For a no-arbitrage equilibrium in the bond market under a parallel shift in the yield curve, the yield y must be a quadratic function of time to maturity τ . If the function is not quadratic, there would be riskless arbitrage in the market. This is a surprising and little understood implication of the assumption of parallel yield curve shifts in the market. Clearly, the implication that yields are quadratic is good news and bad news. The good news is that the yield curve math is simple and can fit a lot of observable yield curves over some portion of their range. The bad news is the basic shape of the quadratic function!

CONCLUSIONS ABOUT THE USE OF DURATION'S PARALLEL SHIFT ASSUMPTIONS

Unfortunately, the assumption that yields move in parallel fashion results in a number of conclusions that do not make sense if we impose the no-riskless-arbitrage

conditions: First, interest rates can have only one “hump” in the yield curve. This hump will occur at the maturity

$$\tau = -\frac{k}{\sigma^2}$$

Much more important, beyond a certain maturity, yields will become *negative* since sigma is positive. Yields will be zero for two values of T defined by setting y equal to zero and solving for the maturities consistent with zero yield:

$$T = -\frac{k}{\sigma^2} \pm \sqrt{\frac{2r}{\sigma^2} + \frac{k^2}{\sigma^4}}$$

For normal values of k , these values of T define a range over which y will be positive. Outside of that range, yields to maturity will be negative. Forward rates are also quadratic and have the form

$$f(\tau) = r - 2k\tau - \frac{3}{2}\tau^2\sigma^2$$

because forward rates are related to yields by the formula

$$f(\tau) = y(\tau) + \tau y'(\tau)$$

as shown in Chapter 5. Zero-coupon bond prices are given by

$$P(\tau) = e^{-y(\tau)\tau} = e^{-r\tau + k\tau^2 + \frac{1}{2}\sigma^2\tau^3}$$

These features of a parallel shift-based term structure model pose very serious problems. If the parallel yield curve shift assumption is used, as it is when one uses the traditional duration approach to hedging, the resulting bond market is either:

1. Inconsistent with the real world because of negative yields if we impose the no-arbitrage conditions
2. Allowing of riskless arbitrage if we let yields have a shape that is not quadratic in the term to maturity

In either case, the model is unacceptable for use by market participants because its flaws are so serious. The most important point of this chapter, however, is to show how we can move in a very systematic way from an assumption about how interest rates move to explicit formulas for forward rates, zero-coupon bond prices, zero-coupon bond yields, and many other complex securities. We have laid the groundwork for extending the models to be consistent with observable yield curves, which we demonstrated with HJM 1-, 2-, and 3-factor models in Chapters 6 through 9. Most importantly, we have laid the groundwork for valuing the Jarrow-Merton put option that quantifies interest rate risk, market risk, liquidity risk, and credit risk in a fully integrated way.

We now use the same sort of analysis to identify some more attractive alternative models of the term structure that were developed prior to the publication of the Heath, Jarrow, and Morton approach (1990a, 1990b, 1992a).

THE VASICEK AND EXTENDED VASICEK MODELS

In this section, we introduce four term structure models that are based on increasingly realistic assumptions about the random movement of interest rates. All of the models are special cases of the HJM framework. As we saw in Chapters 8 and 9, the legacy assumption that the yield curve shifts in a parallel fashion is inconsistent with actual movements of interest rates over 50 years of daily data in the United States. Parallel yield curve shifts can be consistent with a no-arbitrage economy, but only if the yield curve is quadratic and only if yields become negative beyond a given point on the yield curve. For this reason, academics sought a richer set of assumptions about yield curve movements that would allow us to derive a theoretical yield curve whose shape is as close as possible, if not identical, to observable market yields and whose other properties are realistic.

This chapter is too mathematical for many practitioners and too simple for many hard-core “rocket scientists.” We ask both to hang in there as our objective in this volume is to show clearly the links between good theory and best practice in risk management. Along the way, we point the readers to complementary sources that are more practical or more theoretical depending on the point of view of the reader.

We start with the simplest possible model. We use the same assumptions for each of the four models in this chapter. We assume bond prices in all four models depend solely on a single random factor, the short rate of interest. This is a simpler framework than we employed using the two one-factor HJM models in Chapters 6 and 7. We keep it simple for expository purposes. By the end of the chapter, the reader will be grateful that the insights of HJM are so powerful that the kinds of derivations we now turn to are no longer necessary.

As we go through each of these four single-factor term structure models, we go through the same steps as in the prior section to derive the form that zero-coupon bond prices and yields must have for the risk-free bond market to be a no-arbitrage market. The HJM framework, by contrast, derives the no-arbitrage framework for n factors, with one factor as the simplest special case. We go through the normal academic term structure model process:

- We specify the stochastic process for the movements in interest rates.
- We use Ito’s lemma to derive the stochastic movements in zero-coupon bond prices.
- We impose no-arbitrage restrictions to generate the partial differential equation that determines zero-coupon bond prices.
- We solve this equation for bond prices, subject to the boundary condition that the price of a zero-coupon bond with zero years to maturity must be 1.

The four models we discuss assume that changes in the short-term rate of interest are normally distributed. This assumption is a little more questionable in low-rate environments like Japan since the mid-1990s and the United States in the wake of the

2006–2011 credit crisis. The HJM framework does not require this specification, but we use it in the remainder of this chapter for expository reasons. The four models we review here can be summarized like this:

1. The *parallel rate shock model*, where changes in the short-term risk-free rate r are a random walk with no drift over time. This is the same as the continuous time duration model reviewed in the prior section. It was originally introduced by Merton. Because the model has zero drift in rates over time, this model generally will not exactly match the observable yield curve.
2. The *extended parallel rate shock model*, which matches the observable yield curve exactly. Changes in the short-term risk-free rate of interest r are a random walk with non-zero drift, allowing a perfect fit to an observable yield curve. We label this model the “extended” Merton model or Ho and Lee model.
3. A non-parallel rate shift model, in which changes in the risk-free short-term rate of interest r allow for interest rate cycles. Analytically, we would say that changes in r follow the “Ornstein-Uhlenbeck” process with a drift term, and r is reverting to a constant long-term value plus normally distributed shocks, that is, the Vasicek model. It is this “mean reversion” that causes cycles in interest rates.
4. A nonparallel rate shift model, which fits the observable yield curve exactly. In this model changes in r follow the Ornstein-Uhlenbeck process again. The drift in r is such that we match the observable yield curve. This model is the “extended” Vasicek (or Hull and White) model.

Each of these models has their strengths and weaknesses. All are special cases of the HJM approach and less realistic than any of the HJM implementations in Chapters 6 through 9. We discuss each in turn. We use these models to fit the risk-free yield curves in the countries we care about. We build the credit models of Chapters 16 and 17 on top of them to get the yield curves for risky issuers. The term structure models in this chapter should *not* be used for yield curves with credit risk.

THE MERTON TERM STRUCTURE MODEL: PARALLEL YIELD CURVE SHIFTS

One simple assumption about the short-term risk-free interest rate r is that it follows a simple random walk with a zero drift over time.⁴ In the earlier part of this chapter, we applied this assumption to zero-coupon bond yields of all maturities and required that this random shift in rates be the same for all maturities. Merton (1970) applied the random walk assumption to the short-term rate of interest r and derived what this means for the rest of the yield curve. We can write the change in r as

$$dr = \sigma dZ$$

The change in the short rate of interest r equals a constant sigma times a random shock term, where Z represents a standard Wiener process with mean zero and standard deviation of 1. The constant sigma is the instantaneous volatility of interest rates, which we can scale up and down to maximize the goodness of fit versus the observable yield curve. We will be careful to distinguish between interest rate

volatility and the volatility of zero-coupon bond prices throughout the rest of this book. From the previous section, we know that Ito's lemma can be used to write the stochastic process for the random changes in the price of a zero-coupon bond price as of time t with maturity T :

$$dP = P_r dr + \frac{1}{2} P_{rr} (dr)^2 + P_t$$

The change in the bond's price equals its first derivative with respect to r times the change in r , plus a term multiplied by the second derivative of the bond's price and a drift term, the derivative of the bond's price with respect to the current time t . In this random walk model, the variance of the interest r is $(dr)^2 = \sigma^2$, so the expression above can be expanded to read

$$dP = P_r \sigma dZ + \frac{1}{2} P_{rr} \sigma^2 + P_t$$

Our goal is to find the formula for P as a function of the short rate r . Once we do this, we can vary the time to maturity T and derive the shape of the entire yield curve and how the yield curve moves in response to changes in r .

The next step in doing this is to impose the condition that no arbitrage is possible in the bond market. Since we have only one random factor in the economy, we know that we can eliminate this risk factor by hedging with only one instrument (as pointed out in Chapter 12). If we had n independent risk factors, we would need n instruments to eliminate the n risks. Let's assume that an investor holds one unit of a zero-coupon bond, bond 1, with maturity T_1 . The investor forms a portfolio of amount W that consists of the one unit of bond 1 and w units of bond 2, which has a maturity T_2 . The bond with maturity T_2 is the hedging instrument. The value of the portfolio is

$$W = P_1 + wP_2$$

What is the proper hedge ratio w , and what are the implications of a perfect hedge for bond pricing? We know from Ito's lemma that the change in the value of this hedged portfolio is

$$\begin{aligned} dW &= dP_1 + w dP_2 \\ &= P_{1,r} \sigma dZ + \left(\frac{1}{2} P_{1,rr} \sigma^2 + P_{1,t} \right) dt + w \left[P_{2,r} \sigma dZ + \left(\frac{1}{2} P_{2,rr} \sigma^2 + P_{2,t} \right) dt \right] \end{aligned}$$

Gathering the coefficients of the random shock term dZ together gives

$$dW = \left(\frac{1}{2} P_{1,rr} \sigma^2 + P_{1,t} \right) dt + w \left(\frac{1}{2} P_{2,rr} \sigma^2 + P_{2,t} \right) dt + [P_{1,r} + wP_{2,r}] \sigma dZ$$

We repeat the process we just followed to arrive at the zero risk hedge. If we choose w , the hedge amount of bond 2, such that the coefficient of the random shock term dz is zero, then we have a perfect hedge and a riskless portfolio. The value of w for which this is true is

$$w = \frac{-P_{1,r}}{P_{2,r}}$$

Note that this is identical to the hedge ratio we found for the duration-based parallel shift model. If we use this hedge ratio, then the instantaneous return on the portfolio should exactly equal the short-term rate of interest times the value of the portfolio:

$$dW = rWdt = (rP_1 + rwP_2)dt$$

If this is not true, then riskless arbitrage will be possible. Imposing this condition and rearranging the equation above gives us the following relationship:

$$dW = (rP_1 + rwP_2)dt = \left(\frac{1}{2}P_{1,r}\sigma^2 + P_{1,t} \right) dt + w \left(\frac{1}{2}P_{2,r}\sigma^2 + P_{2,t} \right) dt$$

We then substitute the expression for w above into this equation, eliminate the dt coefficient from both sides, rearrange, and divide by interest rate volatility sigma to get the following relationship.

$$-\lambda = \frac{\frac{1}{2}P_{1,r}\sigma^2 + P_{1,t} - rP_1}{\sigma P_{1,r}} = \frac{\frac{1}{2}P_{2,r}\sigma^2 + P_{2,t} - rP_2}{\sigma P_{2,r}}$$

For any two maturities T_1 and T_2 , the no-arbitrage condition requires that this ratio be equal. We call the negative of this ratio λ ,⁵ the market price of risk. For normal risk aversion, the market price of risk should be positive because in a risk-averse market, riskier (longer maturity) bonds should have a higher expected return than the short rate r and because the rate sensitivity of all bonds P_r is negative. Since the choice of T_1 and T_2 is arbitrary, the market price of risk ratio must be constant for all maturities. *It is the fixed income equivalent of the Sharpe ratio* that measures excess return per unit of risk with the following statistic:

$$\text{Sharpe ratio} = \frac{\text{Expected return} - \text{Risk-free return}}{\text{Standard deviation of return}}$$

We will return to the Sharpe ratio and the strong parallels between equilibrium in the fixed income and the equity markets frequently in this book. In our case, the numerator is made up of the drift in the bond's price:

$$\text{Drift} = \text{Expected return} = \frac{1}{2}P_{rr}\sigma^2 + P_t$$

which also equals the expected return since the shock term or stochastic component of the bond's price change on average has zero expected value.

We now solve the equation above for any given bond with maturity T and remaining time to maturity $T - t = \tau$ as of time t , under the assumption that λ (the market price of risk) is constant over time. Rearranging the equation above gives the following partial differential equation

$$\lambda\sigma P_r + \frac{1}{2}P_{rr}\sigma^2 + P_t - rP = 0$$

which must be solved subject to the boundary condition that the price of the zero-coupon bond upon maturity T must equal its principal amount 1:

$$P(r, T, T) = 1$$

We use a common method of solving partial differential equations, the educated guess. We guess that P has the solution

$$P(r, t, T) = P(r, \tau) = e^{rF(\tau) + G(\tau)}$$

where $T - t = \tau$ and we need to solve for the forms of the functions F and G . We know that

$$\begin{aligned} P_r &= F(\tau)P \\ P_{rr} &= F(\tau)^2 P \\ P_t &= -P_\tau = (-rF' - G')P \end{aligned}$$

Substituting these values into the partial differential equation and dividing by the bond price P gives the following:

$$\lambda\sigma F + \frac{1}{2}\sigma^2 F^2 - rF' - G' - r = 0$$

We know from Merton (1970) and the similar equation above that the solution to this equation is

$$\begin{aligned} F(\tau) &= -\tau \\ G(\tau) &= -\frac{\lambda\sigma\tau^2}{2} + \frac{1}{6}\sigma^2\tau^3 \end{aligned}$$

and therefore the formula for the price of a zero-coupon bond is given by

$$P(r, t, T) = P(r, \tau) = e^{-r\tau - \frac{\lambda\sigma^2}{2}\tau^2 + \frac{1}{6}\sigma^2\tau^3}$$

This is nearly the same bond-pricing equation that we obtained under the assumption of parallel shifts in bond prices. It has the same virtues and the same liabilities of the duration approach:

- It is a simple analytical formula.
- Zero-coupon bond yields are a quadratic function of time to maturity.
- Yields turn negative (and zero-coupon bond prices rise above one) beyond a certain point.
- If interest rate volatility σ is zero, zero-coupon bond yields are constant for all maturities and equal to r .

Using continuous compounding,

$$P(r, \tau) = e^{-y(\tau)\tau}$$

and the zero-coupon bond yield formula can be calculated as

$$y = r + \frac{1}{2}\lambda\sigma\tau - \frac{1}{6}\sigma^2\tau^2$$

The last term in this formula ultimately becomes so large as time to maturity increases that yields become negative. Setting $y' = 0$ and solving for the maximum level of y (the “hump” in the yield curve) shows that this occurs at a time to maturity of

$$\tau^* = \frac{3\lambda}{2\sigma}$$

and that the yield to maturity at that point is

$$y(\tau^*) = r + \frac{3}{8}\lambda^2$$

Note that the peak in the yield curve is independent of interest rate volatility, although the location of this hump is affected by rate volatility. As volatility increases, the hump moves to shorter and shorter maturity points on the yield curve. We can also determine the point at which the zero-coupon bond yield equals zero using the quadratic formula. This occurs when the yield curve reaches a time to maturity of

$$\tau^* = \frac{3\lambda}{2\sigma} + \frac{3}{\sigma} \sqrt{\frac{\lambda^2}{4} + \frac{2r}{3}}$$

The market price of risk is normally expected to be positive so this formula shows rates will turn negative at a positive term to maturity τ^* .

The price volatility of zero-coupon bonds is given by

$$\sigma P_r = -\tau \sigma P$$

Another useful tool that we can derive from this term structure model is the percentage change in the price of zero-coupon bonds for small changes in the short rate r . Under the Merton term structure model this sensitivity to r is

$$r\text{-Duration} = -\frac{P_r}{P} = -\frac{-\tau P}{P} = \tau$$

so the rate sensitivity of a zero-coupon bond in this model is its time to maturity in years. We label this percentage change “ r -duration” to contrast it with Macaulay’s measure of price sensitivity discussed in earlier chapters, where price is differentiated with respect to its continuous yield to maturity. We can calculate the r -duration of coupon bearing bonds using the same logic. The hedge ratio w given above for hedging of a bond with maturity T_1 with a bond maturing at T_2 is

$$w = \frac{-P_{1r}}{P_{2r}} = -\frac{\tau_1 P_1}{\tau_2 P_2}$$

The Merton term structure model gives us very important insights into the process of deriving a term structure model. Its simple formulas make it a useful expository tool, but the negative yields that result from the formula are a major concern. It leads to a logical question: Can we extend the Merton model to fit the actual yield curve perfectly? If so, this superficially at least would allow us to avoid the negative yield problems associated with the model. We turn to this task next.

THE EXTENDED MERTON MODEL

Ho and Lee (1986) extended the Merton model to fit a given initial yield curve perfectly in a discrete time framework. The Ho and Lee/Extended Merton model is a special case of HJM framework, where interest volatility at any current time t for future time T in state s_t $\sigma(t, T; s_t)$ is a deterministic constant independent of state s_t : $\sigma(t, T; s_t) = \sigma$. As Jarrow (2002, 288) notes, this implies that forward rates and spot rates can be negative. In this section, we derive the Ho and Lee/Extended Merton model using the same approach above. This basic extension technique is common to all term structure models developed prior to publication of the HJM model. In the extended Merton case, we assume that the short rate of interest is again the single stochastic factor driving movements in the yield curve. Instead of assuming that the short rate r is a random walk, however, we assume that it has a time-dependent drift term $a(t)$:

$$dr = a(t)dt + \sigma dZ$$

As before, Z represents a standard Wiener process that provides random shocks to r ; Z has mean zero and standard deviation of 1. The instantaneous standard deviation of interest rates is again the constant sigma, which we can scale up and down to fit observable market yields. Appendix A of this chapter shows how to derive zero-coupon bond prices in this model using exactly the same procedures as the previous section. The only difference is that we allow the change in the short rate of interest to follow the drift term $a(t)$. When we apply the no-arbitrage conditions and derive zero-coupon bond prices, we get the following formula:

$$P(r, t, T) = e^{-r\tau - \frac{\lambda\sigma^2}{2}\tau^2 + \frac{1}{6}\sigma^2\tau^3 - \int_t^T a(s)(T-s)ds}$$

Note that this is identical to the formula for the zero-coupon bond price in the Merton term structure model with the exception of the last term. The value of the zero-coupon bond is 1 when $t = T$ as the boundary condition demands.

Like the Merton term structure model, the price volatility of zero-coupon bonds is given by⁶

$$\sigma P_r = -\tau\sigma P$$

and the percentage change in price of zero-coupon bonds for small changes in the short rate r is

$$r\text{-Duration} = -\frac{P_r}{P} = -\frac{-\tau P}{P} = \tau$$

The behavior of the yield to maturity on a zero-coupon bond in the extended Merton-Ho and Lee model is always consistent with the observable yield curve since the expression for yield to maturity

$$y(\tau) = r + \frac{\lambda\sigma\tau}{2} - \frac{1}{6}\sigma^2\tau^2 + \frac{1}{\tau} \int_t^T a(s)(T-s)ds$$

contains the extension term

$$\text{Extension}(\tau) = \frac{1}{\tau} \int_t^T a(s)(T-s)ds$$

The function $a(s)$ is chosen such that the theoretical zero-coupon yield to maturity y and the actual zero-coupon yield are exactly the same. The function $a(s)$ is the “plug” that makes the model fit and corrects for model error that would otherwise cause the model to give implausible results. Since any functional form for yields can be adapted to fit a yield curve precisely, it is critical in examining any model for plausibility to minimize the impact of this extension term.⁷ Why? Because the extension term itself contains no economic content.

In the case of the Ho and Lee model, the underlying model would otherwise cause interest rates to sink to negative infinity, just as in the Merton model. The extension term's magnitude, therefore, must offset the negative interest zero-coupon bond yields that would otherwise be predicted by the model. As maturities get infinitely long, the magnitude of the extension term will become infinite in size. This is a significant cause for concern, even in the extended form of the model. One of the most important attributes of term structure model selection is to find one that best captures the underlying economics of the yield curve so that this extension is minimized for the entire maturity spectrum.

THE VASICEK MODEL

Both the Merton model and its extended counterpart, the Ho and Lee model, are based on an assumption about random interest rate movements that imply that, for any positive interest rate volatility, zero-coupon bond yields will be negative at every single instant in time for long maturities beyond a critical maturity τ^* . The extended version of the Merton model, the Ho and Lee model, offsets the negative yields with an “extension” factor that must grow larger and larger as maturities lengthen. Vasicek (1977) proposed a model that avoids the certainty of negative yields and eliminates the need for a potentially infinitely large extension factor. Even more important, the Vasicek model produces realistic interest rate cycles. Vasicek accomplishes this by assuming that the short rate r has a constant interest rate volatility σ like the models above, with an important twist: the short rate exhibits “mean reversion” around the long-run average level of interest rates. The Vasicek model and its extended form are special cases of HJM with a single risk factor and a time (but not state) dependent volatility function which takes the form $\sigma(t, T; s_t) = ke^{-a(T-t)}$ where k and a are constants. Vasicek calls the long-run average level of interest rates μ :

$$dr = \alpha(\mu - r)dt + \sigma dZ$$

The change in the short-term rate of interest r is a function of a drift term and a shock term where

r = the instantaneous short rate of interest

α = the speed of mean reversion

μ = the long-run expected value for r

σ = the instantaneous standard deviation of r

Z is the standard Wiener process with mean zero and standard deviation of 1, exactly like we have used in the first two examples. The difference from previous models comes in the drift term, which gives this stochastic process its name: the Ornstein-Uhlenbeck process. The drift term in the stochastic process proposed by Vasicek pulls the short rate r back toward μ , so μ can be thought of as the long-run level of the short rate. When the short rate r is above μ , the first term tends to pull r downward since α is assumed to be positive. When the short rate r is below μ , r tends to drift upward. The second term of the stochastic process, of course, applies random

shocks to the short rate, which may temporarily offset the tendencies toward mean reversion of the underlying stochastic process. The impact of mean reversion is to create realistic interest rate cycles, with the level of α determining the length and violence of rises and falls in interest rates. What are the implications of this model for the pricing of bonds? We derive the zero-coupon bond-pricing formula in Appendix B of this chapter. The zero-coupon bond-pricing formula is an exponential function with two terms, one of which is a function of the short-term rate of interest r :

$$P(r, t, T) = e^{-F(t, T)r - G(t, T)}$$

$$= \exp \left[-rF(\tau) - \left(\mu + \frac{\lambda\sigma}{\alpha} - \frac{\sigma^2}{2\alpha^2} \right) [\tau - F(\tau)] - \frac{\sigma^2 F^2(\tau)}{4\alpha} \right]$$

where the function F is defined in Appendix B as

$$F(t, T) = F(\tau) = \frac{1}{\alpha} (1 - e^{-\alpha\tau})$$

The continuously compounded zero-coupon bond yield in the Vasicek model uses the definition of function G in Appendix B of this chapter:

$$y(\tau) = -\frac{1}{\tau} \ln[P(\tau)] = -\frac{1}{\tau} [-rF(\tau) - G(\tau)]$$

$$= \frac{F(\tau)}{\tau} r + \frac{G(\tau)}{\tau}$$

As time to maturity gets infinitely long, one can calculate the infinitely long zero-coupon bond yield maturity as

$$y(\infty) = \mu + \frac{\lambda\sigma}{\alpha} - \frac{\sigma^2}{2\alpha^2}$$

This yield to maturity is positive for almost all realistic sets of parameter values, correcting one of the major objections to the Merton term structure model, without the necessity to “extend” the yield curve with a time-dependent drift in the short rate of interest r .

What does duration mean in the context of the Vasicek model? Macaulay defined duration, as we saw in Chapter 12, as the percentage change in bond prices with respect to the parallel shift in yield to maturity in a continuous time context. The parallel shift in the Macaulay model was the single stochastic factor driving the yield curve. In the Vasicek model, the short rate r is the stochastic factor. We define “ r -duration” as above: the percentage change in the price of a bond with respect to changes in the short rate r :

$$r\text{-Duration} = \frac{-P_r}{P} = \frac{F(\tau)P}{P} = F(\tau) = \frac{1}{\alpha} [1 - e^{-\alpha\tau}]$$

This function F is a powerful tool one can use for hedging in the Vasicek model and its extended version, which we cover in the next section. The hedge ratio necessary to hedge one unit of a zero-coupon bond with a remaining maturity of τ_1 using a bond with remaining maturity of τ_2 is

$$w = -\frac{P_{1r}}{P_{2r}} = -\frac{F(\tau_1)P[\tau_1]}{F(\tau_2)P[\tau_2]}$$

a hedge ratio substantially different from that using the Merton or Ho and Lee models:

$$w = -\frac{\tau_1 P[\tau_1]}{\tau_2 P[\tau_2]}$$

In practical use, it is this difference in hedge ratios that allows us to distinguish between different models. The ability to extend a model, as in the previous section, renders all extendible models equally good in the sense of fitting observable data. For example, the HJM framework allows the analyst to fit an observable term structure of zero-coupon bond yields with (almost) any volatility structure. In reality, however, the explanatory power of various term structure models and volatility assumptions can be substantially different. Ultimately, the relative performance of each model's hedge ratios and its ability to explain price movements of traded securities with the fewest parameters are what differentiates the best models from the others.

The stochastic process proposed by Vasicek allows us to calculate the expected value and variance of the short rate at any time in the future s from the perspective of current time t . Denoting the short rate at time t by $r(t)$, the expected value of the short rate at future time s is

$$E_t[r(s)] = \mu + [r(t) - \mu]e^{-\alpha(s-t)}$$

The standard deviation of the potential values of r around this mean value is

$$\text{Standard deviation } [r(s)] = \sqrt{\frac{\sigma^2}{2\alpha} [1 - e^{-2\alpha(s-t)}]}$$

Because $r(s)$ is normally distributed, there is a positive probability that $r(s)$ can be negative. As pointed out by Black (1995), this is inconsistent with a no-arbitrage economy in the special sense that consumers hold an option, which Black assumed is costless, to hold cash instead of investing at negative interest rates. The magnitude of this theoretical problem with the Vasicek model depends on the level of interest rates and the parameters chosen.⁸ The biggest problem with the Vasicek model is not the positive probability of negative rates. Like all of the term structure models described in this chapter, the biggest problem is the assumption that only one random factor drives the yield curve, which is inconsistent with the evidence of Dickler, Jarrow, and van Deventer (2011) that 6 to 10 random factors drive the movements of the risk-free term structure.

THE EXTENDED VASICEK–HULL AND WHITE MODEL

Hull and White (1990b) bridged the gap between the observable yield curve and the theoretical yield curve implied by the one-factor Vasicek model by “extending” or stretching the theoretical yield curve to fit the actual market data. The Hull and White extension was published contemporaneously with the Heath, Jarrow, and Morton (1990a, 1990b, 1992) extension for n -factor models. Hull and White apply the identical logic as the previous section, but they allow the market price of risk term λ to drift over time, instead of assuming it is constant as in the Vasicek model. If we assume this term drifts over time, what are the implications for the pricing formula for zero-coupon bonds?

Hull and White use the time-dependent drift term in interest rates θ , where θ in turn depends on a time-dependent market price of risk:

$$\theta(t) = \alpha\mu + \lambda(t)\sigma$$

Hull and White derive the zero-coupon bond price in the extended Vasicek model using this assumption as shown in Appendix C of this chapter:

$$\begin{aligned} P(r, t, T) &= e^{-F(t, T)r - G(t, T)} \\ &= \exp \left[-rF(\tau) - \int_t^T F(s, T)\theta(s)ds + \left(\frac{\sigma^2}{2\alpha^2} \right) [\tau - F(\tau)] - \frac{\sigma^2 F^2(\tau)}{4\alpha} \right] \end{aligned}$$

The zero-coupon bond price is again an exponential function with two terms, one that is multiplied by the random short-term rate of interest r and the other that contains the functions F and G . Function F is defined the same way as in the previous section. Function G contains the extension term and is defined in Appendix C of this chapter. As in the Vasicek model, the price sensitivity of a zero-coupon bond is given by the formula

$$r\text{-Duration} = \frac{-P_r}{P} = \frac{F(\tau)P}{P} = F(\tau) = \frac{1}{\alpha} [1 - e^{-\alpha\tau}]$$

This is the same formula as in the regular Vasicek model. We now turn to a brief historical review of other term structure models proposed in the academic literature leading up to the publication of the HJM framework.

ALTERNATIVE TERM STRUCTURE MODELS

A complete review of the vast array of term structure models that have been proposed in the academic literature is beyond the scope of this chapter.⁹ So, we will briefly review a small subset of the term structure models of the past five decades, with occasional asides to remind the reader that each of the models mentioned is a special case of Heath, Jarrow, and Morton. None of the models reviewed by the authors cited in the notes has more than the three random factors that we illustrated using the

HJM framework in Chapter 9. Consistent with Dickler, Jarrow, and van Deventer (2011), we remind the reader that 6 to 10 factors are needed for both realism and regulatory purposes, per the Bank for International Settlements regulations cited in prior chapters.

In the remainder of this chapter, we consider alternative term structure models that offer a richer array of potential movements in the term structure of interest rates, while preserving the implicit assumption that default will not occur. In Chapters 16 through 17, we broaden our analysis to include multifactor models of the yield curve that explicitly incorporate the probability of default. For example, the Shimko, Tejima, and van Deventer extension of the Merton model, which we discuss in Chapter 16, is a two-factor term structure model in the guise of a discussion of the valuation of risky debt. The two risky factors are the short rate of interest r , which drives the risk-free yield curve's movements, and the value of the underlying assets being financed.

Alternative One-Factor Interest Rate Models

The Cox, Ingersoll, and Ross (CIR) model, which the authors derive in a general equilibrium framework, has been as popular in the academic community to the same degree as the Vasicek model has been among financial market participants. It is therefore appropriate that we begin our discussion of one-factor model alternatives to the Vasicek model with a review of the CIR approach.

The CIR Model The CIR model has been particularly influential because the original version of the paper was in circulation in the academic community at least since 1977, even though the paper was not formally published until 1985. CIR Cox, Ingersoll, and Ross (1985) assume that the short-term interest rate is the single stochastic factor driving interest rate movements, and that the variance of the short rate of interest is proportional to the level of interest rates. This is a special case of HJM, which Jarrow (2002, 288) labels “nearly proportional volatility such that the interest rate volatility function $\sigma(t, T; s_t) = k(t, T) \min[\log f(t, Y), M]$, subject to an upper bound of a constant M .” Below this upper bound, volatility is proportional to the continuously compounded forward rate so that forward rates are always nonnegative. This implies that interest rate volatility is higher in periods of high interest rates than it is in periods of low interest rates. The realism of this assumption depends on which financial market is being studied, but casual empiricism would lead one to believe that it is a more desirable property when modeling the United States (1978–1985), Brazil, or Mexico than it would perhaps be for Japanese financial markets today. This observation is consistent with the rising levels of interest rate volatility that we found in the U.S. Treasury market from 1962 to 2011 in Chapters 6 through 9, but the CIR assumption is much simpler than the pattern of volatility levels reflected in the actual data.

The authors assume that stochastic movements in the short rate take the form

$$dr = k(\mu - r)dt + \sigma\sqrt{r}dz$$

The change in interest rates will have a mean reverting drift term, which causes interest rate cycles, just like the Vasicek and extended Vasicek models. The interest

rate shock term has the same volatility as the Vasicek model but it is multiplied by the square root of the short-term rate of interest. CIR show that the value of a zero-coupon bond with maturity $\tau = T - t$ takes the form

$$P(r, \tau) = A(\tau)e^{-B(\tau)r}$$

We refer the reader to the original CIR paper for the definitions of A and B . Its authors also derive an analytical solution for the price of a European call option on a zero-coupon bond.

Dothan (1978) provides a model of short rate movements where the short-term riskless rate of interest r follows a geometric Wiener process. The short rate has a lognormal distribution and will therefore always be positive:

$$dr = \sigma r dz$$

The resulting analytical solution for zero-coupon bond prices is quite complex. The Dothan model, while sharing one of the attractive properties of the CIR model, lacks the mean reversion term that causes interest rate cycles, one of the key features necessary in a realistic term structure model.

Longstaff (1989) proposes a model in which the variance of the short rate is proportional to the level of the short rate, like the CIR model, and the mean version of the short rate is a function of its square root:

$$dr = k(\mu - \sqrt{r})dt + \sigma\sqrt{r}dz$$

The resulting pricing equation for zero-coupon bonds is very unique in comparison to other term structure models in that the implied yield to maturity on zero-coupon bonds is a nonlinear function of the short rate of interest:

$$P(r, \tau) = A(\tau)e^{B(\tau)r + C(\tau)\sqrt{r}}$$

A , B , and C are complex functions of the term structure model parameters and are described in Longstaff.

Black, Derman, and Toy (1990) suggest another model that avoids the problem of negative interest rates and allows for time-dependent parameters. The stochastic process specifies the percentage change in the short-term rate of interest is a mean reverting function of the natural logarithm of interest rates:

$$d[\ln(r)] = [\theta(t) - \phi(t)\ln(r)]dt + \sigma(t)dz$$

The model has many virtues from the perspective of financial market participants. It combines the ability to fit the observable yield curve (like the Ho and Lee and extended Vasicek models) with the nonnegative restriction on interest rates and the ability to model the volatility curve observable in the market. The model's liability is the lack of tractable analytical solutions, which are very useful in (1) confirming the accuracy of numerical techniques and (2) valuing large portfolios where speed is essential.

Black and Karasinski (1991) further refine the Black, Derman, and Toy approach with the explicit incorporation of time-dependent mean reversion:

$$d[\ln(r)] = \phi(t)[\ln[\mu(t)] - \ln(r)]dt + \sigma(t)dz$$

This modification allows the model to fit observable cap prices, one of the richest sources of observable market data incorporating interest rate volatility information if one ignores the possibility of manipulation of the short-term LIBOR (which we discuss in Chapter 17). The authors describe in detail how to model bond prices and interest derivatives using a lattice approach. The model, like its predecessor the Black, Derman, and Toy model, has been quite popular among financial market participants.

Two-Factor Interest Rate Models

As we discuss in this section, one of the challenges in specifying a two-factor model is selecting which two factors are the most appropriate. In Chapters 16 and 17, we discuss credit models where the two factors of the risky debt term structure model are the riskless short rate of interest and either (1) the value of the asset being financed (in the Shimko, Tejima, and van Deventer extension of the Merton credit model); or (2) a macroeconomic factor driving default, as in the Jarrow credit model. In this section, we review a number of two-factor models that are based on various assumptions about the two risky factors on the assumption that we are modeling a risk-free yield curve.

The Brennan and Schwartz Model Brennan and Schwartz (1979) introduced a two-factor model where both a long-term rate and a short-term rate follow a joint Gauss-Markov process. The long-term rate is defined as the yield on a consol (perpetual) bond. Brennan and Schwartz assume that the log of the short rate has the following stochastic process:

$$d[\ln(r)] = \alpha[\ln(l) - \ln(p) - \ln(r)]dt + \sigma_1 dz_1$$

The short rate r moves in response to the level of the consol rate l and a parameter p relating the target value of $\ln(r)$ relative to the level of $\ln(l)$. Brennan and Schwartz show that the stochastic process for the consol rate can be written

$$dl = l[l - r + \sigma_2^2 + \lambda_2 \sigma_2]dt + l\sigma_2 dz_2$$

Lambda in this expression is the market price of long-term interest rate risk. Longstaff and Schwartz proceed to test the model on Canadian government bond data, with good results.

Chen and Scott (1992) derive a two-factor model in which the nominal rate of interest I is the sum of two independent variables y_1 and y_2 , both of which follow the stochastic process specified by CIR:

$$dy_i = k_i(\theta_i - y_i)dt + \sigma_i\sqrt{y_i}dz_i$$

Chen and Scott show that the price of a discount bond in this model is

$$P(y_1, y_2, t, T) = A_1 A_2 e^{-B_1 y_1 - B_2 y_2}$$

where A and B have the same definition as in the CIR model, with the addition of the appropriate subscripts. The authors go on to value a wide range of interest rate derivatives using this model. The end result is a powerful model with highly desirable properties and a wealth of analytical solutions.

Hull and White (1994b) show that there is a similar extension for the Vasicek model when the nominal interest rate I is the sum of two factors r_1 and r_2 . The value of a zero-coupon bond with maturity τ is simply the product of two factors P_1 and P_2 , which have exactly the same functional form as the single factor Vasicek model, except that one is driven by r_1 and the other by r_2 :

$$V(r_1, r_2, \tau) = P_1(r_1, \tau) P_2(r_2, \tau)$$

Both stochastic factors are assumed to follow stochastic processes identical to the normal Vasicek model:

$$dr_i = \alpha_i(\mu_i - r_i)dt + \sigma_i dz_i$$

Longstaff and Schwartz (1992a) propose a model in which two stochastic factors, which are assumed to be uncorrelated, drive interest rate movements. The factors x and y are assumed to follow the stochastic processes

$$\begin{aligned} dx &= (\gamma - \delta x)dt + \sqrt{x} dz_1 \\ dy &= (\eta - \nu y)dt + \sqrt{y} dz_2 \end{aligned}$$

Longstaff and Schwartz derive the value of a discount bond in this economy to be

$$V(x, y, \tau) = E_1(\tau) e^{E_2(\tau)x + E_3(\tau)y}$$

Chen's Three-Factor Term Structure Model

As factors are added to a term structure model, it becomes more realistic and more complex. Chen (1994) introduces a three-factor term structure model where the short-term rate of interest is random and mean reverting around a level that is also random. Moreover, Chen also assumes that the volatility of the short rate is random:

$$\begin{aligned} dr &= k(\theta - r)dt + \sqrt{\sigma} \sqrt{r} dz_1 \\ d\theta &= \nu(\bar{\theta} - \theta)dt + \zeta \sqrt{\theta} dz_2 \\ d\sigma &= \mu(\bar{\sigma} - \sigma)dt + \eta \sqrt{\sigma} dz_3 \end{aligned}$$

The volatility of the short rate is written by Chen as $\sqrt{\sigma}$. The long-term level of interest rates θ is itself random and mean-reverting around a long-term level. The

volatility of the short rate is also mean-reverting with a form much like the CIR specification. Chen goes on to derive, with considerable effort, closed-form solutions for this richly specified model.

Chen's model demonstrates that there is a rich array of choices for financial market participants who require the additional explanatory power of a two- or three-factor model and who are willing to incur the costs, which we discuss below, of such models.

REPRISING THE HJM APPROACH

Realism in term structure modeling requires 6 to 10 random factors driving the risk-free term structure of interest rates. The preceding review of 1, 2, and 3 factor models has probably left the reader with the sense that models can quickly become intractable as the number of factors rises without some unifying approach. That unifying approach is the HJM framework. The reason for presenting the HJM approach first, in Chapters 6 through 9, is to show that it is much easier to use in practice than many of the "simpler" term structure models discussed previously that have fewer risk factors and, therefore, less explanatory power. When we say "HJM" in this context, we include its popular variant, the LIBOR market model, which has been widely employed in valuing securities that are a function of the LIBOR-swap curve. (We discuss the LIBOR market model in Chapter 17, where we focus on the modeling of term structures where the underlying yield curve embeds the possibility of default.) For risk-free term structures, the HJM framework is so rich that the expansion of risk factors from the three factors of Chapter 9 to 4, 5, 5, 10, or 20 factors is routine. That is what a sophisticated risk-management software solution is for—to remove the "routine" from risk management so a talented risk manager can spend his or her time on the art and judgment element of the discipline.

In the next chapter, we turn to term structure implementation questions. What volatility function is best if we want to employ a multifactor HJM interest rate modeling framework?

APPENDIX A: DERIVING ZERO-COUPON BOND PRICES IN THE EXTENDED MERTON/HO AND LEE MODEL

Ito's lemma gives us the same formula for changes in the value of a zero-coupon bond P , except for substitution of the slightly more complex expression for dr :

$$\begin{aligned} dP &= P_r dr + \frac{1}{2} P_{rr} (dr)^2 + P_t \\ &= P_r [a(t) + \sigma dZ] + \frac{1}{2} P_{rr} \sigma^2 + P_t \end{aligned}$$

We form a no-arbitrage portfolio with value W as in the first section in this chapter so that the coefficient of the dz term is zero. We get the same hedge ratio as given in the original Merton model:

$$w = \frac{-P_{1_r}}{P_{2_r}}$$

By applying the no-arbitrage condition that $dW = rW$, we are led to a no-arbitrage condition closely related to that of the original Merton model:

$$-\lambda = \frac{P_{1_r}a(t) + \frac{1}{2}P_{1_{rr}}\sigma^2 + P_{1_t} - rP_1}{\sigma P_{1_r}} = \frac{P_{2_r}a(t) + \frac{1}{2}P_{2_{rr}}\sigma^2 + P_{2_t} - rP_2}{\sigma P_{2_r}}$$

In the Ho and Lee model, the market price of risk must be equal for any two zero-coupon bonds with arbitrary maturities T_1 and T_2 . In the Ho and Lee case, the market price of risk again is the fixed income counterpart of the Sharpe ratio, with expected return on each bond equal to

$$\text{Drift} = \text{Expected return} = P_r a(t) + \frac{1}{2}P_{rr}\sigma^2 + P_t$$

Now the value of the zero-coupon bond price P for any given maturity is fixed by the shape of the yield curve, which Ho and Lee seek to match perfectly. Our mission in solving the partial differential equation in the Ho and Lee model is to find the relationship between the drift term $a(t)$ and the bond price P . The partial differential equation, which must be solved, comes from rearranging the no-arbitrage condition above, and our continued assumption that λ is constant:

$$P_r[a(t) + \lambda\sigma] + \frac{1}{2}P_{rr}\sigma^2 + P_t - rP = 0$$

which must hold subject to the fact that the value of a zero-coupon bond must equal 1 at maturity.

$$P(r, T, T) = 1$$

We use an educated guess to postulate a solution and then see what must be true for our guess to be correct (if it is possible to make it correct). We guess that the solution P is closely related to the Merton model:

$$P(r, t, T) = e^{-r\tau + G(t, T)}$$

We seek to find the function G , which satisfies the partial differential equation. We take the partial derivatives of P and substitute them into the partial differential equation:

$$\begin{aligned} P_r &= -\tau P \\ P_{rr} &= \tau^2 P \\ P_t &= (r + G')P \end{aligned}$$

When we substitute these partial derivatives into the partial differential equation and simplify, we get the following relationship:

$$- [a(t) + \lambda\sigma]\tau + \frac{1}{2}\sigma^2\tau^2 + G' = 0$$

We solve this differential equation in G by translating the derivative of G with respect to the current time t (not τ) into this expression:

$$G(t, T) = -\frac{\lambda\sigma\tau^2}{2} + \frac{1}{6}\sigma^2\tau^3 - \int_t^T a(s)(T-s)ds$$

Therefore, the value of a zero-coupon bond in the extended Merton/Ho and Lee model is given by the equation

$$P(r, t, T) = e^{-r\tau - \frac{\lambda\sigma\tau^2}{2} + \frac{1}{6}\sigma^2\tau^3 - \int_t^T a(s)(T-s)ds}$$

APPENDIX B: DERIVING ZERO-COUPON BOND PRICES IN THE VASICEK MODEL

We use the notation of Chen (1992) to answer this question using the same process as in the Merton/Ho and Lee models. By Ito's lemma, movements in the price of a zero-coupon bond is

$$\begin{aligned} dP &= P_r dr + \frac{1}{2}P_{rr}(dr)^2 + P_t \\ &= P_r[\alpha(\mu - r)dt + \sigma dZ] + \frac{1}{2}P_{rr}\sigma^2 + P_t \end{aligned}$$

Using exactly the same no-arbitrage argument that we used previously, we can eliminate the dz term by choosing the hedge ratio necessary to eliminate interest rate risk (which is what dz represents). The partial differential equation consistent with a no-arbitrage bond market in the Vasicek model is

$$P_r[\alpha(\mu - r) + \lambda\sigma] + \frac{1}{2}P_{rr}\sigma^2 + P_t - rP = 0$$

and, as above, it must be solved subject to the boundary condition that a zero-coupon bond's price at maturity equals its principal amount, 1:

$$P(r, T, T) = 1$$

The market price of risk λ is assumed to be constant. As a working assumption, we guess that the zero-coupon bond price has the solution

$$P(r, t, T) = P(r, \tau) = e^{-rF(\tau) - G(\tau)}$$

where F and G are unknown functions of $\tau = T - t$. If our working assumption is correct, we will be able to obtain solutions for F and G . We know that

$$\begin{aligned} P_r &= -FP \\ P_{rr} &= F^2P \\ P_t &= (-rF_t - G_t)P \end{aligned}$$

By replacing the derivatives of the zero-coupon bond price P in the partial differential equation above and rearranging, we know that the following relationship must hold:

$$r[\alpha F - F_t - 1] + \left[\frac{1}{2} F^2 \sigma^2 - F(\alpha \mu + \lambda \sigma) - G_t \right] = 0$$

This relationship must hold for all values of r , so the coefficient of r must be zero.

$$\alpha F - F_t - 1 = 0$$

We can solve this partial differential equation by rearranging it until we have

$$\frac{\alpha F_t}{1 - \alpha F} = -\alpha$$

We then take the integral of both sides such that

$$\int_t^T \frac{\alpha F_t(s, T)}{1 - \alpha F(s, T)} ds = - \int_t^T \alpha ds$$

Evaluating the integrals on both sides of the equation leaves the relationship

$$\ln[1 - \alpha F(t, T)] = -\alpha(T - t)$$

Calculating the exponential of both sides defines F :

$$F(t, T) = F(\tau) = \frac{1}{\alpha} (1 - e^{-\alpha \tau})$$

To determine the value of the function G , we must solve the partial differential equation

$$\frac{1}{2}F^2\sigma^2 - F(\alpha\mu + \lambda\sigma) - G_t = 0$$

or

$$G_t = \frac{1}{2}F^2\sigma^2 - F(\alpha\mu + \lambda\sigma)$$

We can calculate that

$$\int_t^T F(s, T)ds = \frac{1}{\alpha}[\tau - F(\tau)]$$

and that

$$\int_t^T F^2(s, T)ds = \frac{1}{\alpha^2}[\tau - F(\tau)] - \frac{F^2(\tau)}{2\alpha}$$

We can take the integral of both sides of the equation above to solve for G :

$$\begin{aligned} \int_t^T G_s(s, T)ds &= \int_t^T \left[\frac{1}{2}F(s, T)^2\sigma^2 - F(s, T)(\alpha\mu + \lambda\sigma) \right] ds \\ &= \frac{1}{2}\sigma^2 \left[\frac{1}{\alpha^2}(\tau - F(\tau)) - \frac{F^2(\tau)}{2\alpha} \right] - \frac{\alpha\mu + \lambda\sigma}{\alpha}[\tau - F(\tau)] \\ &= \left[\frac{\sigma^2}{2\alpha^2} - \mu - \frac{\lambda\sigma}{\alpha} \right] [\tau - F(\tau)] - \frac{\sigma^2}{4\alpha}F^2(\tau) \end{aligned}$$

Since

$$\int_t^T G_s(s, T)ds = G(T, T) - G(t, T)$$

and since $G(T, T)$ must be zero for the boundary condition that $P(T, T) = 1$,

$$G(t, T) = G(\tau) = \left[\mu + \frac{\lambda\sigma}{\alpha} - \frac{\sigma^2}{2\alpha^2} \right] [\tau - F(\tau)] + \frac{\sigma^2}{4\alpha}F^2(\tau)$$

This means that the value of a zero-coupon bond in the Vasicek model is

$$P(r, t, T) = e^{-F(t, T)r - G(t, T)}$$

$$= \exp \left[-rF(\tau) - \left(\mu + \frac{\lambda\sigma}{\alpha} - \frac{\sigma^2}{2\alpha^2} \right) [\tau - F(\tau)] - \frac{\sigma^2 F^2(\tau)}{4\alpha} \right]$$

APPENDIX C: VALUING ZERO-COUPON BONDS IN THE EXTENDED VASICEK MODEL

The partial differential equation changes only very slightly from the one we used in the third section of this chapter:

$$P_r[\alpha(\mu - r) + \lambda(t)\sigma] + \frac{1}{2}P_{rr}\sigma^2 + P_t - rP = 0$$

subject to the usual requirement that the bond's price equals 1 at maturity:

$$P(r, T, T) = 1$$

We can rewrite the partial differential equation as

$$P_r[(\alpha\mu + \lambda(t)\sigma) - \alpha r] + \frac{1}{2}P_{rr}\sigma^2 + P_t - rP = 0$$

and using the definition

$$\theta(t) = \alpha\mu + \lambda(t)\sigma$$

we can simplify the partial differential equation to the point that it looks almost identical to that of the Ho and Lee model earlier in the chapter:

$$P_r[\theta(t) - \alpha r] + \frac{1}{2}P_{rr}\sigma^2 + P_t - rP = 0$$

As in the third section in this chapter, we assume that the solution to the pricing of zero-coupon bonds takes the form

$$P(r, t, T) = P(r, \tau) = e^{-rF(\tau) - G(\tau)}$$

where F and G are unknown functions of $\tau = T - t$. Using the partial derivatives of P under the assumption that we have guessed the functional form correctly, we know that the following relationship must hold:

$$r[\alpha F - F_t - 1] + \left[\frac{1}{2} F^2 \sigma^2 - F\theta(t) - G_t \right] = 0$$

From Appendix B, we can prove

$$F(t, T) = F(\tau) = \frac{1}{\alpha} (1 - e^{-\alpha\tau})$$

We now must solve for G given the following equation:

$$G_t = \frac{1}{2} F^2 \sigma^2 - F\theta(t)$$

We do this by taking the integral of both sides and making use of the integral of F^2 , which was given in Appendix B, to arrive at the solution for G :

$$G(t, T) = G(\tau) = \int_t^T F(s, T)\theta(s)ds - \frac{\sigma^2}{2\alpha^2} [\tau - F(\tau)] + \frac{\sigma^2}{4\alpha} F^2(\tau)$$

Therefore, under the extended Vasicek model, the value of a zero-coupon bond is given by the formula

$$\begin{aligned} P(r, t, T) &= e^{-F(t, T)r - G(t, T)} \\ &= \exp \left[-rF(\tau) - \int_t^T F(s, T)\theta(s)ds + \left(\frac{\sigma^2}{2\alpha^2} \right) [\tau - F(\tau)] - \frac{\sigma^2 F^2(\tau)}{4\alpha} \right] \end{aligned}$$

NOTES

1. For more on Ito's lemma, see Shimko (1992) for implementations in applied finance.
2. If there are two random factors driving movements in interest rates, two bonds would be necessary to eliminate the risk of these random factors. n bonds would be necessary for an n -factor model. This parallels the HJM hedging example in Chapter 12.
3. We have implicitly assumed in this chapter that we are modeling the risk-free term structure. We will relax this assumption in Chapter 17.
4. To be precise, a Gaussian random walk.
5. We insert the minus sign in front of λ to be consistent with other authors' notation and subsequent chapters.
6. We could change the sign of this expression to make it positive since the sign of the Wiener process dz is arbitrary.
7. For example, consider this term structure model, where zero-coupon bond yields are a linear function of years to maturity: $y(\tau) = .05 + 0.1\tau$. It is clearly a ridiculous model, but it can be made to fit the yield curve exactly in a manner similar to the Ho and Lee model.
8. The same objection applies to the Merton/Ho and Lee models and a wide range of other models, which assume a constant volatility of interest rates, regardless of the level

of short-term interest rates. Since the storage of cash is not costless, as we discussed earlier in this book, negative rates are becoming increasingly common in countries like Japan and the United States, which are recovering from serious financial crises.

9. We refer the serious student of term structure model history to these excellent references by Dai, Duffie, Kan, Singleton, and Jarrow:

- Qiang Dai and Kenneth J. Singleton, "Specification Analysis of Affine Term Structure Models," *Journal of Finance* 55, no. 5 (October 2000).
- Qiang Dai and Kenneth J. Singleton, "Term Structure Dynamics in Theory and Reality," *Review of Financial Studies* 16, no. 3 (Fall 2003).
- Darrell Duffie and Rui Kan, "Multi-Factor Term Structure Models," *Philosophical Transactions: Physical Sciences and Engineering* 347, no. 1684, Mathematical Models in Finance (June 1994).
- Darrell Duffie and Kenneth J. Singleton, "An Econometric Model of the Term Structure of Interest-Rate Swap Yields," *Journal of Finance* 52, no. 4 (September 1997).
- Darrell Duffie and Kenneth J. Singleton, "Modeling Term Structures of Defaultable Bonds," *Review of Financial Studies* 12, no. 4 (1999).
- Robert A. Jarrow, "The Term Structure of Interest Rates," *Annual Review of Financial Economics* 1 (October 2009).

Estimating the Parameters of Interest Rate Models

In this chapter, we continue from our theoretical overview of term structure models to their practical application. Here we fit the parameters for the term structure model that is most consistent with the assets and liabilities being modeled and the marketplace in which they are traded.

REVISITING THE MEANING OF NO ARBITRAGE

As we discussed in prior chapters, the meaning of *no-arbitrage* assumptions should be the same for finance academics and financial market participants. Sadly, however, that is not always the case. Academics often assume a stylized market of zero-coupon bonds in which one can buy or sell them in unlimited amounts without changing the price. Models that produce zero-coupon bonds, and whose theoretical valuations match their input values, are judged to be no arbitrage. This is a basic characteristic of the one-, two-, and three-factor Heath Jarrow, and Morton (HJM) models discussed and illustrated in Chapters 6 through 9. This narrow view of the meaning of no arbitrage is necessary but not sufficient. We need all securities whose value should stem from the relevant risk-free yield curve to have a theoretical valuation that matches observed prices as well. If a perfect match is not possible, the best match one can obtain, within the model structure chosen, is the objective. We obtain this “best match” by selecting the model parameters that produce the best match.

A FRAMEWORK FOR FITTING TERM STRUCTURE MODELS

In Chapters 6 through 9, we took the initial zero-coupon bond prices with maturities in one, two, three, and four years as observable. Most likely, they were obtained from using the advanced yield curve smoothing techniques of Chapter 5 on the observable prices of coupon-bearing bonds (which we discuss in Chapter 17). As we saw in each of these chapters, the observable (or derived) zero-coupon bonds are not sufficient to determine what interest rate volatility assumptions are best. In fact, we showed in Chapters 6 through 9 that four different assumptions about interest rate volatility can

be implemented using the same zero-coupon bond prices, and all of them are no arbitrage in the academic sense.

Now we add additional securities, which provide us with the option of using two different approaches to fitting interest rate volatility:

1. *Historical volatility*: Of course we always have the opportunity to use interest rate volatility that is derived from statistical analysis of past interest rate movements. This is what we did in Chapters 6 through 9 using U.S. Treasury data from 1962 to 2011. This exercise need not be repeated here.¹
2. *Implied volatility*: In this chapter, we focus on the second alternative in which we use market prices of other securities to “imply” the best-fitting interest rate volatility assumptions in order to ensure that these other securities are priced in a no-arbitrage sense as well.

The process for implying the interest rate volatility functions is very similar to that of implying stock return volatility in the Black-Scholes (1973) model for options on common stock. The difference is simple. In the Black-Scholes model, a single formula for options value is assumed to be true, and one equation in one unknown is solved to obtain the stock return volatility consistent with an observed option price. As is well known, the theory implies that the volatility should be the same for all strike prices at a given options maturity, but this is not supported by empirical analysis in which one minimizes the sum of squared prices errors on n options with a single-stock return volatility. In optimizing interest rate volatility, however, we can employ whatever interest rate volatility function we choose as long as the volatility functions are consistent with no arbitrage in the HJM sense as discussed in Chapter 6.

This allows interest rate volatility take on any number of forms: to be a constant, to take on a specific functional form like most of the one-factor models in Chapter 13, to be a function of multiple parameters, and to vary with time to maturity and the level of interest rates in a complex way. We illustrate the process of parameter fitting in a series of increasingly complex examples, with more observable securities that are priced off the zero-coupon yield curve.

FITTING ZERO-COUPON BOND PRICES AND VOLATILITY PARAMETERS JOINTLY

In many markets, like the market for mortgage-backed securities in the United States, there are no observable bonds that are noncallable from which zero-coupon bond prices can be derived using the smoothing techniques in Chapter 5. Instead, the observable securities are either individual mortgage loans that exhibit complex prepayment behavior or packages of these individual mortgage securities. In this case, it is necessary and desirable to optimize best-fitting zero-coupon bond prices, prepayment parameters, default parameters, and interest rate volatility simultaneously. This is easy to do even in common spreadsheet software and the process by which one does so is the subject of the rest of this book. In the remainder of this chapter, we take the zero-coupon bond prices as given and derive the relevant volatility functions from observable prices of other securities.

STEPS IN FITTING THE INTEREST RATE VOLATILITY ASSUMPTIONS

In each of the interest rate volatility fitting examples that follow, we pursue the same process:

1. We select the number of risk factors driving the risk-free yield curve.
2. We select the functional form, which describes how volatility varies (if it varies) by interest rate level and maturity.
3. We obtain the prices of other securities in addition to the four observable zero-coupon bond prices used in Chapters 6 through 9.
4. We iteratively solve for the interest rate volatility parameters, which minimize the sum of squared pricing errors of the other securities in order to ensure that we come as close as possible to a no-arbitrage interest rate framework.

In this chapter, we restrict the nature of “other securities” to a simple type: bonds that pay a fixed annual interest rate with a single option to prepay prior to maturity; that is, bonds with an embedded European option to prepay. In later chapters, we show how to use the HJM framework for a wide variety of security types. All of them can be used to derive interest rate volatility in a manner similar to what we do in this chapter on similar securities. In each of the examples that follow, we use the three-factor HJM example of Chapter 9 as our base case.

Example 1: Fitting Interest Rate Volatility When Six Callable Bonds Are Observable

In this example and in Examples 2 and 3, we use Chapter 9’s three-factor HJM example as a base case. For Example 1, we make these assumptions:

1. There are three factors driving random interest rates.
2. These risk factors take on various values at different levels of interest rates:
 - *Risk factor 1* is a constant (σ_1) for all maturities of forward rates and all interest rate levels.
 - *Risk factor 2* takes on three values. The first value (σ_2 low) is for all levels of the one-period spot rate equal to 2.00 percent or below. The second value (σ_2 mid) is for all levels of the one-period spot rate from 2.01 to 3.99 percent. The third value (σ_2 high) is for all levels of the one-period spot rate of 4.00 percent or higher.
 - *Risk factor 3* takes on one value (σ_3 low) for all levels of the one-period spot rate equal to 2.99 percent or below and a second value (σ_3 high) for all levels of the one-period spot rate above that level.
3. There are six observable callable bonds outstanding. They all pay annual coupons of 2 percent, 3 percent, or 4 percent and they are callable either at year 2 only or year 3 only. In the case of the call dates, it is assumed that the bond holders receive the coupon scheduled for the call date. The assumed observable prices are given in Exhibit 14.1.

Note that we have included for the reader’s convenience the actual values of noncallable bonds with annual coupons of 2 percent, 3 percent, and 4 percent. These

EXHIBIT 14.1 Bonds Maturing in Year 4

Name	Bonds Maturing in Year 4	Call Provisions	Actual Price
Base	Implied Value of \$100 Bond with Annual \$3 Coupon	Noncall	104.707100
Base	Implied Value of \$100 Bond with Annual \$2 Coupon	Noncall	100.833236
Base	Implied Value of \$100 Bond with Annual \$4 Coupon	Noncall	108.580963
Security 1	Implied Value of \$100 Bond with Annual \$3 Coupon	Call in Year 3	104.350000
Security 2	Implied Value of \$100 Bond with Annual \$2 Coupon	Call in Year 3	100.750000
Security 3	Implied Value of \$100 Bond with Annual \$4 Coupon	Call in Year 3	107.600000
Security 4	Implied Value of \$100 Bond with Annual \$3 Coupon	Call in Year 2	103.950000
Security 5	Implied Value of \$100 Bond with Annual \$2 Coupon	Call in Year 2	100.750000
Security 6	Implied Value of \$100 Bond with Annual \$4 Coupon	Call in Year 2	106.250000

values are simple to calculate using the formulas of Chapter 4. The zero-coupon bond prices that we used as inputs in Chapters 6 through 9 remain unchanged, and the bushy trees that are created in this example and examples 2 and 3 preserve these values. We reproduce in Exhibit 14.2 the cash flow table for the noncallable bond with a 3 percent coupon. For any interest rate volatility assumptions, that bond continues to have the same price we found in Chapters 6 through 9.

What happens if this bond is callable? The cash flows at time $T = 1$ and time $T = 2$ are unchanged. At time $T = 3$, we are certain to receive the interest payment of \$3. The only question is whether to call the bond and pay \$100 at that time instead of paying \$103 at time $T = 4$. A rational bond issuer will call the bond whenever the present value of paying \$103 at time $T = 4$ is more than \$100 paid at time $T = 3$ from the perspective of time $t = 3$. (We revisit this issue in detail in Chapter 21 for European call provisions and Chapter 27 for American call provisions.) For now, we remind the reader that this decision depends on the one-period spot rate that prevails at time $t = 3$ on all 64 nodes of the three-factor bushy tree that we build using the previous three risk factors, exactly as we did in Chapter 9. Exhibit 14.3 is an example evaluating the decision to call or not call the bond at each of these 64 scenarios using the best-fitting values for risk factor volatility that we solve for next.

We construct a modified cash flow table for each of the six callable securities that looks exactly like Exhibit 14.3. We now need to solve for the single parameter for risk factor 1, the three parameters for risk factor 2, and the two parameters for risk factor 3 that provide the best fit. This solution can be achieved using the “solver” function in common spreadsheet software or the advanced nonlinear optimization

EXHIBIT 14.2 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.00	3.00	3.00	3.00	103.00
2		3.00	3.00	3.00	103.00
3		3.00	3.00	3.00	103.00
4		3.00	3.00	3.00	103.00
5			3.00	3.00	103.00
6			3.00	3.00	103.00
7			3.00	3.00	103.00
8			3.00	3.00	103.00
9			3.00	3.00	103.00
10			3.00	3.00	103.00
11			3.00	3.00	103.00
12			3.00	3.00	103.00
13			3.00	3.00	103.00
14			3.00	3.00	103.00
15			3.00	3.00	103.00
16			3.00	3.00	103.00
17				3.00	103.00
18				3.00	103.00
19				3.00	103.00
20				3.00	103.00
21				3.00	103.00
22				3.00	103.00
23				3.00	103.00
24				3.00	103.00
25				3.00	103.00
26				3.00	103.00
27				3.00	103.00
28				3.00	103.00
29				3.00	103.00
30				3.00	103.00
31				3.00	103.00
32				3.00	103.00
33				3.00	103.00
34				3.00	103.00
35				3.00	103.00
36				3.00	103.00
37				3.00	103.00
38				3.00	103.00
39				3.00	103.00
40				3.00	103.00
41				3.00	103.00
42				3.00	103.00

EXHIBIT 14.2 (Continued)

Row Number	Current Time				
	0	1	2	3	4
43				3.00	103.00
44				3.00	103.00
45				3.00	103.00
46				3.00	103.00
47				3.00	103.00
48				3.00	103.00
49				3.00	103.00
50				3.00	103.00
51				3.00	103.00
52				3.00	103.00
53				3.00	103.00
54				3.00	103.00
55				3.00	103.00
56				3.00	103.00
57				3.00	103.00
58				3.00	103.00
59				3.00	103.00
60				3.00	103.00
61				3.00	103.00
62				3.00	103.00
63				3.00	103.00
64				3.00	103.00
Risk-Neutral Value =					104.7070997370

techniques found in Chapter 36 when we discuss optimal capital structure. We take the former approach with example 1. We go through these steps in an automated fashion when using the “solver” functionality:

1. Guess the six parameter values.
2. Value each of the callable bonds.
3. Calculate the sum of squared pricing errors compared to observable prices.
4. If the guess can be improved, go to step 1 and repeat. If the guess cannot be improved, we produce the final interest rate volatility values.

In step 2, valuation of the callable bonds is done by taking the modified cash flow table for that parameter set for each bond and multiplying those cash flows by the probability-weighted discount factors like those from Chapter 9. The best-fitting parameter values that emerge from this process are as shown in Exhibit 14.4.

These volatility parameters result in the probability-weighted discount factors in the table in Exhibit 14.5.

EXHIBIT 14.3 Maturity of Cash Flow Received

Row Number	0	1	2	3	4	Call Flag (1 = yes)	NPV Factor	NPV of Year 4 Payment	NPV Benefit If Called
1	0	3	3	3	103	0	0.853834418	87.94494504	
2	0	3	3	3	103	0	0.890549119	91.72655926	
3	0	3	3	103	0	1	1.006463481	103.6657385	3.665738549
4	0	3	3	3	103	0	0.936308227	96.43974743	
5	0	0	3	3	103	0	0.859071384	88.48435252	
6	0	0	3	3	103	0	0.896011273	92.28916115	
7	0	0	3	103	0	1	1.012636592	104.301569	4.301569011
8	0	0	3	3	103	0	0.942051044	97.03125751	
9	0	0	3	3	103	0	0.969558868	99.8645634	
10	0	0	3	103	0	1	0.975505626	100.4770795	0.477079479
11	0	0	3	103	0	1	0.98424679	101.3774193	1.377419324
12	0	0	3	103	0	1	0.977833561	100.7168568	0.716856753
13	0	0	3	3	103	0	0.920704487	94.83256215	
14	0	0	3	3	103	0	0.960294587	98.91034247	
15	0	0	3	103	0	1	0.977697321	100.7028241	0.70282408
16	0	0	3	3	103	0	0.95828644	98.70350332	
17	0	0	0	3	103	0	0.918750002	94.63125017	
18	0	0	0	3	103	0	0.958256059	98.70037412	
19	0	0	0	103	0	1	0.975621851	100.4890506	0.489050621
20	0	0	0	3	103	0	0.956252175	98.49397404	
21	0	0	0	3	103	0	0.941218292	96.94548412	
22	0	0	0	3	103	0	0.946991224	97.54009612	
23	0	0	0	103	0	1	0.981658004	101.1107744	1.110774365
24	0	0	0	3	103	0	0.962168488	99.10335429	
25	0	0	0	3	103	0	0.962657985	99.15377249	
26	0	0	0	3	103	0	0.968562417	99.76192895	
27	0	0	0	103	0	1	0.977241365	100.6558606	0.655860582

28	0	0	0	3	103	0	0.970873782	99.9999996	
29	0	0	0	3	103	0	0.943464404	97.17683364	
30	0	0	0	3	103	0	0.949251113	97.77286461	
31	0	0	0	103	0	1	0.98400062	101.3520639	1.352063882
32	0	0	0	3	103	0	0.964464595	99.33985332	
33	0	0	0	3	103	0	0.943863066	97.21789576	
34	0	0	0	3	103	0	0.949652219	97.81417858	
35	0	0	0	103	0	1	0.98441641	101.3948902	1.394890241
36	0	0	0	3	103	0	0.96487213	99.38182942	
37	0	0	0	3	103	0	0.962657985	99.15377249	
38	0	0	0	3	103	0	0.968562417	99.76192895	
39	0	0	0	103	0	1	0.977241365	100.6558606	0.655860582
40	0	0	0	3	103	0	0.970873782	99.9999996	
41	0	0	0	103	0	1	0.971284026	100.0422546	0.042254641
42	0	0	0	103	0	1	0.977241365	100.6558606	0.655860582
43	0	0	0	103	0	1	0.985998082	101.5578024	1.557802422
44	0	0	0	103	0	1	0.979573442	100.8960645	0.896064497
45	0	0	0	3	103	0	0.964955261	99.39039184	
46	0	0	0	3	103	0	0.970873782	99.9999996	
47	0	0	0	103	0	1	0.979573442	100.8960645	0.896064497
48	0	0	0	103	0	1	0.973190664	100.2386384	0.238638374
49	0	0	0	3	103	0	0.920942495	94.85707703	
50	0	0	0	3	103	0	0.96054283	98.93591148	
51	0	0	0	103	0	1	0.977950063	100.7288565	0.728856463
52	0	0	0	3	103	0	0.958534164	98.72901886	
53	0	0	0	3	103	0	0.943464404	97.17683364	
54	0	0	0	3	103	0	0.949251113	97.77286461	

(Continued)

EXHIBIT 14.3 (Continued)

Row Number	0	1	2	3	4	Call Flag (1 = yes)	NPV Factor	NPV of Year 4 Payment	NPV Benefit If Called
55	0	0	0	103	0	1	0.98400062	101.3520639	1.352063882
56	0	0	0	3	103	0	0.964464595	99.33985332	
57	0	0	0	3	103	0	0.964955261	99.39039184	
58	0	0	0	3	103	0	0.970873782	99.9999996	
59	0	0	0	103	0	1	0.979573442	100.8960645	0.896064497
60	0	0	0	103	0	1	0.973190664	100.2386384	0.238638374
61	0	0	0	3	103	0	0.945715876	97.40873525	
62	0	0	0	3	103	0	0.951516394	98.00618858	
63	0	0	0	103	0	1	0.986348827	101.5939292	1.593929208
64	0	0	0	3	103	0	0.966766182	99.57691673	
Risk-Neutral Value =									104.3315124

EXHIBIT 14.4 Input Variables

Sigma 1	0.0274%
Sigma 3 Low	0.1529%
Sigma 3 High	1.0525%
Sigma 2 Low	0.4235%
Sigma 2 Mid	1.3792%
Sigma 2 High	5.0703%

EXHIBIT 14.5 Probability-Weighted Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1	0.1246257233	0.0152422660	0.0018050848	0.0015412435
2		0.1246257233	0.0152422660	0.0018050848	0.0016075167
3		0.2492514466	0.0304845321	0.0036101696	0.0036335038
4		0.4985028932	0.0609690641	0.0072203391	0.0067604629
5			0.0153357540	0.0018161562	0.0015602078
6			0.0153357540	0.0018161562	0.0016272964
7			0.0306715080	0.0036323124	0.0036782125
8			0.0613430159	0.0072646248	0.0068436474
9			0.0309463446	0.0037652815	0.0036506621
10			0.0309463446	0.0037652815	0.0036730533
11			0.0618926891	0.0075305630	0.0074119325
12			0.1237853783	0.0150611261	0.0147272745
13			0.0614894041	0.0073810537	0.0067957693
14			0.0614894041	0.0073810537	0.0070879859
15			0.1229788083	0.0147621074	0.0144328729
16			0.2459576165	0.0295242149	0.0282926548
17				0.0018523209	0.0017018198
18				0.0018523209	0.0017749977
19				0.0037046418	0.0036143295
20				0.0074092836	0.0070851435
21				0.0018636820	0.0017541316
22				0.0018636820	0.0017648905
23				0.0037273641	0.0036589968
24				0.0074547282	0.0071727045
25				0.0037607637	0.0036203292
26				0.0037607637	0.0036425344
27				0.0075215274	0.0073503477
28				0.0150430548	0.0146049075
29				0.0074725181	0.0070500548
30				0.0074725181	0.0070932961

(Continued)

EXHIBIT 14.5 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
31				0.0149450361	0.0147059248
32				0.0298900722	0.0288279164
33				0.0037713312	0.0035596203
34				0.0037713312	0.0035814531
35				0.0075426625	0.0074251207
36				0.0150853249	0.0145554096
37				0.0037944626	0.0036527697
38				0.0037944626	0.0036751738
39				0.0075889252	0.0074162116
40				0.0151778503	0.0147357769
41				0.0076569268	0.0074370507
42				0.0076569268	0.0074826656
43				0.0153138537	0.0150994304
44				0.0306277074	0.0300020887
45				0.0152140705	0.0146808974
46				0.0152140705	0.0147709422
47				0.0304281411	0.0298065989
48				0.0608562821	0.0592247656
49				0.0074446886	0.0068561301
50				0.0074446886	0.0071509423
51				0.0148893772	0.0145610674
52				0.0297787545	0.0285439535
53				0.0074903504	0.0070668790
54				0.0074903504	0.0071102234
55				0.0149807008	0.0147410188
56				0.0299614015	0.0288967110
57				0.0151149376	0.0145852385
58				0.0151149376	0.0146746966
59				0.0302298751	0.0296123828
60				0.0604597502	0.0588388645
61				0.0300329010	0.0284025913
62				0.0300329010	0.0285767977
63				0.0600658021	0.0592458334
64				0.1201316041	0.1161391723

The zero-coupon bond prices that we used as input are valued in exactly the same way as in Chapter 9, and the theoretical values and input values match exactly. For the callable bonds that were observable, the fit is very good but not exactly perfect (Exhibit 14.6).

We could make the fit perfect by either decreasing the size of the time step (in which case the value will converge to the theoretical value of a callable European bond that we discuss in Chapter 21) or by decreasing the number of observable bonds or interest rate volatility parameters. The former strategy is the better strategy, because limiting the

EXHIBIT 14.6 Bonds Maturing in Year 4

Name	Bonds Maturing in Year 4	Call Provisions	Model Price	Actual Price	Squared Error
Base	Implied Value of \$100 Bond with Annual \$3 Coupon	Noncall	104.707100	104.707100	
Base	Implied Value of \$100 Bond with Annual \$2 Coupon	Noncall	100.833236	100.833236	
Base	Implied Value of \$100 Bond with Annual \$4 Coupon	Noncall	108.580963	108.580963	
Security 1	Implied Value of \$100 Bond with Annual \$3 Coupon	Call in Year 3	104.331512	104.350000	0.000342
Security 2	Implied Value of \$100 Bond with Annual \$2 Coupon	Call in Year 3	100.748886	100.750000	0.000001
Security 3	Implied Value of \$100 Bond with Annual \$4 Coupon	Call in Year 3	107.619247	107.600000	0.000370
Security 4	Implied Value of \$100 Bond with Annual \$3 Coupon	Call in Year 2	103.950321	103.950000	0.000000
Security 5	Implied Value of \$100 Bond with Annual \$2 Coupon	Call in Year 2	100.760819	100.750000	0.000117
Security 6	Implied Value of \$100 Bond with Annual \$4 Coupon	Call in Year 2	106.237123	106.250000	0.000166 0.000996

realism of the volatility model and the number of observable bonds used as inputs decreases the realism of the model, even if the model fits observable prices exactly. We show that in examples 2 and 3. The volatility assumptions that we have derived above produce a realistic “bushy tree” of short rate values as shown in Exhibit 14.7.

The yield curve shifts implied by the interest rate volatility parameters we have derived are given in Exhibit 14.8.

We have four options for further improving the realism of projected interest rate movements:

1. Find more observable securities prices to use as input
2. Combine the existing inputs with historical analysis like that of Chapter 7 or Adrian, Crump, and Moench (2012)
3. Use judgment
4. Some combination of 1, 2, and 3




EXHIBIT 14.7 Evolution of a One-Year U.S. Treasury Spot Rate, Three-Factor HJM Model with Rate- and Maturity-Dependent Volatility




EXHIBIT 14.8 Kamakura Corporation HJM Zero-Coupon Yield Curve Movements

A Note on Parameter Fitting in Example 1 Whether one uses advanced nonlinear optimization like that discussed in Chapter 36 or common spreadsheet software, it is very common for optimization problems to have multiple local optima. The user can “guide” the solution to be the most realistic answer with an educated guess at the starting values used in the optimization. We encourage readers to do so for every volatility model one might employ as a model risk assessment of parameter stability. Working hard to add more observable securities is always a good thing, particularly if the additional data reveals that the assumptions of the model are too simple to provide an accurate fit to the input data. A classic example is the Black-Scholes options model’s “volatility smile,” which is more aptly named a “volatility frown.” The frown is due to the fact that the Black-Scholes model assumes that stock return volatility is constant, but the frown proves that it is not. Fitting the frown is not the right solution. The correct solution is to use a volatility model with a richer structure, as the HJM approach permits.

Example 2: The Consequences of Fewer Inputs

We now ask this question: What are the consequences of having less input data? Assume that securities 4 through 6, the bonds callable in year 2, are no longer observable. The first implication of this is that we have only three equations, the valuation for the three observable bonds, so that we can solve for only three unknown parameters. We arbitrarily set the sigma two parameters (sigma 2 low, sigma 2 mid, and sigma 2 high) to zero and solve for the best-fitting values of the other three parameters (Exhibit 14.9).

We obtain a good fit to the three remaining observable bonds, but the theoretical values for the other three bonds drift away from the values that we could observe in Example 1.

Example 3: The Case of One Input

What if we have only one observable callable bond? In that case, we can solve for only one input. We set all of the sigma 2 and sigma 3 parameters to zero and solve for our

EXHIBIT 14.9 Bonds Maturing in Year 4

Name	Bonds Maturing in Year 4	Call Provisions	Model Price	Actual Price	Squared Error
Base	Implied Value of \$100 Bond with Annual \$3 Coupon	Noncall	104.707100	104.707100	
Base	Implied Value of \$100 Bond with Annual \$2 Coupon	Noncall	100.833236	100.833236	
Base	Implied Value of \$100 Bond with Annual \$4 Coupon	Noncall	108.580963	108.580963	
Security 1	Implied Value of \$100 Bond with Annual \$3 Coupon	Call in Year 3	104.326910	104.350000	0.000533

(Continued)

EXHIBIT 14.9 (Continued)

Name	Bonds Maturing in Year 4	Call Provisions	Model Price	Actual Price	Squared Error
Security 2	Implied Value of \$100				
	Bond with Annual \$2 Coupon	Call in Year 3	100.720693	100.750000	0.000859
Security 3	Implied Value of \$100				
	Bond with Annual \$4 Coupon	Call in Year 3	107.651271	107.600000	0.002629
Security 4	Implied Value of \$100				
	Bond with Annual \$3 Coupon	Call in Year 2	103.871326	103.950000	
Security 5	Implied Value of \$100				
	Bond with Annual \$2 Coupon	Call in Year 2	100.587339	100.750000	
Security 6	Implied Value of \$100				
	Bond with Annual \$4 Coupon	Call in Year 2	106.226133	106.250000	
					0.004021
Input variables		Sigma 1			0.0000%
		Sigma 3 Low			0.6596%
		Sigma 3 High			2.6493%
		Sigma 2 Low			0.0000%
		Sigma 2 Mid			0.0000%
		Sigma 2 High			0.0000%
1 million times sum of squared errors					4,021

only remaining interest volatility parameter. This case is the equivalent of a one-factor Ho and Lee (1986) implementation in the HJM framework. We assume that our only observable callable bond is the bond with a 3 percent coupon that is callable only at time $T = 3$. Our best-fitting parameter values are as shown in Exhibit 14.10.

While we get an exact fit for the only observable bond, our model values for the other bonds that we could observe in Example 1 are less accurate. Even more important, the simulated yield curve shifts implied by Example 3 are much less realistic and comprehensive than those we found in Example 1 (Exhibit 14.11).

We conclude with some brief notes on implementation of parameter fitting on a fully automated basis.

INTEREST RATE PARAMETER FITTING IN PRACTICAL APPLICATION

In this chapter, we noted that more data, more historical analysis, and more judgment are the ingredients necessary to fully exploit the HJM approach (and any other high-quality no-arbitrage modeling framework) for accurate risk measurement and risk management. In the chapters that remain in this book, we take on that task. There are a number of barriers to accomplishing that objective that we need to overcome:

EXHIBIT 14.10 Bonds Maturing in Year 4

Name	Bonds Maturing in Year 4	Call Provisions	Model Price	Actual Price	Squared Error
Base	Implied Value of \$100 Bond with Annual \$3 Coupon	Noncall	104.707100	104.707100	
Base	Implied Value of \$100 Bond with Annual \$2 Coupon	Noncall	100.833236	100.833236	
Base	Implied Value of \$100 Bond with Annual \$4 Coupon	Noncall	108.580963	108.580963	
Security 1	Implied Value of \$100 Bond with Annual \$3 Coupon	Call in Year 3	104.350000	104.350000	0.000000
Security 2	Implied Value of \$100 Bond with Annual \$2 Coupon	Call in Year 3	100.731641	100.750000	0.000859
Security 3	Implied Value of \$100 Bond with Annual \$4 Coupon	Call in Year 3	107.748447	107.600000	0.002629
Security 4	Implied Value of \$100 Bond with Annual \$3 Coupon	Call in Year 2	103.753160	103.950000	
Security 5	Implied Value of \$100 Bond with Annual \$2 Coupon	Call in Year 2	100.527566	100.750000	
Security 6	Implied Value of \$100 Bond with Annual \$4 Coupon	Call in Year 2	106.194074	106.250000	
					0.000000
Input variables		Sigma 1			0.7155%
		Sigma 3 Low			0.0000%
		Sigma 3 High			0.0000%
		Sigma 2 Low			0.0000%
		Sigma 2 Mid			0.0000%
		Sigma 2 High			0.0000%

1. We have been assuming a risk-free yield curve, but most yield curves involve an issuer of securities who can default. We begin the introduction to credit risk analytics in Chapters 15 and 16.
2. When we deal with an issuer who can default, we note that the credit spread is the product of supply and demand for the credit of that particular issue, and yet the default probabilities and recovery rates of that issuer are but a few of the many determinants of what the credit spread will be. We deal with that issue extensively in Chapter 17.




EXHIBIT 14.11 Kamakura Corporation HJM Zero-Coupon Yield Curve Movements

3. Many of the yield curves for credit risky issuers involve observable data from markets like the credit default swap market and the LIBOR arena where lawsuits and investigations of potential market manipulation are under way as we write this chapter.
4. Many of the yield curves are composites or indices of risky issuers, like the credit index at the heart of JPMorgan Chase's losses of \$2 billion announced on May 8, 2012. The LIBOR rate is another index. The credit risk of indices is unique because of a number of complications:
 - a. A potential default of one's counterparty
 - b. A potential default by a firm that is a component of the index, like a bank that is a member of the U.S. dollar LIBOR panel
 - c. A potential re-benchmarking of the index, such as a change in the LIBOR panel, because of credit risk concerns about one of the original panel members or because of concerns by a panel member about the process of setting the index itself
 - d. The potential manipulation of the index by market participants
 - e. A compounding of the four issues above, which is now thought of as the norm by many market participants

We now tackle these challenges with an introduction to credit risk.

NOTE

1. For an excellent example of historical term structure modeling for a five-factor model, see Adrian, Crump, and Moench, "Pricing the Term Structure with Linear Regression," Working Paper, Federal Reserve Bank of New York, May 9, 2012.

PART

Three

Risk Management Techniques for Credit Risk Analytics

An Introduction to Credit Risk

Using Market Signals in Loan Pricing and Performance Measurement

In the first 14 chapters of this book, the fundamental emphasis has been to lay the foundation for an integrated treatment of credit risk, market risk, liquidity risk, and interest rate risk. We now introduce the analysis of credit risk, building on the foundation, which we established with our prior focus on the term structure of interest rates in a multifactor Heath, Jarrow, and Morton (HJM) framework.

MARKET PRICES FOR CREDIT RISK

Our primary focus in this chapter is to introduce the use of market prices in credit risk analysis. In subsequent chapters, we discuss the state-of-the-art models that are used with this market data to construct an integrated approach to market risk, interest rate risk, liquidity risk, and credit risk. As you will see, in those chapters, as well as this one, we will emphasize best practice. That said, it is important to add a word of warning about “past practice” and some of its dangers.

For a large insurance company or pension fund, the use of market signals in credit risk analysis is accepted without question because so much of the fixed income portfolio consists of bonds issued by large, highly rated issuers. This is definitely not the case in the banking industry, especially with the demise of the purely wholesale-oriented commercial banks such as First Chicago, Continental Illinois, the old JP Morgan, and the Industrial Bank of Japan. Even today, outside of the investment portfolio, the dominant majority of transactions on the balance sheet of banks are extensions of credit to small businesses and retail clients who obviously don’t have ratings and don’t issue debt in their own name that is readily traded in the marketplace. For these institutions, using market prices of credit risk in the management of the bank was a larger leap of faith. To put the numbers in perspective, a large bank in Beijing at the time of this writing has 700 million transactions on its balance sheet, but less than 1,000 are traded with a CUSIP or ISIN number that you could look up on a common market data service.

Using market prices for credit risk analysis has enormous benefits to all financial institutions and we will spend a lot of time in subsequent chapters on the execution of these ideas. Some of the major benefits can be summarized as follows:

- *Increased accuracy in pricing:* Financial institutions create value added for their owners by skillful asset selection and asset pricing. Market prices of credit risk can be used to sharply improve the accuracy of pricing on credit extended to all counterparties from retail clients to sovereigns.
- *Increased clarity in corporate strategy:* A skillful use of market prices allows senior management to see more clearly where value has been created in asset selection. More importantly, it allows senior management to see where future opportunities lie. Most importantly, senior management can also clearly see which areas to avoid.
- *Increased sophistication in risk management:* Market prices for credit risk provide invaluable insights about the correlation of risks among borrowers and about the macroeconomic factors driving the correlation in default. This allows a sophisticated financial institution to hedge its portfolio of risks with respect to changes in these macro factors. This reduces the sharp cyclicity of the credit cycle.
- *Increased precision in measuring the safety and soundness of financial institutions:* Market prices of credit risk have helped regulators and shareholders alike more quickly and more accurately measure the safety and soundness of financial institutions.

We turn to each of these benefits in turn after a brief review of the kinds of market data that are critical in this regard. We close the chapter with words of caution about the market for credit-related securities, potential market manipulation, and the May 8, 2012, \$2 billion loss announced by JPMorgan Chase.

CRITICAL SOURCES OF MARKET DATA ON CREDIT RISK

Not too many years ago, financial institutions weighted credit risk analysis very heavily toward the use of history:

- Historical rates of default
- Historical movements in credit ratings
- Historical transition matrices of movements from one ratings category to another

We discuss these historical data techniques with a critical eye in the next chapter. It is impossible to avoid these techniques in a large financial institution, but it is critical to supplement them heavily with market data from sources such as those described in the following sections.

Bond Prices

There are many vendors who sell very large databases of daily bond prices spanning most of the developed countries and literally more than 100,000 different bond issues. An institution that does not use this data extensively in setting corporate

strategy, asset pricing, and risk management policies is destined to be a subpar performer and more likely to end up as a troubled institution. The use of these bond prices should be very pervasive around the institution, not just restricted to one or two specialists in an analytical group. In recent years, the TRACE (for Trade Reporting and Compliance Engine) system in the United States has made public each significant trade in corporate and municipal bonds. While this data is sporadic, it is the raw material from which a daily database of “on the run” credit spreads for each major bond issuer can be constructed and analyzed.

Credit Default Swap Prices

The introduction of credit default swaps in the past decade has provided financial institutions with another view of the credit risk of the top credits in the world on a daily basis. This data has been overhyped and oversold, so a sophisticated analyst has to use it with great caution. We discuss the risks in the use of credit default swap data at the end of this chapter. The use of credit default swap pricing has many potential benefits; however, the following considerations exist:

- *Transparency*: The pricing of the credit default swap itself makes the market’s assessment of the reference name’s credit risk starkly apparent in the simplest of ways—it is the insurance premium for credit risk. There’s no need to use the advanced techniques for measuring credit spread that we discuss in Chapter 17. That is, if the quote were real, not just a posting intended to mislead a potential investor in bonds. We investigate this issue next.
- *“On the run” maturities*: Credit default swaps are most frequently quoted at five-year maturities, but also at maturities from one- to 10-year maturities for the most popular reference names. These quotes are the maturities at which many new transactions are done in lending markets, so they provide a concentrated view of credit risk at key maturities. Bond prices on existing issues, by contrast, are scattered at many maturities, and data at key “new issue” maturities is sparse by definition, since all outstanding issues now have shorter “off the run” maturities.
- *Cyclicalities*: The history of credit default swaps shows the cyclicalities of credit risk at a constant maturity very clearly. We return to the drivers of this cyclicalities in depth the next chapter when discussing reduced form credit models.
- *Recovery rates*: Credit default swaps are traded in two forms. The digital default swap, a structure rarely traded, essentially pays \$1 upon default. The more common credit default swap structure pays only the “loss given default” as carefully defined in ISDA contract language (see www.isda.org for model credit default swap contracts for various types of counterparties). Comparing the pricing for these two structures allows financial institutions to get the market’s view on the expected magnitude of “loss given default,” supplementing the notoriously sparse loss given default data maintained by most financial institutions and the legacy rating agencies.

First to Default Swaps

A brief boom, now over, in “first to default swaps” provided the market’s view on the implicit “correlation” between the (standard) five reference credits underlying the

first to default swap. The reference names might be IBM, Ford, BP PLC, the Republic of Korea, and Toyota. If all of the reference names are likely to default at exactly the same time and if they all have similar probabilities of default, the pricing on the first to default swap will be very close to the pricing on a credit default swap for any one of the reference names. If, by contrast, there is very little correlation between the reference names from a credit risk perspective, first to default swap pricing will approach five times the cost of the credit default swap on any one of the reference names since you are five times more likely to have to pay on the swap.

The market's view on this degree of correlation is very valuable, when it is real, and we discuss it at length in later chapters, particularly in the context of counterparty credit risk. There is one important caveat emptor with respect to first to default swaps—those who truly believe that the pairwise correlation between any two counterparties is equal for all pairs will be eaten alive in the marketplace!

Collateralized Debt Obligations

The CDO market has collapsed in the wake of the 2006–2011 credit crisis, the classic case of a financial fad with no fundamental economic reason for being. Pricing in the collateralized debt obligation market, when it was active, could potentially provide the same kind of implied correlation estimates among larger portfolios of 100 or so reference names that typically make up the collateral pool on a CDO. Price transparency in this market was always very poor, in part because of the typical “arbitrage” cycle on Wall Street where “smart people take money from dumb people” early in the life of a new and complex instrument.¹ An increase in price transparency offers similar benefits as the first to default swap market in understanding the market's implied correlation in defaults among the reference collateral. With the market's demise, that transparency is highly unlikely to come about in the near future.

Interest Rate Swap Prices

Very large international financial institutions are typically the counterparties on the dominant proportion of interest rate swaps in major currencies. Historical movements in interest rate swap spreads, then, potentially reflect (directly and indirectly) credit spreads of these large financial institutions. A major international pension fund, for instance, has shown the median default probabilities for large international financial banks are highly statistically significant in explaining movements in interest rate swap spreads.² This is true, even in the era of very highly rated “special purpose vehicles” designed to insulate the swap market from credit risk. (We discuss the impact of potential manipulation of LIBOR on the interest rate swap market in Chapter 17. We also discuss the reasons why interest rate swap spreads to U.S. Treasuries have turned negative at 30 years [and less often at 10 years], much to the surprise of most market participants.)

Equity Prices

Equity prices contain extremely valuable information on credit quality on a much larger universe of companies than one can find in the bond markets or the markets

for over-the-counter credit derivatives. (We review this in Chapter 18 on legacy credit risk techniques like the Merton model of risky debt and in Chapters 16 and 17 on reduced form models.) Different modeling technologies make use of this equity market information in different ways, but the explanatory power of equity markets is undeniable at this point among sophisticated market participants.

We now return to the practical use of this market credit information in credit risk analytics.

INCREASED ACCURACY IN PRICING

Some of the uses of market credit information are very basic and yet they substantially improve financial institutions' performance. For example, one major lender in the Asia-Pacific region was lending to a borrower in Japan at yen LIBOR plus 1.50 percent when the bonds of the borrower were trading at 9 percent (yes, 9 percent) over the Japanese government bond rate. The maturities of the bond and the loan were similar, and the loan was unsecured just like the bond. The head of international lending sheepishly admitted they mispriced the loan and that the bank would have been much better off just buying the bonds. As simple as this comparison is, it is shocking how few financial institutions make this comparison in their lending activities—even though it's standard in their bond-trading activities. Hedge funds have arisen in large numbers to arbitrage these price discrepancies, but the problems in the credit default swap market, which we discuss in an upcoming section, are large enough that price discrepancies can still be found.

In Chapter 36, we revisit the topic of performance measurement in detail, but a preview of that chapter is useful here. In a presentation to a risk management group, two executives of JPMorgan Chase explained that the credit default swap market had changed the bank's thinking on measuring the shareholder value added on lending transactions.³ A loan, they explained, was the economic equivalent of buying a Treasury bond with the same maturity and entering into a credit default swap to provide credit protection on the borrower. Therefore, the bank measured the shareholder value added on a loan as the excess of the loan's interest rate over the sum of the Treasury yield and the pricing on the credit default swap. (Very simple. Very accurate. No capital allocation. No bureaucracy needed. More on this in Chapter 36.) It is all the more ironic that the \$2 billion credit derivatives loss that JPMorgan Chase reported on May 8, 2012, was due to trades in an instrument that provided no incremental risk-adjusted value by this measure from JPMorgan Chase itself 11 years earlier. A risk manager's work is never done.

INCREASED CLARITY IN CORPORATE STRATEGY

From a corporate strategy point of view, one of the most critical decisions faced by senior management is the proper asset allocation for the institution. Where is shareholder value added the greatest? As we saw in the previous section, credit default swap pricing is providing new insights on a small subset of major corporate and sovereign names. More than that, however, we can use market prices of debt, equity, and credit derivatives to compare asset classes and to answer questions like these:

- How much of the credit spread (see Chapter 17) is the loss component and how much is the liquidity component (i.e., everything other than the risk-adjusted potential loss)?
- How correlated are credit risk-adjusted returns in various sectors?
- What are the drivers of credit risk in each of these sectors?

Van Deventer and Imai (2003) analyze these issues in great length and we will devote considerable time to them in later chapters, as well.

INCREASED SOPHISTICATION IN RISK MANAGEMENT

Financial institutions have always needed to know how to simulate true credit-adjusted portfolio performance for every counterparty from retail borrowers to major sovereign and corporate credits. They have always needed to know the default probability for each borrower and the macroeconomic factors that drive that default probability and that cause correlation in events of default. If we don't use the market prices of credit risk discussed previously, we have to make some very simple assumptions to approximate true credit risk. Representative legacy examples of this are the Moody's Investors Service "diversity score" concept and the simple correlation and default probability assumptions in Standard & Poor's CDO Evaluator.⁴

If we use the credit information discussed previously, however, we can make some substantial leaps forward in risk assessment from the transaction level to the portfolio level, the full balance sheet of a large financial institution. We do this in detail in Chapters 36 through 41.

INCREASED PRECISION IN MEASURING THE SAFETY AND SOUNDNESS OF FINANCIAL INSTITUTIONS

International supervisory agencies have worked for years on capital regulations labeled Basel II, Basel III, and Solvency II. The proposed capital accords are well intentioned but highly convoluted, and many of their fundamental assumptions (such as the "risklessness" of sovereign debt) have been proven false. More important, they result in financial ratios that still have not yet been proven as predictors of financial institutions failure. The 2006–2011 credit crisis has prompted many critics of the capital rules to call for their repeal and replacement with outright prohibitions on risky trading activities by financial institutions that are supported by government deposit insurance, such as the proposed "Volcker Rule" in the United States.

The market signals on credit risk discussed above have been proposed as key indicators of the risk of major international financial institutions:

- Their bond credit spreads are visible in the marketplace.
- Their risk as embedded in traditional and digital credit default swaps is, in theory, visible in the marketplace. More on this controversial proposal next.
- Their default probabilities can be derived from stock prices and other data as we describe in Chapters 16 and 17.

All of these measures outperform the Basel II and III capital calculations as predictors of financial institutions' failure. In part, this is the obvious result of the fact that the Basel process constrains the risk assessment to one financial ratio (or two or three in Basel III). The reduced form credit modeling approach (see Chapter 16) allows *any* financial ratio to be added as an explanatory variable, an unconstrained approach. If the Basel ratio has predictive power, the reduced form modeling approach can employ it along with all the other key explanatory variables like stock prices and macroeconomic factors. The reduced form modeling approach to credit risk assessment cannot be outperformed by any single financial ratio, because that ratio can in turn be used as an input to the reduced form approach.

For that reason, efficient use of all available market-based credit information is the key to success in assessing the safety and soundness of financial institutions worldwide.

We plunge into that task in detail in the next three chapters. Before that, however, we warn the reader of the dangers of taking credit default swap quotes at face value as “observable market prices.”

CREDIT DEFAULT SWAPS: THE DANGERS OF MARKET MANIPULATION

Financial market participants are often trained to leap into number crunching without questioning the integrity of financial market data that is visible on a computer screen. This tendency can and has been exploited ruthlessly by market participants who have the ability to artificially move the numbers on the screen. The most primitive example is the allegations of LIBOR manipulation that we discuss in Chapter 17 on the credit spreads of risky debt issuers. In the remainder of this chapter, we discuss potential manipulation of credit default swap (CDS) spreads by the suppression of competition and control of price dissemination. We applaud the efforts of governments and financial exchanges around the world to combat these anticompetitive practices.

In the rest of this chapter, we look at weekly credit default swap data since the week ended July 16, 2010, to measure how much of existing single-name CDS trading is the exchange of positions among a small group of dealers and how much is true end-user trading. We analyze data from the Depository Trust and Clearing Corporation (DTCC) and made available from www.dtcc.com.

We are interested in how much of the live trades in the CDS trade warehouse maintained by DTCC represents transactions between two dealers and how much involves nondealer end users on at least one side of the trade. The DTCC trade warehouse explanation above does not identify the dealers supplying data to DTCC at the current time. On June 2, 2010, however, DTCC listed the 14 “families” of contributing dealers as follows:

- Bank of America Merrill Lynch
- Barclays
- BNP Paribas
- Citibank
- Credit Suisse
- Deutsche Bank

- Goldman Sachs
- HSBC
- JPMorgan Chase
- Morgan Stanley
- Nomura
- Royal Bank of Scotland
- UBS
- UniCredit

We presume the list of contributing dealers is somewhat longer than it was in 2010, but it is not clear that the highly concentrated nature of CDS trading has changed during that time interval. We analyze that issue below.

The table in Exhibit 15.1 is released weekly by the DTCC. It lists the number of contracts and notional values that are still “live” in the trade warehouse at that time. DTCC is careful to note that the data may not be complete if the inclusion of a particular trade would identify the parties to the transaction. We report here the composition of trades by four categories of seller–buyer combinations in the trading of single-name credit default swaps.

Based on the number of trades reported by DTCC, the breakdown among these four buyer–seller combinations is as shown in Exhibit 15.2.

In percentage terms, the dealer–dealer share of single-name CDS contracts traded has ranged from 79.86 percent to 88.47 percent with an average of 81.68 percent. Trades between nondealer parties were only 0.12 percent of total trades on average. The results are similar if one looks at the dealer–dealer share of trades by notional principal. Dealer–dealer trades ranged from 73.40 percent to 88.36 percent of the total, with an average of 76.28 percent. Trades involving two nondealers were only 0.16 percent of total notional principal traded. There is anecdotal evidence that the dealer–dealer share of total trades is overstated by hedge funds that execute CDS trades via a CDS dealer in prime broker role, but we have no alternative, until the DTCC makes more open disclosure, but to use the DTCC definitions.

Well-intentioned academics and journalists, plus many politicians (well intentioned or not), discuss trading levels on single-name credit default swaps as if such numbers were gifts from heaven itself. Instead, we see that roughly 76 to 82 percent of all single-name credit default swaps are trades between Bill Smith at Goldman Sachs and John Smith at JPMorgan Chase or other dealer firms. As a sophisticated observer of financial markets, should an investor take these traded prices as meaningful information on the credit worthiness of the underlying reference name? Is it safe to assume that Bill Smith and John Smith do not talk to each other?

EXHIBIT 15.1 Contracts and Notional Values “Live” in Trade Warehouse

Combination of Seller and Buyer	Seller	Buyer
Dealer Dealer	Dealer	Dealer
Dealer Nondealer	Dealer	Nondealer
Nondealer Dealer	Nondealer	Dealer
Nondealer Nondealer	Nondealer	Nondealer

EXHIBIT 15.2 Dealer and Nondealer Single-Name Transactions Outstanding in DTCC CDS Warehouse

Week Ended	Number of Contracts by Seller/Buyer Combination						Differential
	Dealer Dealer	Dealer Nondealer	Nondealer Dealer	Nondealer Nondealer	Subtotal	Reported Total	
20100716	1,780,463	182,384	150,791	2,795	2,116,433	2,116,433	0
20100723	1,780,076	185,012	153,225	2,795	2,121,108	2,121,108	0
20100730	1,782,563	186,941	155,006	2,755	2,127,265	2,127,265	0
20100806	1,787,765	188,671	157,025	2,743	2,136,204	2,136,204	0
20100813	1,782,549	192,840	160,952	2,747	2,139,088	2,139,088	0
20100820	1,779,499	194,426	162,957	2,760	2,139,642	2,139,642	0
20100827	1,772,297	196,213	165,346	2,764	2,136,620	2,136,620	0
20100903	1,831,602	198,254	167,417	2,771	2,200,044	2,200,044	0
20100910	1,847,971	199,310	168,649	2,786	2,218,716	2,218,716	0
20100917	1,860,287	198,839	168,728	2,775	2,230,629	2,230,629	0
20100924	1,801,130	195,451	166,213	2,660	2,165,454	2,165,454	0
20101001	1,802,108	196,497	167,436	2,655	2,168,696	2,168,696	0
20101008	1,811,610	197,018	169,094	2,661	2,180,383	2,180,383	0
20101015	1,808,791	198,179	171,116	2,662	2,180,748	2,180,748	0
20101022	1,841,753	199,019	173,292	2,677	2,216,741	2,216,741	0
20101029	1,837,673	199,810	174,510	2,643	2,214,636	2,214,636	0
20101105	1,825,202	201,591	176,823	2,649	2,206,265	2,206,265	0
20101112	1,808,864	202,721	177,462	2,644	2,191,691	2,191,691	0
20101119	1,784,575	204,428	178,984	2,632	2,170,619	2,170,619	0
20101126	1,770,485	204,364	178,211	2,629	2,155,689	2,155,689	0
20101203	1,766,147	204,963	178,725	2,650	2,152,485	2,152,485	0
20101210	1,769,277	205,987	179,788	2,646	2,157,698	2,157,698	0
20101217	1,767,108	205,634	180,193	2,654	2,155,589	2,155,589	0
20101224	1,704,570	201,612	175,678	2,562	2,084,422	2,084,422	0

(Continued)

EXHIBIT 15.2 (Continued)

Week Ended	Number of Contracts by Seller/Buyer Combination						Differential
	Dealer Dealer	Dealer Nondealer	Nondealer Dealer	Nondealer Nondealer	Subtotal	Reported Total	
20101231	1,701,714	201,191	175,291	2,563	2,080,759	2,080,759	0
20110107	1,843,074	131,034	106,671	2,567	2,083,346	2,083,346	0
20110114	1,848,471	131,801	108,391	2,586	2,091,249	2,091,249	0
20110121	1,850,185	131,923	109,154	2,604	2,093,866	2,093,866	0
20110128	1,707,914	204,737	181,179	2,628	2,096,458	2,096,458	0
20110204	1,713,247	206,161	183,531	2,628	2,105,567	2,105,567	0
20110211	1,719,386	206,554	185,363	2,626	2,113,929	2,113,929	0
20110218	1,715,585	207,507	186,799	2,632	2,112,523	2,112,523	0
20110225	1,718,783	208,780	187,808	2,638	2,118,009	2,118,009	0
20110304	1,724,481	209,898	189,062	2,605	2,126,046	2,126,046	0
20110311	1,719,551	211,080	190,595	2,620	2,123,846	2,123,846	0
20110318	1,729,472	210,623	191,208	2,615	2,133,918	2,133,918	0
20110325	1,696,420	206,646	187,273	2,504	2,092,843	2,092,843	0
20110401	1,701,147	209,337	189,958	2,516	2,102,958	2,102,958	0
20110408	1,698,877	210,121	191,680	2,558	2,103,236	2,103,236	0
20110415	1,691,980	214,881	196,208	2,607	2,105,676	2,105,676	0
20110422	1,767,950	215,877	197,448	2,613	2,183,888	2,183,888	0
20110429	1,768,403	216,455	199,153	2,635	2,186,646	2,186,646	0
20110506	1,767,951	221,487	204,053	2,668	2,196,159	2,196,159	0
20110513	1,775,603	221,418	204,421	2,677	2,204,119	2,204,119	0
20110520	1,776,098	222,736	206,106	2,702	2,207,642	2,207,642	0
20110527	1,776,285	223,582	206,802	2,709	2,209,378	2,209,378	0
20110603	1,777,392	224,485	207,867	2,714	2,212,458	2,212,458	0
20110610	1,781,903	225,460	208,261	2,717	2,218,341	2,218,341	0
20110617	1,795,267	224,688	208,761	2,710	2,231,426	2,231,426	0
20110624	1,734,126	222,474	206,059	2,607	2,165,266	2,165,266	0

20110701	1,737,889	219,851	204,321	2,610	2,164,671	2,164,671	0
20110708	1,733,924	216,168	201,352	2,606	2,154,050	2,154,050	0
20110715	1,741,867	218,505	203,424	2,609	2,166,405	2,166,405	0
20110722	1,743,470	220,040	204,686	2,619	2,170,815	2,170,815	0
20110729	1,747,442	220,074	203,771	2,631	2,173,918	2,173,918	0
20110805	1,748,126	219,130	204,122	2,634	2,174,012	2,174,012	0
20110812	1,757,042	219,698	206,192	2,643	2,185,575	2,185,575	0
20110819	1,747,003	221,249	207,548	2,666	2,178,466	2,178,466	0
20110826	1,722,290	222,923	208,693	2,685	2,156,591	2,156,591	0
20110902	1,727,885	223,004	208,856	2,693	2,162,438	2,162,438	0
20110909	1,731,784	221,123	206,972	2,658	2,162,537	2,162,537	0
20110916	1,743,133	220,175	207,120	2,659	2,173,087	2,173,087	0
20110923	1,716,596	212,569	200,507	2,550	2,132,222	2,132,222	0
20110930	1,727,811	213,012	200,692	2,559	2,144,074	2,144,074	0
20111007	1,732,673	212,648	202,830	2,555	2,150,706	2,150,706	0
20111014	1,735,075	211,921	202,804	2,555	2,152,355	2,152,355	0
20111021	1,739,757	211,990	202,523	2,553	2,156,823	2,156,823	0
20111028	1,742,992	210,870	201,771	2,553	2,158,186	2,158,186	0
20111104	1,746,669	210,765	200,437	2,540	2,160,411	2,160,411	0
20111111	1,741,147	212,293	201,296	2,556	2,157,292	2,157,292	0
20111118	1,744,612	214,899	203,250	2,569	2,165,330	2,165,330	0
20111125	1,754,536	216,581	203,995	2,587	2,177,699	2,177,699	0
20111202	1,758,625	216,953	203,838	2,597	2,182,013	2,182,013	0
20111209	1,742,749	217,496	203,376	2,605	2,166,226	2,166,226	0
20111216	1,747,392	215,732	202,668	2,562	2,168,354	2,168,354	0

Sources: Kamakura Corporation; DTCC.

Let's look at examples from other markets in which the same dealer firms have participated for many decades:

1. *London Interbank Offered Rate Market*: A lawsuit by Charles Schwab, which we discuss in more detail in Chapter 17, accuses 12 firms of colluding to manipulate LIBOR. Nine of the 12 firms are also dealers listed by DTCC in single-name credit default swaps.⁵
2. *Municipal bond market*: A recent article on Reuters.com summarizes the repeated legal violations stemming from collusion in the market for municipal bonds in the United States. Six of the eight firms named are also dealers in single-name CDS that were listed by DTCC in 2010.⁶

Given that the majority of dealers listed by DTCC in single-name credit default swaps are, in the words of the Reuters article, "serial offenders," a sophisticated observer should assume that both traded CDS spreads and quoted spreads are highly likely to have been affected by collusion. Credit risk modelers who use manipulated CDS spreads as inputs to their models will find their fate determined by the CDS dealers posting these CDS quotes.

We now analyze CDS trading volume for the 1,090 reference names for which CDS trades were reported by DTCC during the 77-week period from July 16, 2010, to December 30, 2011. The weekly trade information is from the Section IV reports from DTCC. The data is described this way in the DTCC document "Explanation of Trade Information Warehouse Data" (May 2011):

Section IV (Weekly Transaction Activity) provides weekly activity where market participants were engaging in market risk transfer activity. The transaction types include new trades between two parties, a termination of an existing transaction, or the assignment of an existing transaction to a third party. Section IV excludes transactions which did not result in a change in the market risk position of the market participants, and are not market activity. For example, central counterparty clearing, and portfolio compression both terminate existing transactions and re-book new transactions or amend existing transactions. These transactions still maintain the same risk profile and consequently are not included as "market risk transfer activity."

Our emphasis is not on gross trading volume. As mentioned previously, dealer-dealer volume is 81.68 percent in the single-name credit default swap market and it would be nearly costless for dealers to inflate gross trading volume by trading among themselves. Instead, we focus on end-user trading where at least one of the parties to a trade is not a dealer. Accordingly, we make the following adjustments to the weekly number of trades reported by DTCC for each reference name:

1. We divide each weekly number of trades by five to convert volume to an average daily volume for that week.
2. From that gross daily average number of trades, we classify 81.68 percent of trades as dealer-dealer trades, using the average dealer-dealer share of trades in the DTCC trade warehouse during the 77-week period studied.

3. The remaining 18.32 percent is classified as daily average nondealer volume, the focus of the reporting below.

DAILY NONDEALER TRADING VOLUME FOR 1,090 REFERENCE NAMES

Before reporting on the 1,090 reference names in the DTCC warehouse, it is important to note a very important fact. Kamakura Risk Information Services' default probability service produces daily default probabilities on 30,764 public firms around the world, including 5,222 in the United States. The DTCC weekly reports were designed to report the top 1,000 reference names ranked by daily trading volume. In none of the 77 weeks, however, were there as many as 1,000 reference names reported by DTCC. In essence, the weekly reports contain every reference name in which there was nondealer trading volume in the DTCC. For all other reference names, then, trading volume was zero. This means that more than 29,600 of the public firms in KRIS had zero credit default swap trading volume over the entire 77-week period studied. We now turn to the 1,090 reference names for which there was at least one credit default swap trade during the 77-week period.

We first analyzed daily average nondealer CDS trading volume by calculating the daily average, as described above, for all 1,090 reference names for the entire 77-week period. The distribution of the reference names by nondealer daily average trading volume contains some very surprising results (Exhibit 15.3).

EXHIBIT 15.3 Kamakura Corporation Distribution of Average Nondealer Single-Name CDS Trades per Day for 1,090 Reference Names, Week Ending July 16, 2010, to December 30, 2011

Number of Contracts				
Minimum Trades per Day	Maximum Trades per Day	Number of Reference Names	Percent of All Reference Names	Cumulative Percent of All Reference Names
0.00	0.25	332	30.5%	30.5%
0.25	0.50	211	19.4%	49.8%
0.50	0.75	174	16.0%	65.8%
0.75	1.00	114	10.5%	76.2%
1.00	2.00	189	17.3%	93.6%
2.00	3.00	47	4.3%	97.9%
3.00	4.00	12	1.1%	99.0%
4.00	5.00	6	0.6%	99.5%
5.00	6.00	0	0.0%	99.5%
6.00	7.00	2	0.2%	99.7%
7.00	8.00	0	0.0%	99.7%
8.00	9.00	0	0.0%	99.7%
9.00	10.00	2	0.2%	99.9%
10.00	11.00	0	0.0%	99.9%

(Continued)

EXHIBIT 15.3 (Continued)

Number of Contracts		Number of Reference Names	Percent of All Reference Names	Cumulative Percent of All Reference Names
Minimum Trades per Day	Maximum Trades per Day			
11.00	12.00	1	0.1%	100.0%
12.00	13.00	0	0.0%	100.0%
13.00	14.00	0	0.0%	100.0%
14.00		0	0.0%	100.0%
Total		1090	100.0%	

Sources: Kamakura Corporation; DTCC.

Among the surprising conclusions are these facts:

- 76.2 percent of the reference names had one or fewer nondealer contracts traded per day.
- 93.6 percent of reference names had two or fewer nondealer contracts traded per day.
- 99.0 percent of reference names had four or fewer nondealer contracts traded per day.
- Only five reference names averaged more than five nondealer contracts traded per day.
- Only one reference name averaged more than 10 nondealer contracts traded per day.

These trading volumes are miniscule to say the least, and it is astonishing that financial journalists have not made the appropriate disclaimers when quoting “CDS prices” (many of which are quotes, not trades) for specific reference names.

In order to better understand the data, we next analyzed each week of data for all 1,090 reference names. In aggregate, there should be $77 \times 1,090 = 83,930$ “reference name-weeks” of data reported, but there was no data reported on 14,632 occasions because there were no trades in that week for that reference name. Dividing these 14,632 observations by 1,090 reference names shows that, on average over the 77-week period, there were 13 weeks for each reference name in which there were zero CDS trades. The daily average nondealer trade volumes for these 83,930 observations have the distribution shown in Exhibit 15.4.

Again, this analysis allows us to draw some more surprising conclusions about all weekly trading in single-name CDS for all 1,090 reference names, a total of 83,930 observations:

- 17.43 percent of the observations showed zero nondealer daily average CDS trading.
- 77.48 percent of the observations showed less than one daily average nondealer CDS contract traded.

EXHIBIT 15.4 Kamakura Corporation Distribution of Nondealer Weekly Average Single-Name CDS Trades per Day for All Reference Names and All Weeks, Week Ending July 16, 2010, to December 30, 2011

Number of Contracts		Number of Observations	Percent of All Observations	Cumulative Percent of All Observations
Minimum Trades per Day	Maximum Trades per Day			
0	0	14,632	17.4%	17.43%
0.01	0.25	22,427	26.7%	44.15%
0.25	0.5	13,679	16.3%	60.45%
0.5	0.75	8,480	10.1%	70.56%
0.75	1	5,813	6.9%	77.48%
1	2	11,046	13.2%	90.64%
2	3	4,112	4.9%	95.54%
3	4	1,782	2.1%	97.67%
4	5	853	1.0%	98.68%
5	6	412	0.5%	99.17%
6	7	245	0.3%	99.47%
7	8	133	0.2%	99.62%
8	9	71	0.1%	99.71%
9	10	58	0.1%	99.78%
10	11	34	0.0%	99.82%
11	12	30	0.0%	99.85%
12	13	20	0.0%	99.88%
13	14	25	0.0%	99.91%
14	15	14	0.0%	99.92%
15	16	8	0.0%	99.93%
16	17	4	0.0%	99.94%
17	18	7	0.0%	99.95%
18	19	8	0.0%	99.96%
19	20	6	0.0%	99.96%
20	21	9	0.0%	99.97%
21	22	6	0.0%	99.98%
22	23	0	0.0%	99.98%
23	24	0	0.0%	99.98%
24	25	2	0.0%	99.98%
25	26	1	0.0%	99.98%
26	27	1	0.0%	99.99%
27	28	2	0.0%	99.99%
28	29	1	0.0%	99.99%
29	30	1	0.0%	99.99%
30	31	1	0.0%	99.99%
31	32	2	0.0%	99.99%
32	33	1	0.0%	100.00%
33	34	0	0.0%	100.00%

(Continued)

EXHIBIT 15.4 (Continued)

Number of Contracts		Number of Observations	Percent of All Observations	Cumulative Percent of All Observations
Minimum Trades per Day	Maximum Trades per Day			
34	35	1	0.0%	100.00%
35	36	0	0.0%	100.00%
36	37	1	0.0%	100.00%
37	38	0	0.0%	100.00%
38	39	0	0.0%	100.00%
39	40	1	0.0%	100.00%
40	41	0	0.0%	100.00%
41	42	0	0.0%	100.00%
42	43	0	0.0%	100.00%
43	44	0	0.0%	100.00%
44	45	0	0.0%	100.00%
45	46	0	0.0%	100.00%
46	47	1	0.0%	100.00%
47		0	0.0%	100.00%
Total		83,930	100.0%	

Sources: Kamakura Corporation; DTCC.

- 90.64 percent of the observations showed less than two daily average nondealer CDS contracts traded.
- 99.17 percent of the observations showed less than six daily average nondealer CDS contracts traded.
- There were only 10 observations of the 83,930 where more than 28 daily average nondealer CDS contracts traded.

We now turn to the trading volume of single-name credit default swaps for major financial institutions.

On December 28, 2011, Scott Richard published an article in the *Financial Times* advocating the use of credit default swap spreads in setting of deposit insurance rates. This section explains why such an idea, attractive in theory, would be extremely dangerous in practice. Richardson argued in his article “A Market-Based Plan to Regulate Banks” that the credit default swap market should be used to price deposit insurance. He states, “How can the FDIC determine the correct price for its insurance? The answer is to use the large and active market in bank insurance via credit default swaps.”

Unfortunately, there are several very practical reasons why the credit default swap market cannot be used as Richard suggests.

1. There is no “large and active market” in bank credit default swaps. Using data reported by the DTCC during the 77 weeks ending December 30, 2011, there were credit default swaps traded on only 13 reference names among U.S. banking firms:

Ally Financial
American Express
Bank of America
Capital One Bank
Capital One Financial Corporation
Citigroup
Citigroup Japan
Istar Financial
JPMorgan Chase
MetLife
Morgan Stanley
Goldman Sachs
Wells Fargo

These 13 reference names represent 11 consolidated corporations, four of which would not be considered banking firms by most observers prior to the recent credit crisis (American Express, Goldman Sachs, MetLife, and Morgan Stanley). During the 77 weeks of data on all live trades in the DTCC credit default swap trade warehouse, there were no other trades on any other of the 6,453 banks insured by the FDIC in the United States as of March 31, 2011.

2. Even on the firms listed above, there was an average of only 2.3 nondealer credit default swap trades per day during the 77 weeks ended December 30, 2011. None of the banks listed above averaged more than five nondealer trades per day over the 77-week period studied.
3. Six of the 11 firms listed are in a conflict of interest position as major dealers in the credit default swap market: Bank of America, Citigroup, JPMorgan Chase, Morgan Stanley, Goldman Sachs, and Wells Fargo. Dealer–dealer trades made up 81.68 percent of live trades in the DTCC over the 77-week period studied. The dealers would be setting deposit insurance rates for themselves if Richard’s proposals were adopted.
4. Credit default swaps, to the extent they trade, represent supply, demand, probability of default, and the probability of rescue, not the probability of failure alone. Kamakura Corporation’s September 29, 2009 blog on www.kamakuraco.com, “Comparing Default Probabilities and Credit Default Swap Quotes: Insights from the Examples of FNMA and Citigroup,” showed how credit default swap quotes for both firms were far below their probability of failure because senior debt holders correctly anticipated that the depositors of Citibank and the senior debt holders of Citigroup and FNMA would be rescued by the U.S. government. This rendered the credit default swap spreads on both failing firms’ severe underestimates of their probability of failure.
5. Because of the thinness of nondealer trading and the very small number of dealing firms worldwide, there is high risk of collusion. Indeed, on April 29, 2011, the European Union launched an investigation of possible collusion among dealers in the credit default swap market.

For these reasons, Richard’s proposal to use credit default swaps to price deposit insurance would not work given current market conditions and the lack of competition in the market for credit default swaps.

CREDIT DEFAULT SWAP TRADING VOLUME IN MUNICIPALS AND SUB-SOVEREIGNS

We now look at weekly credit default swap trading volume for sub-sovereigns and municipals among those 1,090 reference names. We find, unfortunately, that (in the words of Gertrude Stein) “there is no there there.”

Of the 1,090 reference names for which DTCC reported credit default swap trades in the 77-week period, only nine were sub-sovereigns of any type:

Hong Kong Special Administrative Region
State of California
State of Florida
State of Illinois
State of New Jersey
State of New York
State of Texas
The City of New York
The Commonwealth of Massachusetts

Eight of the nine reference names were in the United States and seven of the nine reference names were U.S. states. The only cities on which credit default swaps had any trades in the 77 weeks ended December 30, 2011, were the City of New York and Hong Kong. We can summarize the trading volume in these nine reference names as follows:

- There were $9 \times 77 = 693$ weekly observations, but there were only 309 weeks in which trades were reported. That means in 55.4 percent of the 77 weeks, on average, there would be no trades on these nine reference names.
- The average number of nondealer trades per day on all nine reference names over the 77-week period was 0.11 trades per day.
- The median number of nondealer trades per day over the 77-week period among the nine reference names was 0.06 trades per day.
- The highest average number of nondealer trades per day for the full 77-week period among the nine reference names was 0.36 trades per day.
- The lowest average number of nondealer trades per day for the full 77-week period among the nine reference names was 0.00 (rounding to two decimal places).
- The highest number of gross trades in one week was 84, which is the equivalent of 16.8 gross trades per day and 3.1 nondealer trades per day.

The obvious conclusion is that there is minimal trading volume in credit default swaps on sub-sovereign and municipal reference names, and that any pricing indications should be regarded with a high degree of skepticism.

CREDIT DEFAULT SWAP TRADING VOLUME IN SOVEREIGN CREDITS

We now turn to weekly credit default swap trading volume for sovereigns among those 1,090 reference names. We find that, in a small subset of sovereign names, there is regular trading, but in modest volume.

Of the 1,090 reference names for which DTCC reported credit default swap trades in the 77-week period, only 57 were sovereigns:

Abu Dhabi
Arab Republic of Egypt
Argentine Republic
Bolivarian Republic of Venezuela
Commonwealth of Australia
Czech Republic
Dubai
Federal Republic of Germany
Federative Republic of Brazil
French Republic
Hellenic Republic
Ireland
Japan
Kingdom of Belgium
Kingdom of Denmark
Kingdom of Norway
Kingdom of Saudi Arabia
Kingdom of Spain
Kingdom of Sweden
Kingdom of Thailand
Kingdom of The Netherlands
Lebanese Republic
Malaysia
New Zealand
People's Republic of China
Portuguese Republic
Republic of Austria
Republic of Bulgaria
Republic of Chile
Republic of Colombia
Republic of Croatia

Republic of Estonia
Republic of Finland
Republic of Hungary
Republic of Iceland
Republic of Indonesia
Republic of Italy
Republic of Kazakhstan
Republic of Korea
Republic of Latvia
Republic of Lithuania
Republic of Panama
Republic of Peru
Republic of Poland
Republic of Slovenia
Republic of South Africa
Republic of The Philippines
Republic of Turkey
Romania
Russian Federation
Slovak Republic
Socialist Republic of Vietnam
State of Israel
State of Qatar
United Kingdom of Great Britain and Northern Ireland
United Mexican States
United States of America

No credit default swap trades were reported in the 77 weeks ending December 30, 2011, for any other sovereign credit.

We first analyze the 77-week averages for the 57 sovereigns for which CDS trading volume was greater than zero during the 77 weeks ending December 31, 2011. The daily average nondealer trading volume, calculated as described previously, was distributed as shown in Exhibit 15.5.

The conclusions that can be drawn from this table are summarized here:

- 52.6 percent of the 57 sovereigns averaged less than one nondealer CDS trade per day.
- 93.0 percent of the 57 sovereigns averaged less than five nondealer CDS trades per day.
- The remaining four sovereigns had the highest levels of nondealer CDS trades per day of the 1,090 reference names reported by DTCC, but those figures were

EXHIBIT 15.5 Kamakura Corporation Distribution of Average Nondealer Single-Name CDS Trades per Day for 57 Sovereigns, Week Ending July 16, 2010, to December 30, 2011

Number of Contracts				
Minimum Trades per Day	Maximum Trades per Day	Number of Reference Names	Percent of All Reference Names	Cumulative Percent of All Reference Names
0.00	0.25	10	17.5%	17.5%
0.25	0.50	7	12.3%	29.8%
0.50	0.75	6	10.5%	40.4%
0.75	1.00	7	12.3%	52.6%
1.00	2.00	5	8.8%	61.4%
2.00	3.00	6	10.5%	71.9%
3.00	4.00	7	12.3%	84.2%
4.00	5.00	5	8.8%	93.0%
5.00	6.00	0	0.0%	93.0%
6.00	7.00	1	1.8%	94.7%
7.00	8.00	0	0.0%	94.7%
8.00	9.00	0	0.0%	94.7%
9.00	10.00	2	3.5%	98.2%
10.00	11.00	0	0.0%	98.2%
11.00	12.00	1	1.8%	100.0%
12.00	13.00	0	0.0%	100.0%
13.00	14.00	0	0.0%	100.0%
14.00		0	0.0%	100.0%
Total		57	100.0%	

Sources: Kamakura Corporation; DTCC.

only 11.28, 9.52, 9.32, and 6.61 trades per day. The sovereigns with the three highest average number of nondealer trades per day were all members of the European Union.

- The average 77-week daily average number of nondealer trades per day for the 57 sovereigns was 2.08 trades per day.
- The median 77-week daily average number of nondealer trades per day was 0.92 trades per day.
- The U.S. 77-week average nondealer trades per day was approximately equal to the median (see disclosure restrictions imposed by DTCC discussed previously).

We conclude that trading volume for the most active sovereigns is higher than it is for the most active corporations, which is only logical given that such sovereigns issue more debt than the most active corporations. The correlation between trading volume and debt outstanding is weak, however, with the United States and Japan well down on the ranking of trade volume even though those two countries have the largest amount of debt outstanding.

We now analyze all 77 weeks of data, not just the average over that period, for all 57 sovereigns for which DTCC reported nonzero trade volume. There were 4,389 ($= 57 \times 77$) observations on CDS trading volume for these 57 sovereigns, and there were no trades during 265 observations, 6 percent of the total. The distribution of nondealer trades per day over these 4,389 observations is summarized in the chart in Exhibit 15.6.

EXHIBIT 15.6 Kamakura Corporation Distribution of Nondealer Weekly Average Single-Name CDS Trades per Day for 57 Sovereigns and All Weeks, Week Ending July 16, 2010, to December 20, 2011

Number of Contracts		Number of Observations	Percent of All Observations	Cumulative Percent of All Observations
Minimum Trades per Day	Maximum Trades per Day			
0	0	265	6.0%	6.04%
0.01	0.25	931	21.2%	27.25%
0.25	0.5	560	12.8%	40.01%
0.5	0.75	331	7.5%	47.55%
0.75	1	244	5.6%	53.11%
1	2	662	15.1%	68.19%
2	3	452	10.3%	78.49%
3	4	297	6.8%	85.26%
4	5	191	4.4%	89.61%
5	6	127	2.9%	92.50%
6	7	73	1.7%	94.17%
7	8	50	1.1%	95.31%
8	9	32	0.7%	96.04%
9	10	30	0.7%	96.72%
10	11	24	0.5%	97.27%
11	12	19	0.4%	97.70%
12	13	12	0.3%	97.97%
13	14	19	0.4%	98.41%
14	15	13	0.3%	98.70%
15	16	5	0.1%	98.82%
16	17	4	0.1%	98.91%
17	18	5	0.1%	99.02%
18	19	7	0.2%	99.18%
19	20	6	0.1%	99.32%
20	21	9	0.2%	99.52%
21	22	6	0.1%	99.66%
22	23	0	0.0%	99.66%
23	24	0	0.0%	99.66%
24	25	1	0.0%	99.68%
25	26	1	0.0%	99.70%
26	27	1	0.0%	99.73%
27	28	2	0.0%	99.77%

EXHIBIT 15.6 (Continued)

Number of Contracts		Number of Observations	Percent of All Observations	Cumulative Percent of All Observations
Minimum Trades per Day	Maximum Trades per Day			
28	29	1	0.0%	99.79%
29	30	1	0.0%	99.82%
30	31	1	0.0%	99.84%
31	32	2	0.0%	99.89%
32	33	1	0.0%	99.91%
33	34	0	0.0%	99.91%
34	35	1	0.0%	99.93%
35	36	0	0.0%	99.93%
36	37	1	0.0%	99.95%
37	38	0	0.0%	99.95%
38	39	0	0.0%	99.95%
39	40	1	0.0%	99.98%
40	41	0	0.0%	99.98%
41	42	0	0.0%	99.98%
42	43	0	0.0%	99.98%
43	44	0	0.0%	99.98%
44	45	0	0.0%	99.98%
45	46	0	0.0%	99.98%
46	47	1	0.0%	100.00%
47		0	0.0%	100.00%
Total		4,389	100.0%	

Sources: Kamakura Corporation; DTCC.

One can draw the following conclusions over 4,389 weekly observations:

- 53.11 percent of the observations showed one nondealer trade per day or less.
- 92.5 percent of the observations showed six nondealer trades per day or less.
- 96.72 percent of the observations showed 10 nondealer trades per day or less.
- Only 0.32 percent of the observations were for more than 25 nondealer trades per day.
- The highest volume week featured 1,271 gross trades per week, 254.2 gross trades per day, and 46.6 average nondealer trades per day.
- Just 10 sovereigns account for 52 percent of the total trading volume in credit default swaps over the 77-week period ending December 30, 2011.

IMPLICATIONS OF CDS TRADING VOLUME DATA

Although, in theory, the credit default swap market should provide major price discovery benefits, trading is very thin and highly concentrated among a small

number of dealers who are under investigation for manipulation of LIBOR and the U.S. municipal bond market. We conclude that use of CDS data in default modeling would make the modeler a potential accomplice in creating the appearance of legitimacy for a market that, at least at the present time, is “not there.”

For that reason, in the next chapter, we introduce state-of-the-art credit models that use the much more transparent market for common stock as key market inputs. We turn to that task now.

NOTES

1. Quote from a young trader at Barings Securities after his first year on the trading floor, circa 1998.
2. Specifically, the Kamakura Risk Information Services “Jarrow Chava” default probabilities, KDP-jc, provided by Kamakura Corporation.
3. Presentation to the North American Asset and Liability Management Association, September 2001, Toronto.
4. See the websites for each of the rating agencies for more detail on these concepts.
5. See *Charles Schwab Bank, N.A., etc. vs. Bank of America Corporation, etc.*, United States District Court, Northern District of California, San Francisco Division, Case no. CV114187, www.lieffcabraser.com/media/pnc/4/media.904.pdf.
6. See Cate Long, “Muniland’s Serial Offenders,” Reuters, January 3, 2012, <http://blogs.reuters.com/muniland/2012/01/03/munilands-serial-offenders/>.

Reduced Form Credit Models and Credit Model Testing

For decades, researchers have looked for a more comprehensive, all-encompassing framework for risk management than the legacy approaches to credit risk that we review in Chapter 18, our “credit risk museum.” Sophisticated financial practitioners wanted a synthesis between models of default for major corporations and the credit scoring technology that has been used for decades to rate the riskiness of retail and small business borrowers. The framework long sought by bankers, the reduced form modeling framework, has blossomed dramatically in the years since the original Jarrow and Turnbull (1995) paper, “Pricing Derivatives on Financial Securities Subject to Credit Risk.” On December 10, 2003, the Federal Deposit Insurance Corporation (FDIC) announced that it was adopting the reduced form modeling framework for its new Loss Distribution Model (see the FDIC website for details of the model) after many years of studying the Merton model (which we describe in Chapter 18) as an alternative.

In this chapter, we introduce the basic concepts of the reduced form modeling technology and the reasons that it has gained such strong popularity.¹

THE JARROW-TURNBULL MODEL

Many researchers trace the first reduced form model to Jarrow and Turnbull (1995), and we start our review of the reduced form model with that paper. Jarrow and Turnbull’s objective was to outline a modeling approach that would allow valuation, pricing, and hedging of derivatives where:

- The asset underlying the derivative security might default.
- The writer of the derivative security might default.

This kind of compound credit risk is essential to properly valuing the Jarrow-Merton put option on the assets of a financial institution, which we discussed as the best single integrated measure of the total risk of all types that a financial institution has. It is also essential for a more basic objective: measuring risk accurately regardless of the risk measure that one chooses to employ. This compound approach to default is also at the very core of the tsunami of regulatory initiatives (Basel II, Basel II.5,

Basel III, Solvency II, and so on) over the past decade because of the prevalence of compound default risks in typical loan products, guarantees, and in derivatives markets. Analyzing compound credit risk is a substantial departure from the popular Black-Scholes approach that is all-pervasive in the marketplace. The Black-Scholes model assumes:

- The company whose common stock underlies an option on the common stock will not default.
- The firm who has written the call option and has an obligation to pay if the option is “in the money” will not default.

Ignoring this kind of compound default risk can lead to very large problems, like the potential loss of \$2 billion that JPMorgan Chase had as of this writing due to its large credit default swap position made public on May 8, 2012. The legacy approaches to credit risk that we discuss in Chapter 18 were not flexible enough or accurate enough to provide the kind of framework that Jarrow and Turnbull were seeking.²

The Jarrow-Turnbull Framework

Jarrow and Turnbull assume a frictionless economy where two classes of zero-coupon bonds trade.³ The first class of bonds is risk-free, zero-coupon bonds $P(t, T)$, where t is the current time and T is the maturity date. They also introduce the idea of a “money market fund” $B(t)$, which represents the compounded proceeds of \$1 invested in the shortest-term, risk-free asset starting at time zero. We used the money market fund concept in modeling the risk-free yield curve using the Heath, Jarrow, and Morton (HJM) approach in Chapters 6 through 9.

The second class of bonds trading is the risky zero-coupon bonds of XYZ company, which potentially can default. These zero-coupon bonds have a bond price $v(t, T)$. For expositional purposes, Jarrow and Turnbull assume that interest rates and the probability of default are statistically independent.⁴ They assume that the time of bankruptcy T^* for firm XYZ is exponentially distributed with parameter λ . They assume that there is a Poisson bankruptcy process driven by the parameter μ , which they assume to be a positive constant less than 1.⁵ Jarrow and Turnbull assume that the risk-free interest rates are random and consistent with the extended Merton term structure model (the continuous time version of the Ho and Lee model) that we studied in Chapter 13, which presented special cases of the HJM framework. Jarrow and Turnbull also assume that in the event of bankruptcy, \$1 of XYZ zero-coupon debt pays a recovery rate $\delta < 1$ on the scheduled maturity date of the zero (not at the time of default). When bankruptcy occurs, the value of XYZ zero-coupon debt drops to $\delta P(T^*, T)$ because the payment at maturity drops from \$1 to δ , but the payment of δ is now certain and risk free. The value of δ can be different for debt of different seniorities. For expositional purposes, Jarrow and Turnbull assume the recovery rate is constant. We relax that assumption in the next section.

Jarrow and Turnbull show that the value of a defaultable zero-coupon bond in this model is

$$v(t, T) = [e^{-\lambda\mu(T-t)} + (1 - e^{-\lambda\mu(T-t)})\delta]P(t, T)$$

We can interpret this formula in a practical way by noting a few points. The product $(\lambda\mu)$ is the continuous probability of default at any instant in time. The quantity

$$1 - e^{-\lambda\mu(T-t)}$$

is the (risk-neutral) probability that XYZ company defaults between time t and maturity of the zero-coupon bond T . The quantity

$$e^{-\lambda\mu(T-t)}$$

is the (risk-neutral) probability that company XYZ does *not* default between time t and maturity of the zero-coupon bond at time T . With this in mind, we can see that Jarrow and Turnbull have reached a very neat and, with hindsight, simple conclusion that the value of a zero-coupon bond in the Jarrow-Turnbull model is the probability-weighted present value of what the bond holder receives if we know there is no default (in which case we have the equivalent of a risk-free bond $P(t, T)$) and what we receive if we know for certain that default will occur, the recovery rate times the risk-free, zero-coupon bond price $\delta P(t, T)$.

This clever result depends on the assumption that default and interest rates are not correlated. We relax this assumption in the next section. Jarrow and Turnbull then go on to derive closed-form solutions for a number of important securities prices:

- European call options on XYZ Company's defaultable zero-coupon bond
- Defaultable European call option on XYZ's defaultable zero-coupon bond (where the writer of the call option might default, assuming its default probability is independent of both interest rates and the bankruptcy of XYZ)
- European call option on XYZ Company's defaultable common stock

The Jarrow-Turnbull modeling framework is a very important extension of the risk management analytics we have discussed so far because it provides an integrated and flexible framework for analyzing interest rate risk and credit risk in an integrated way, the central objective of this volume. We have one task left, and that is to discuss the extension of the approach above to interest rate-dependent and macro-economic factor-dependent default.

THE JARROW MODEL

The Jarrow (1999–2001) model is an important extension of the Jarrow-Turnbull (1995) model, which was perhaps the first reduced form model to experience widespread industry acceptance. Many other researchers, most notably Duffie and Singleton (1999), have done very important work on reduced form models, but in the remainder of this chapter we will concentrate on Jarrow's work because of the continuity with the Jarrow-Turnbull model we discussed in the previous section.

In the early days of the credit derivatives markets, dealers logically turned to the legacy Merton model (1974) for pricing purposes. Dealers soon realized that the credit default swap market prices were consistently different from those indicated by

the Merton model for many of the reasons we discuss in Chapter 18. Dealers were under considerable pressure from both management, risk managers, and external auditors to adopt a modeling technology with better pricing accuracy. Dealers' instinct for self-preservation was an added incentive. The Jarrow-Turnbull model was the first model that allowed the matching of market prices and provided a rational economic basis for the evolution of market prices of everything from corporate debt to credit derivatives.

The original Jarrow-Turnbull model assumes that default is random but that default probabilities are nonrandom, time-dependent functions. The Jarrow model extends the Jarrow-Turnbull model in a large number of significant respects. First, default probabilities are assumed to be random, with explicit dependence on random interest rates and an arbitrary number of lognormally distributed risk factors. These lognormally distributed risk factors are extremely important in modeling correlated default, a critical factor in the credit crisis of 2006–2011 (driven by home price declines) and the two “lost decades” in Japan (1990–2010) in which commercial real estate prices, home prices, and stock prices all suffered very large declines.

Jarrow goes beyond the Jarrow-Turnbull framework by explicitly incorporating a liquidity factor that affects the prices of bonds, but not the prices of common stock. This liquidity factor can be random and is different for each bond issuer. In addition, the liquidity parameter can be a function of the same macro risk drivers that determine the default intensity. We discuss this issue in detail in Chapter 17. Although we use the term “liquidity factor,” we intend the term to describe everything affecting bond prices above and beyond potential losses due to default, the intersection of supply and demand for the credit of each borrower. This can include things like bid-offered price spreads, the impact on market prices of the sale of large holdings of a specific bond issue (the Long-Term Capital Management phenomenon), general risk aversion in the fixed income market, size of the firm, location of the firm, and so on. We give specific examples in Chapter 17.

The Jarrow model also allows for the calculation of the implied recovery given default δ_i . This parameter, which can be random and driven by other risk factors, is defined by Jarrow as the fractional recovery of δ where ν is the value of risky debt a fraction of an instant before bankruptcy at time τ and the subscript i denotes the seniority level of the debt.

$$\text{Recovery} = \delta_i(\tau)\nu(\tau-, T : i)$$

This is a different recovery structure than Jarrow-Turnbull, where the recovery rate was specified as a fraction of the principal value to be received at maturity T .

Duffie and Singleton (1999) were the first to use this specification for the recovery rate. They recognized that the traditional bankers' thinking on the recovery rate, expressed as percentage of principal, made the mathematical derivation of bond and credit derivatives prices more difficult. They also recognized that thinking of recovery as a percentage of principal was too narrow—what is the principal on an “in the money” interest rate cap? On a foreign exchange option? On an undrawn loan commitment? Expressing recovery as a percentage of value one instant before bankruptcy is both more powerful and more general. We can easily convert it back to the traditional “recovery rate as a percentage of principal” to maximize user-friendliness of the concept.

The hazard rate, or default intensity in the Jarrow model is given by a linear combination of three terms:

$$\lambda(t) = \lambda_0 + \lambda_1 r(t) + \lambda_2 Z(t)$$

This is a much simpler expression than the Merton credit model discussed in Chapter 18, and it is much richer than the structure of Jarrow-Turnbull. The first term, λ_0 , can be made time dependent, extending the credit model exactly like we extended term structure models in Chapters 6 through 9 and Chapter 13. The term $Z(t)$ is the “shock” term with mean zero and standard deviation of one, which creates random movements in the macro factors, such as home prices during the 2006 to 2011 credit crisis, which then drives default for that particular company.⁶ Movements in this macro factor are generally written in this form:

$$dM(t) = M(t)[r(t)dt + \sigma_m dZ(t)]$$

The change in the macro factor is proportional to its value and drives upward at the random risk-free interest rate $r(t)$, subject to random shocks from changes in $Z(t)$, multiplied by the volatility of the macro factor, σ_m . While the default intensity in the Jarrow model as we have written it describes interest rates and one macro factor as drivers of default, it is easy to extend the model to an arbitrary number of macro factors. This is because the sum of a linear combination of normally distributed variables such as $Z(t)$ is still normally distributed. In this volume, we use the one-factor notation for expositional purposes. Please note that the incorporation of these macroeconomic drivers of default mean that the default probabilities (and default intensities) are correlated due to their dependence on common factors. This is a very important step forward from Jarrow-Turnbull where the default probabilities of the two companies in that example were uncorrelated with each other and with interest rates.

In the Jarrow model, interest rates are random, but the term structure model chosen is a more realistic model than the extended Merton model. The term structure model assumed for the riskless rate of interest is a special case of the HJM framework, the Hull-White or extended Vasicek model, which we reviewed in Chapter 13. For practical use, however, it is critical that the risk-free interest rate assumptions be completely consistent with the current observable yield curve (and any other interest rate derivatives on that yield curve) as discussed in Chapter 14. The extended Vasicek model contains an extension from the original Vasicek model that allows the theoretical yield curve to exactly match the actual yield curve. In this model, as we discussed in Chapter 13, the short-term rate of interest $r(t)$ drifts over time in a way consistent with interest rate cycles, subject to random shocks from the Brownian motion $W(t)$, which like $Z(t)$ has mean zero and standard deviation of 1:⁷

$$dr(t) = a[\bar{r}(t) - r(t)]dt + \sigma_r dW(t)$$

ZERO-COUPON BOND PRICES IN THE JARROW MODEL

Jarrow shows that the price of a zero-coupon bond can be derived that explicitly incorporates the interactions between the default intensity, interest rates, and the

macro factor. The results are simple enough to be modeled in spreadsheet software, especially the case where the coefficients of the default intensity formula for short-term interest rates and the macro factors are zero. We discuss practical implementation of this version of the Jarrow model below.

The Jarrow Model and the Issue of Liquidity in the Bond Market

One of the many virtues of reduced form models is the ability to fit parameters to the model from a wide variety of securities prices. Of course the type of data most directly related to the interests of lenders is bond prices or credit derivatives prices. Since a large company like IBM has only one type of common stock outstanding but could have 10, 20, or 30 bond issues outstanding, it is logical to expect that there is less liquidity in the bond market than in the market for common stock. Jarrow makes an explicit adjustment for this in his model by introducing a very general formulation for the impact of liquidity on bond market prices. This is a very powerful feature of the model that reduces the question of bond market liquidity to a scientific question (how best to fit a parameter that captures the liquidity impact) rather than a philosophical or religious question of whether bond prices can be used to calculate the parameters of a credit model. Note that there can be either a liquidity discount or liquidity premium on the bonds of a given issuer, and there is no implicit assumption by Jarrow that the liquidity discount is constant—it can be random and is flexible in its specifications, consistent with the research we discuss in Chapter 17.

Jarrow's original (1999–2001) model shows how the Jarrow credit model can be fit to bond prices and equity prices simultaneously. It can also be fit to bond prices alone, using the techniques of Chapter 17. Jarrow and Yildirim (2002) show how the model can be fit to credit derivatives prices and Jarrow and Chava (2004) fit the model to historical data on defaults.

The Jarrow-Merton Put Option as a Risk Index and a Practical Hedge

The Jarrow model is much better suited to hedging credit risk on a portfolio level than the Merton model we discuss in Chapter 18 because the link between the (N) macro factor(s) M and the default intensity is explicitly incorporated in the model. Take the example of Exxon, whose probability of default is driven by interest rates and oil prices, among other things. If $M(t)$ is the macro factor oil prices, it can be shown that the size of the hedge that needs to be bought or sold to hedge one dollar of risky zero-coupon debt with market value v under the Jarrow model is given by

$$\partial v_i(t, T : i) / \partial M(t) = -[\partial \gamma_i(t, T) / \partial M(t) + \lambda_2(1 - \delta_i)(T - t) / \sigma_m M(t)] v_l(t, T : i).$$

The variable v is the value of risky zero-coupon debt and gamma is the liquidity discount function representing the illiquidities often observed in the debt market. There are similar formulas in the Jarrow model for hedging coupon-bearing bonds, defaultable caps, floors, credit derivatives, and so on.

Van Deventer and Imai (2003) show that the steps in hedging the macro factor risk for any portfolio are identical to the steps that a trader of options has been taking for 30 years (hedging his net position with a long or short position in the common stock underlying the options):

- Calculate the change in the value (including the impact of interest rates on default) of all retail credits with respect to interest rates.
- Calculate the change in the value (including the impact of interest rates on default) of all small business credits with respect to interest rates.
- Calculate the change in the value (including the impact of interest rates on default) of all major corporate credits with respect to interest rates.
- Calculate the change in the value (including the impact of interest rates on default) of all bonds, derivatives, and other instruments
- Add these delta amounts together.
- The result is the global portfolio delta, on a default-adjusted basis, of interest rates for the entire portfolio.
- Choose the position in interest rate derivatives with the opposite delta.
- This eliminates interest rate risk from the portfolio on a default-adjusted basis.

We can replicate this process for any macro factor that impacts default, such as home prices, exchange rates, stock price indices, oil prices, the value of class A office buildings in the central business district of key cities, and the like. We should know how to do this using logistic regression to estimate default probabilities in a manner consistent with the Jarrow framework.

Most importantly:

- We can measure the default-adjusted transaction level and portfolio risk exposure with respect to each macro factor.
- We can set exposure limits on the default-adjusted transaction level and portfolio risk exposure with respect to each macro factor.
- We know how much of a hedge would eliminate some or all of this risk.

It can be shown that all other risk, other than that driven by macro factors, can be diversified away. This hedging and diversification capability is a powerful step forward from where most financial institutions found themselves at the close of the twentieth century.

FITTING THE JARROW MODEL TO BOND PRICES, CREDIT DERIVATIVES PRICES, AND HISTORICAL DEFAULT DATABASES

One of the many virtues of the reduced form modeling approach is its rich array of closed form solutions. These include closed form solutions for zero-coupon bond prices, coupon bearing bond prices, credit default swap prices, first to default swap prices, and many others. This means that there are many alternatives for fitting the Jarrow model parameters. We discuss a few of them in this section.

Fitting the Jarrow Model to Debt Prices

In Chapter 17, we take the reader through an example of how to fit the Jarrow model to observable prices of risky debt. Third-party bond price vendors offer databases with daily or weekly prices on more than 100,000 different risky debt issues. Prices can be obtained for almost every issuer with a rating from the legacy rating agencies.

Fitting to Current Price Data and Historical Price Data

If one is willing to accept the hypothesis that the impact of interest rates and macroeconomic factors on the default probability of a particular company is fairly stable over time, one can improve the quality of parameter estimates by adding historical data to the estimation procedure. Let's assume that the analyst has five years of daily price data for, say, IBM bonds. The following procedure takes advantage of this data:

1. Assume a structure of the liquidity function in the Jarrow model. For example, a liquidity function that is effectively linear in years to maturity could be adopted. Say the assumed structure contains two parameters a and b .
2. Collect the $N(i)$ observable bond prices on each observation date i , for the whole sample $i = 1, M$ observations.
3. Using a nonlinear equation solver over the entire sample, find the combination of a , b , λ_0 , λ_1 , and λ_2 that minimizes the sum of squared errors on the observable bonds. This parallels the interest rate parameter fitting we did in Chapter 14 for the risk-free term structure.
4. λ_1 and λ_2 are the coefficients of interest rates and the macro factors on default. We hold these parameters constant over the whole sample and then go back over the sample and refit the parameters a , b , and λ_0 for each day in the sample, allowing them to change daily.
5. Eliminate any remaining bond pricing errors on a given day by making λ_0 a time (to maturity) dependent variable. The original value of λ_0 for that day, as above, is useful to obtain because it is the value that minimizes the amount of extension.

Fitting the Jarrow Model to Credit Derivatives Prices

The procedure for fitting the Jarrow model to credit derivatives prices is exactly the same as fitting the Jarrow model to bond prices, except of course that the closed-form solution that is being used is the credit derivatives formula in the Jarrow model, not the bond price formula. Because most counterparties with traded credit derivatives will have only one to three maturities at which prices are visible, using a longer history of credit derivatives prices as described in the second method above will be even more helpful than it is in the bond market.

Fitting the Jarrow Model to a Historical Database of Defaults

In the first part of this chapter, the continuous time default intensity was the object of study, just like the continuous time interest rates we studied in Chapters 6 through 13. In both the interest rate arena and the credit arena, most applied risk analysis is done using discrete rates, either interest rates or default probabilities. For the remainder of this chapter, the default probabilities we discuss all have a discrete time to maturity. In this section, we discuss how to fit a default probability function driven by macroeconomic factors to a discrete historical data set of observations on public firms at various points in time, some of which constitute a "defaulting observation" and some of which constitute a nondefaulting observation.

Fitting the reduced form model to historical bankruptcy data was explored extensively by Shumway (2001); Jarrow and Chava (2004); Campbell, Hilscher, and Szilagyi (2008, 2011); and Bharath and Shumway (2008). Each of these authors used public firm data, but the identical approach is also applied to the construction of default models for all types of counterparties: retail borrowers, small business borrowers, large unlisted borrowers, municipal borrowers, and sovereign borrowers. This fitting process involves advanced hazard rate modeling, also called *logistic regression*, an extension of the credit scoring technology that has been used for retail and small business default probability estimation for many decades. The distinction between credit scoring and hazard rate modeling is very important. Typically, credit scoring analysis uses the most recent information available on a set of nondefaulting counterparties and the last set of information available on the defaulting counterparties. Hazard rate modeling makes use of all information about all counterparties at every available time period, in order to explain, for example, why Lehman Brothers defaulted in September 2008 instead of at some earlier time.

We illustrate the creation and testing of discrete default probability models using data from the Kamakura Risk Information Services (KRIS) public firm default model, first offered in 2002. The Kamakura Corporation's KRIS default probability database includes 1.76 million monthly observations on over 8,000 companies for North America alone. In the North American sample, 2,046 companies have failed since 1990. A company failure is denoted as a "1" in the bankruptcy variable, and the lack of failure is denoted as a "0." Since the entire history of a company is included, Lehman Brothers will have many observations that are zeros and only one observation that is a 1, the September 2008 observation (recorded as of the end of August 2008). Exhibit 16.1 is a sample of a simple logistic regression database where the explanatory variables are market leverage, three-month stock return volatility, the net income to total assets ratio, and the unemployment rate (percent). The variable that we are seeking to explain is the probability of default (the default flag). The default flag is set to zero if the firm does not default in the next month and it is set to one if the company does default over that time period. While the following table includes entries for only Lehman Brothers, all of the companies are modeled together and a set of explanatory variables are gathered that best predict default.

The result is a set of coefficients (the alpha and betas) of the logistic regression formula, which predicts the unannualized probability of default at a given time t $P[t]$ as a function of n explanatory variables:

$$P[t] = \frac{1}{1 + e^{-\alpha - \sum_{i=1}^n \beta_i X_i}}$$

In addition to interest rates and macro factors that are part of the theory of the Jarrow model, other explanatory variables, which are likely to be important, are added as appropriate. These include financial ratios, other macroeconomic variables, multiple inputs from the stock price and its history for each company, multiple inputs from the comprehensive stock market index for the relevant country, company size, industry variables, seasonality, and so on. As we will see in the model-testing section of this chapter, the result is a very powerful model highly capable of discriminating among companies by both long-term and short-term risk. Logistic

EXHIBIT 16.1 Logistic Regression Default Database Sample Entries for Lehman Brothers Holdings

Ticker	Date	Company Name	Default Flag	Market Leverage	3 Month Stock Volatility	Net Income to Total Assets	Unemployment Rate (Percent)
lehmq	1/31/2008	Lehman Brothers Holdings	0	0.9495	0.6966	0.0013	5.4
lehmq	2/29/2008	Lehman Brothers Holdings	0	0.9610	0.6244	0.0013	5.2
lehmq	3/31/2008	Lehman Brothers Holdings	0	0.9699	1.5194	0.0006	5.2
lehmq	4/30/2008	Lehman Brothers Holdings	0	0.9689	1.5422	0.0006	4.8
lehmq	5/30/2008	Lehman Brothers Holdings	0	0.9740	1.5566	0.0006	5.2
lehmq	6/30/2008	Lehman Brothers Holdings	0	0.9858	1.0291	−0.0043	5.7
lehmq	7/31/2008	Lehman Brothers Holdings	0	0.9846	1.4784	−0.0043	6.0
lehmq	8/29/2008	Lehman Brothers Holdings	1	0.9857	1.6163	−0.0043	6.1

regression is the tool that is used to estimate the parameters or coefficients multiplying each input variable. The appendix of this chapter and Jarrow and Chava (2004) discuss the use of default probabilities derived from a discrete database in a continuous time model, just like interest rates smoothed from discrete observations are used in a continuous time term structure model. For a detailed review of logistic regression and its applications, see van Deventer and Imai (2003) and Hosmer and Lemeshow (2000).

Kamakura Risk Information Services produces more than a simple one-month default probability. KRIS provides default probabilities for a full 10-year horizon, constructed from default probabilities projected for the next 120 months. We call the unannualized monthly default probability for the first month $P[1]$, for the second month (conditional on surviving the first month) $P[2]$, and the n th month (conditional on surviving the first $N - 1$ months) $P[N]$. How are these default probabilities combined to form a term structure of default probabilities? (See the appendix of this chapter for details on changing the periodicity of default probabilities for background.)

If one has a time horizon of N months, the probability that a company will *not* go bankrupt during those N months is

$$Q[N] = (1 - P[1])(1 - P[2])(1 - P[3]) \dots (1 - P[N - 1])(1 - P[N])$$

The probability that the company *will* go bankrupt by the end of the n th month is

$$P[N] = 1 - Q[N]$$

This is the cumulative probability of default for an N -month time horizon. Since credit spreads and credit default swaps are all discussed or displayed on an annualized basis, we can annualize the cumulative probability of default for an N -month horizon. The annualized default probability for an N -month horizon is

$$A[N] = 1 - (1 - P[N])^{12/N}$$

It is these annualized default probabilities that are displayed on the KRIS service. Exhibit 16.2 shows the annualized KRIS term structure of default for Bank of America on February 27, 2009, which was a dangerous period in the 2006 to 2011 credit crisis. The annualized three-month default probability for Bank of America on this date was 24.27 percent.

Exhibit 16.3 shows the same default risk for Bank of America, this time displayed on a cumulative basis. The graph shows that the 10-year cumulative failure risk of the bank was 37.89 percent over a 10-year period, even after an injection of \$45 billion into the organization under the U.S. government's Troubled Asset Relief Program (TARP).

Another firm that figured prominently in the 2006–2011 credit crisis was the aggressive mortgage lender Washington Mutual (“WMI Holdings” in the graphs in Exhibit 16.4). Washington Mutual's one-year default probability hit 56.45 percent on the eve of September 25, 2008, the day that the Federal Deposit Insurance



EXHIBIT 16.2 Bank of America Corporation




EXHIBIT 16.3 Bank of America Corporation

Corporation (FDIC) seized its main banking subsidiary and sold it in an assisted rescue to JPMorgan Chase.

On that day, the cumulative default risk of Washington Mutual was 80.45 percent for a time horizon of 10 years as shown in Exhibit 16.5.

Fitting the Jarrow Model to Retail, Small Business, and Governmental Counterparties

While the Jarrow-Turnbull and Jarrow models have been discussed in the context of a publicly listed corporate counterparty, there is no reason to restrict the models’ application to corporations alone. They can be applied with equal effectiveness to




EXHIBIT 16.4 WMI Holdings Corp.




EXHIBIT 16.5 WMI Holdings Corp.

retail client default modeling, small business default modeling, and to the modeling of defaults by sovereign entities. The only difference is the inputs to the logistic regression model like that which we discussed in the preceding sections. Other than that, all of the analysis is identical.

For example, consider sovereigns that issue debt but for whom there are no observable prices (say, Kazakhstan). Historical estimation using logistic regression is the normal fallback estimation procedure. KRIS offers a sovereign default service developed in exactly this way and launched in the fall of 2008.

In many countries, there is a very large amount of bond issuance by cities, towns, municipal special purpose entities, and other government entities besides the national

government itself. The prices of the bonds issued by these entities are sometimes heavily affected by special tax features, such as the exemption of municipal bonds from federal income tax in the United States. Jarrow default probabilities can be fitted to this group as well, both from bond prices (if they are high quality) and by logistic regression from a historical default database.

CORRELATIONS IN DEFAULT PROBABILITIES

In the same way, the Jarrow models fitted to historical default databases illustrate the parallels between default probabilities and interest rates:

- Default probabilities have a term structure just like interest rates.
- Changes in interest rates and macroeconomic factors make the term structure of default probabilities rise and fall over time.
- Like interest rates, short-term default probabilities rise and fall by much more than long-term default probabilities, which may be the effective average of default probabilities over many credit cycles.
- The difference between a strong credit and a weak credit is not their short-term default probability when times are good. In good times, both will have a low probability of default. The difference can be seen most clearly when times are bad.

Macro factors cause correlated movements in the short-term default probabilities for both strong and weak credits.

A few examples will illustrate this phenomenon. The first example, Exhibit 16.6, shows the one-year default risk of Bank of America and Citigroup from 2008 to 2010.

The simple five-year monthly correlation in their one-year default probabilities for the period ending in June 2012 of their Jarrow default probabilities is 91.87 percent. That contrasts sharply with the 40.50 percent correlation in the one-year default probabilities of Apple and Bank of America over the same five-year period. Exhibit 16.7 shows the one- and five-year default probabilities for Apple from 2007 to 2012. While Apple's default risk was a small fraction of that of Bank of America, Apple also experienced a rise in one-year default risk in late 2008 and early 2009 that was very similar to that of Bank of America. This correlated shift in risks, with a large




EXHIBIT 16.6 BAC (JC5) vs. C (JC5)

**EXHIBIT 16.7** Apple Inc.

difference in magnitude, is typical of the change in default risk for good credits versus bad credits in a credit portfolio.

THE JARROW AND JARROW-TURNBULL MODELS: A SUMMARY

In this section, we have introduced the Jarrow-Turnbull and Jarrow models of risky debt. They offer powerful closed-form solutions for practical issues that are central to the integrated treatment of interest rate risk, market risk, liquidity risk, and credit risk. We have shown how to fit the Jarrow model to historical default databases. In Chapter 17, we show how to fit a credit model to observable spreads in the bond markets and derivatives markets. We have shown how the reduced form credit modeling approach, with different explanatory variables, can be applied to:

- Retail counterparties
- Small business counterparties
- Municipalities
- Sovereigns

as well as to public firms.

Finally, we have shown through some very practical examples that the Jarrow reduced form models shed dramatic light on both the term structure of default probabilities for any type of counterparty and on the correlation of default probabilities for any two counterparties. This kind of correlated risk is what drives the “fat tails” in losses that endanger financial institutions. An understanding of this issue is essential to both the integrated measure of risk via the Jarrow-Merton put option and to the hedging of portfolio-level credit risk with macro hedges using stock price index futures and derivatives on key macroeconomic factors.

In the next section, we discuss in detail how to measure the effectiveness of default models, both in terms of absolute accuracy and relative accuracy compared to other models.

TESTS OF CREDIT MODELS USING HISTORICAL DATA

This section is one of the key sections in this book. A debate about credit models without a factual basis for deciding which model is best is just as much of a beauty contest as the debate about smoothing methods that we discussed in Chapter 5. In this section, we replace the beauty contest with science. We define criteria for “best model,” as we did in Chapter 5. We then measure candidate models by these criteria to reach our conclusions about which are best.

This approach is very consistent with the Basel II and Basel III capital regulations that have emerged in recent years from the Basel Committee on Banking Supervision. The Basel capital regulations place very heavy emphasis on the role of credit risk in measuring the safety and soundness of financial institutions. Solvency II for the insurance industry has a similar emphasis on model accuracy. For that reason, it is extremely critical from both a shareholder value point of view and a regulatory point of view that financial institutions have a clear understanding of the accuracy and performance of the credit models employed for that purpose.

Without knowledge of the accuracy of the models:

- How can regulators know whether the financial institution’s capital is too high or too low?
- How can shareholders know whether the reserve for loan losses is adequate or not?
- How can the CEO certify the financial statements of the company as accurate?
- How can the institution have confidence that a hedge of its portfolio credit risk will work?
- How can the institution have confidence in its own survival?

These are critical questions, to say the least, and they are the major motivators of the emphasis on credit model testing in this section. The Basel Accords’ contribution to credit risk modeling has been twofold. First, the Accords stress the need to apply credit risk models to the entire balance sheet, not just to credit-related derivatives such as credit default swaps. Second, the Accords stress the need for quantitative proof that the models perform well. Performance measurement is critical both across credit risk categories and over time; that is, through the peaks and valleys of credit cycles.⁸

The result of these key developments has been a surge of interest in reduced form models and increased implementation of these models around the world in financial institutions of all types, with applications to retail clients, private companies, listed companies, and sovereigns. A prominent example is the December 10, 2003, announcement by the FDIC that it has adopted a reduced form model-related Loss Distribution Model that we mentioned in Chapter 15.⁹ The FDIC announcement is just another confirmation of a wholesale movement away from the older structural models of credit risk like the original Merton (1974) risky debt model and more

recent extensions like Shimko, Tejima, and van Deventer (1993), which we discuss in Chapter 18. We do something atypical in this book. Instead of explaining all of the candidate models first, like we did with yield curve smoothing, and then measuring “who’s best,” we do something different. As the reader will see in the following pages, the credit model test results in both the academic and commercial public domain are so compelling that we reverse the order. After explaining the reduced form approach, but not legacy approaches, we then measure accuracy. Having seen “who’s best” in this way justifies our relegation of older approaches to the “legacy” collection of credit models in Chapter 18, much like the legacy term structure models that we reviewed in Chapter 13.

Given the plethora of credit risk models now available, how and why did the FDIC conclude that the reduced form model offered the most promise? More generally, how should a banker, fund manager, or insurance executive seeking to use credit risk models in credit risk management assess their relative performance and then implement them in practice? The answer to this difficult question is the subject of this section.

AN INTRODUCTION TO CREDIT MODEL TESTING

One of the dilemmas facing the Basel Committee on Banking Supervision as it compiled the New Capital Basel Capital Accord (May 31, 2001) was the (then) paucity of data on credit model performance. In the following decade, however, and the years immediately prior to the publication of the New Capital Accord, an avalanche of studies on relative performance of credit models has reached the public domain. We refer the series reader to Shumway (2001); van Deventer and Imai (2003); Jarrow and Chava (2004); Bharath and Shumway (2008); and Campbell, Hilscher, and Szilagyi (2008 and 2011) for some notable examples. Almost all of these papers were in wide circulation for four or five years prior to their formal publication.

Credit model testing is not a trivial exercise, for a number of reasons. The first is the considerable expense and expertise needed, both in terms of finance and computer science, to assemble the data and computer coding to provide a consistent methodology for testing credit models in a way that the original Basel II accord (and good corporate governance) requires. Van Deventer and Imai (2003) note that the Basel II requires that banks must prove to their regulatory supervisors that the credit models they use perform “consistently and meaningfully.”¹⁰ Typically, the only institutions who have the capability to assemble these kinds of databases are extremely large financial institutions and commercial vendors of default probabilities. Prior to the commercialization of default probabilities by KMV (acquired by Moody’s in 2002) in the early 1990s, studies of default were based on a very small number of defaulting observations. Falkenstein and Boral (2000) cite academic papers by Altman (1968) (33 defaults), Altman (1977) (53 defaults), and Blum (1974) (115 defaults) to illustrate the relatively small sample sizes used to draw inferences about bankruptcy probabilities prior to 1998. By way of contrast, Kamakura Corporation’s commercial default database for public firms includes more than 2,000 failed company observations, and its research database, which spans a longer period, contains more than 2,400 failed companies for North America alone as of 2012.

For major financial institutions that have incurred the expense of a large default database, the results of model testing are highly valuable and represent a significant competitive advantage over other financial institutions who do not have the results of credit model performance tests. For example, there is a large community of arbitrage investors actively trading against users of the Merton default probabilities when the arbitrage investors perceive the signals sent by the Merton model to be incorrect.

Among the vendor community, the majority of vendors offer a single default probability model. This presents a dilemma for potential consumers of commercial default probabilities. A vendor of a single type of credit model has two reasons not to publish quantitative tests of performance. The first reason is that the tests may prove that the model is inferior and ultimately may adversely affect the vendor's commercial prospects. Perhaps for this reason, most vendors require clients to sign license agreements that forbid the clients from publicizing any results of the vendor's model performance. The second reason is more subtle. Even if quantitative performance tests are good, the fact that the vendor offers only one model means that the vendor's tests will be perceived by many as biased in favor of the model that the vendor offers.

Four former employees of Moody's Investors Service have set the standard for quantitative model test disclosure in a series of papers: Andrew Boral, Eric Falkenstein, Sean Keenan, and Jorge Sobehart. The authors respect the important contributions of Boral, Falkenstein, Keenan, and Sobehart to the integrity of the default probability generation and testing process. Much of the work in this section can be attributed to their influence and the influence of the authors mentioned previously.

The need for such tests is reflected in the frequently heard comments of default probability users who display a naiveté with respect to credit models that triggered the massive losses from trading in collateralized debt obligations during the 2006–2011 credit crisis. We present some samples in the next section that illustrate the need for better understanding of credit model testing.

MISUNDERSTANDINGS ABOUT CREDIT MODEL TESTING

A commonly heard comment on credit model performance goes like this:

I like Model A because it showed a better early warning of the default of Companies X, Y, and Z.

Many users of default probabilities make two critical mistakes in assessing default probability model performance. They choose a very small sample (in this case three companies) to assess model performance and use a naïve criteria for good performance. Assessing model performance on only three companies or 50 or even 100 in a universe of 8,000 to 10,000 in the total universe of U.S. public firms needlessly exposes the potential user to: (1) an incorrect conclusion just because of the noise in the small sample; and (2) to the risk of data mining by the default probability vendor, who (like a magician doing card tricks) can steer the banker to the 3 or 50 or 100 examples that show the model in the best light. A test of the whole sample eliminates these risks. Analysts should demand this of both internal models and models purchased from third parties.

The second problem this banker's quote has is the performance criteria. The implications of his comment are twofold:

- I can ignore all false predictions of default and give them zero weight in my decision.
- If model A has higher default probabilities than model B on a troubled credit, then model A must be better than model B.

Both of these implications should be grounds for a failing grade by banking and insurance supervisors and the audit committee of the firm's board of directors. The first comment, ignoring all false positives, is sometimes justified by saying "I sold Company A's bonds when its default probabilities hit 20 percent and saved my bank from a loss of \$1.7 million, and I don't care if other companies which don't default have 20 percent default probabilities because I would never buy a bond with a 20 percent default probability anyway." Why, then, did the bank have the bond of Company A in its portfolio? And what about the bonds that were sold when default probabilities rose, only to have the bank miss out on gains in the bond's price that occurred after the sale? Without knowledge of the gains avoided, as well as the losses avoided, the banker has shown a striking "selection bias" in favor of the model he is currently using. This selection bias will result in any model being judged good by a true believer. We give some examples next.

The second implication exposes the banker and the vendor to a temptation that can be detected by the tests we discuss here. The vendor can make any model show better early warning than any other model simply by raising the default probabilities. If the vendor of model B wants to win this banker's business, all he has to do is multiply all of his default probabilities by six or add an arbitrary scale factor to make his default probabilities higher than model A. The banker making this quote would not be able to detect this moral hazard because he does not use the testing regime mentioned below. The authors can't resist pointing out that using the summer temperature in Fahrenheit as a credit model will outpredict any credit model by this criterion if the company defaults in July or August in the Northern Hemisphere.

Eric Falkenstein and Andrew Boral (2000) of Moody's Investors Service address the issue of model calibration directly:

Some vendors have been known to generate very high default rates, and we would suggest the following test to assess those predictions. First, take a set of historical data and group it into 50 equally populated buckets (using percentile breakpoints of 2%, 4%, . . . 100%). Then consider the mean default prediction on the x-axis with the actual, subsequent bad rate on the y-axis. More often than not, models will have a relation that is somewhat less than 45% (i.e., slope < 1), especially at these very high risk groupings. This implies that the model purports more power than it actually has, and more importantly, it is miscalibrated and should be adjusted. (46)

We present the results of the Falkenstein and Boral test later in this section. We also present a second type of test to detect this kind of bias in credit modeling. If a model has a bias to levels higher than actual default rates, it is inappropriate for

Basel II, Basel III, or Solvency II use because it will be inaccurate for pricing, hedging, valuation, and portfolio loss simulation.

Another quotation illustrates a similar point of view that is inconsistent with Basel II compliance in credit modeling:

That credit model vendor is very popular because they have correctly predicted 10,000 of the last 10,500 small business defaults.

Again, this comment ignores false predictions of default and assigns zero costs to false predictions of default. If any banker truly had that orientation, the Basel II credit supervision process will root them out with a vengeance because the authors hereby propose a credit model at zero cost that outperforms the commercial model referred to previously:

100 percent Accurate Prediction of Small Business Defaults

The default probability for all small businesses is 100 percent.

This naïve model correctly predicts 10,500 of the last 10,500 defaults. It is free in the sense that assigning a 100 percent default probability to everyone requires no expense or third-party vendor, since anyone can say the default probability for everyone is 100 percent. And, like the banker quoted above, it is consistent with a zero weight on the prediction of false positives. When pressed, most financial institutions admit that false positives are important. As one major financial institution comments, one model “correctly predicted 1,000 of the last three defaults.”

Once this is admitted, there is a reasonable basis for testing credit models.

THE TWO COMPONENTS OF CREDIT MODEL PERFORMANCE

The performance of models designed to predict a yes or no—zero or one—event, like bankruptcy, has been the focus of mathematical statistics for more than 50 years.¹¹ These performance tests, however, have only been applied to default probability modeling in recent years. Van Deventer and Imai (2003) provide an overview of these standard statistical tests for default probability modeling. Related papers are those cited at the beginning of this section: Shumway (2001); Jarrow and Chava (2004); Bharath and Shumway (2008); and Campbell, Szilagyi, and Hilscher (2008, 2011). Jarrow, van Deventer, and Wang (2003) and van Deventer and Imai (2003) apply an alternative hedging approach to measuring credit model performance. We discuss each of these alternative testing procedures in turn.

Basel II and Basel III require that financial institutions have the capability to test credit model performance and internal ratings to ensure that they consistently and meaningfully measure credit risk. There are two principal measures of credit risk model performance. The first is a measure of the correctness of the ordinal ranking of the companies by riskiness. For this measure, we use the so-called “receiver operating characteristics” (ROC) accuracy ratio, whose calculation is reviewed briefly in the next section. The second is a measure of the consistency of the predicted default probability with the actual default probability, which Falkenstein and Boral (2000)

call “calibration.” This test is necessary to ensure the accuracy of the model for pricing, hedging, valuation, and portfolio simulation. Just as important, it is necessary to detect a tendency for a model to bias default probabilities to the high side as Falkenstein and Boral note, which overstates the predictive power of a model by the naïve criteria of the first quote in the introduction.

Measuring Ordinal Ranking of Companies by Credit Risk

The standard statistic for measuring the ordinal ranking of companies by credit riskiness is the ROC accuracy ratio. The ROC accuracy ratio is closely related to, but different than, the cumulative accuracy profiles used by Jorge Sobehart and colleagues (2000), formerly at Moody’s Investors Service and now at Citigroup, in numerous publications in recent years.

The ROC curve was originally developed in order to measure the signal-to-noise ratio in radio receivers. The ROC curve has become increasingly popular as a measure of model performance in fields ranging from medicine to finance. It is typically used to measure the performance of a model that is used to predict which of two states will occur (sick or not sick, defaulted or not defaulted, etc.). Van Deventer and Imai (2003) go into extensive detail on the meaning and derivation of the ROC accuracy ratio, which is a quantitative measure of model performance.

In short, the ROC accuracy ratio is derived in the following way:

- Calculate the theoretical default probability for the entire universe of companies in a historical database that includes both defaulted and nondefaulted companies.
- Form all possible pairs of companies such that the pair includes one defaulted “company” and one nondefaulted “company.” To be very precise, one pair would be the December 2001 defaulted observation for Enron and the October 1987 observation for General Motors, which did not default in that month. Another pair would include defaulted Enron, December 2001, and nondefaulted Enron, November 2001, and so on.
- If the default probability technology correctly rates the defaulted company as more risky, we award one point to the pair.
- If the default probability technology results in a tie, we give half a point.
- If the default probability technology is incorrect, we give zero points.
- We then add up all the points for all of the pairs, and divide by the number of pairs.¹²

The results are intuitive and extremely clear on model rankings:

- A perfect model scores 1.00 or 100 percent accuracy, ranking every single one of the defaulting companies as more risky than every nondefaulting company
- A worthless model scores 0.50 or 50 percent, because this is a score that could be achieved by flipping a coin
- A score in the 90 percent range is extremely good
- A score in the 80 percent range is very good, and so on

Van Deventer and Imai provide worked examples to illustrate the application of the ROC accuracy ratio technique. The ROC accuracy ratio can be equivalently summarized in one sentence: *It is the average percentile ranking of the defaulting companies in the universe of nondefaulting observations.*

We turn now to the Kamakura Corporation database, which we shall use to apply the ROC accuracy ratio technology. Results will be summarized for the North American database, which includes all listed companies for which the explanatory variables were available, with no exceptions for any other reason.

The Predictive ROC Accuracy Ratio: Techniques and Results

The KRIS database is a monthly database. Other researchers have used annual data, including the work of Sobehart and colleagues noted previously. Annual data has been used in the past because of the researchers' interest in long-term bankruptcy prediction. One of the purposes of this chapter is to show how monthly data can be used for exactly the same purpose, and to show how accuracy changes as the prediction period grows longer.

The basic Jarrow-Chava (2004) model uses logistic regression technology to combine equity market and accounting data for default probability prediction consistent with the reduced form modeling approach. Jarrow and Chava provide a framework that shows these default probability estimates are consistent with the best practice reduced form credit models of Jarrow (1999, 2001) and Duffie and Singleton (1999). We turn now to the ROC accuracy ratios for the Jarrow-Chava reduced form modeling approach.

THE PREDICTIVE CAPABILITY OF THE JARROW-CHAVA REDUCED FORM MODEL DEFAULT PROBABILITIES

This section shows that the technology for measuring credit model predictive capability is transparent and straightforward. We show that the reduced form modeling approach has a high degree of accuracy even when predicting from a distant horizon.

The standard calculation for the ROC accuracy ratio is always based on the default measure in the period of default (which could be a month if using monthly data, or a year if using yearly data). In order to address Basel and financial institutions management questions about the "early warning" capability of default probability models, it is useful to know the accuracy of credit models at various time horizons, not just in the period of default.

We know that, in their distant past, companies that defaulted were at some point the same as other companies and their default probabilities were probably no different from those of otherwise similar companies. In other words, if we took their default probability in that benign era and used that to calculate the ROC accuracy ratio, we would expect to see an accuracy ratio of 0.50 or 50 percent, because at that point in their history the companies are just "average" and so are their default probabilities.

Therefore, our objective is to see how far in advance of default the ROC accuracy ratio for the defaulting companies rises from the 50 percent level on its way to achieving the standard ROC accuracy ratio in the month of default.

Measuring the Predictive ROC Accuracy Ratio

We calculate the time series of predictive ROC accuracy ratios in the following way:

- We again form all possible pairs of one defaulted company and one non-defaulted company. Instead of defining the default date and associated default probability for Enron as December 2001, however, we step back one month prior to default and use the November time period for Enron as the default date and the November 2001 default probability as the “one-month predictive” default probability.
- We calculate the ROC accuracy ratio on this basis.
- We lengthen the prediction period by one more month again.
- We recalculate which default probability we will use in the ROC comparison for each of the defaulted companies (in the second iteration, we will use the October 2001 default probability for Enron, not November).
- We repeat the process until we have studied the predictive period of interest and have calculated the predictive ROC accuracy ratio for each date.

REDUCED FORM MODEL VS. MERTON MODEL PERFORMANCE

We now supplement the results of Bharath, Campbell, Chava, Hilscher, Jarrow, Shumway, and Szilagyi with results from the KRIS testing results for version 5.0, which uses monthly data spanning the period from 1990 to 2008. We compare the accuracy of the ordinal ranking of credit risk for three credit models distributed by the Kamakura Corporation:

Reduced Form Credit Model

KDP-jc, Kamakura Default Probabilities from an advanced form of the Jarrow-Chava (2004) approach. Version 5 of the model is denoted KDP-jc5.

Structural Credit Models

KDP-ms, Kamakura Default Probabilities using the “best” Merton econometric approach. This model is labeled KDP-ms5 in version 5.

Hybrid Credit Models

KDP-jm, Kamakura Default Probabilities combining the Jarrow and Merton approaches in a hybrid model within the logistic regression framework. A transformation of the KDP-ms Merton default probability is added as an additional explanatory variable to the Jarrow-Chava variables in KDP-jc to form KDP-jm, denoted KDP-jm5 for version 5.

The chart in Exhibit 16.8 shows the ROC accuracy ratios for version 5 of all three of the models based on monthly data from 1990 to 2008, which includes the most dramatic failures in the 2007 to 2010 credit crisis. For comparison, we also plot the accuracy of an older reduced form model, version 4, labeled KDP-JC4, and a pure theory implementation of the Merton model. The chart shows clearly that the




EXHIBIT 16.8 ROC Accuracy Ratio by Model and Months 120 Forward Monthly Database, 1990–2008

accuracy of the credit risk ranking of both the KDP-jc5 Jarrow-Chava reduced form model and the hybrid KDP-jm5 model are consistently superior to the KDP-ms Merton model for all forecasting horizons for 120 months prior to default. In the final month before default (“month 1” to model credit modelers), the ROC accuracy ratios are as follows: 95.10 for KDP-JC5 (Jarrow-Chava reduced form model version 5), 94.93 for KDP-JM5 (Jarrow-Merton hybrid model version 5), 94.69 for KDP-JC4 (Jarrow-Chava reduced form model version 4), 92.09 for KDP-MS5 (Merton structural model, econometric version 5), and 82.59 for the Merton theoretical model.¹³ (The Merton model implementation strategies are explained in Chapter 18.) The differential of more than 12 percentage points versus the Merton theoretical model comes as no surprise as it is consistent with the early work of Jorge Sobehart and colleagues and Bohn et al. (2005), as well as the many papers by Chava and Jarrow, Bharath and Shumway, and Campbell et al. Please note that the ROC accuracy ratio for month N is the accuracy of predicting default in month N , conditional on the firm not defaulting in months one through $N - 1$. This is a much more difficult prediction than predicting if the firm defaulted in months 1 through N . Accuracy ratios for that cumulative default prediction would be much higher for all of the models.

With this kind of precision in ranking of models by accuracy for any user-defined forecasting horizon, there should be no question about the superiority of model A vs. model B when it comes to the Basel II and III requirements for credit model testing.

CONSISTENCY OF ESTIMATED AND ACTUAL DEFAULTS

Falkenstein and Boral (2000) correctly emphasize the need to do more than measure the correctness of the ordinal ranking of companies by riskiness. One needs to determine whether a model is correctly “calibrated” (in the words of Falkenstein and Boral), that is, whether the model has default probabilities that are biased high or low. As noted in the introduction’s example, a naïve user of credit models can be convinced a model has superior performance just because it gives higher default probabilities for some subset of a sample. A test of consistency between actual and expected defaults is needed to see whether this difference in default probability levels is consistent with actual default experience or just an ad hoc adjustment or noise. Anecdotal evidence suggests that default models for public firms in common use are biased high by a factor of three or four times actual default levels.

RECENT RESULTS FROM NORTH AMERICA

In this and subsequent sections, we report test results for the KRIS version 5.0 default probabilities on a sample of 1.76 million observations from North America from 1990 to 2008. We plot actual defaults, month by month, over this period versus the expected default produced by the model. Exhibit 16.9 shows the results for the Jarrow-Chava KDP-JC5 version 5 model.

If we run the regression

$$\text{Actual defaults} = A + B(\text{Expected defaults})$$

as suggested by van Deventer and Wang (2003), the adjusted R -squared of the regression is a good quantitative measure of the consistency between actual and expected defaults. In Exhibit 16.10, we see that the figure for the KDP-JC5 Jarrow Chava model is 74.30 percent, well above the Merton model, which explains less than 43 percent of the month to month variation in actual defaults.

THE FALKENSTEIN AND BORAL TEST

Eric Falkenstein and Andrew Boral (2000), formerly of Moody’s Investors Service, suggested another consistency test between actual and expected defaults. They suggest the following test:

- Order the universe of all default probability observations from lowest to highest default probability level.
- Create N “buckets” of these observations with an equal number of observations in each bucket.
- Measure the default probability boundaries that define the low end and high end of each bucket.
- Measure the actual rate of default in each bucket.

- The actual default rate should lie between the lower and upper boundary of the bucket for most of the N buckets.

Falkenstein and Boral propose this test because the implied “early warning” of a credit model can be artificially “increased” by multiplying the default probabilities by some arbitrary ratio like 5 for the reasons discussed previously. If the model was properly calibrated before this adjustment, the result of multiplication by five will result in an expected level of defaults that is five times higher than actual defaults,




EXHIBIT 16.9 Kamakura Risk Information Services, Actual vs. Expected Public Firm Defaults, 1990–2008

EXHIBIT 16.10 Actual Defaults Measured as a Function of Predicted Defaults by the Model
Actual Defaults = $A + B * \text{Expected Defaults}$

Version	Model	Adjusted R ² Data Set
Version 5.0	KDP-jc5	74.30% Version 5.0, January
	Jarrow-Chava Model	1990-December 2008, Monthly
Version 5.0	KDP-jm5	73.88% Version 5.0, January
	Jarrow-Merton Hybrid	1990-December 2008, Monthly
Version 5.0	KDP-ms5	42.38% Version 5.0, January
	Merton Logistic Model	1990-December 2008, Monthly

which would be unacceptable both to management and to financial institutions regulators under the New Capital Accords and final version of Basel II, Basel III, and Solvency II.

If there is an accidental or intentional bias in default probabilities, the adjusted R^2 test in the prior section will detect date-related or credit-cycle-related biases. If the bias is related to the absolute level of the default probabilities, the Falkenstein and Boral test applies. The graph in Exhibit 16.11 applies the Falkenstein and Boral test using 100 time buckets for version 5.0 of Kamakura Corporation's KDP-JC5 Jarrow-Chava reduced form model. The top line defines the upper default probability in each of the 100 buckets. Each bucket has 17,642 observations. The lower line defines the lower default probability bound of each bucket. The kinked middle line is the actual frequency of defaults. The graph shows a very high degree of consistency in the KDP-jc modeling of actual versus expected defaults. The graph in Exhibit 16.12 displays the same results on a logarithmic scale.

The Falkenstein and Boral test does not work well for the lower default probability buckets due to small sample size issues having to do with the discreteness of a default event. As indicated in this graph, it is impossible for the actual default frequency to fall between the upper and lower bounds of the bucket because 0 defaults out of 17,642 is below the lower bound and 1 default out of 17,642 is above the upper bound. Consequently, this test is most useful for the intermediate and higher buckets. The Jarrow-Chava version 5 model shows good consistency, when measured by default probability level, between actual and expected defaults.




EXHIBIT 16.11 Falkenstein and Boral Test, Kamakura Risk Information Services, Actual Default Rate vs. Upper and Lower Bounds in 100 Percentile Buckets




EXHIBIT 16.12 Falkenstein and Boral Test, Kamakura Risk Information Services, KDP-jc5, Jarrow-Chava Verion 5 Log Scale

Another test suggested by Eric Falkenstein¹⁴ is to plot the distribution of actual defaults by the percentile of the default probabilities. The 1st percentile (the lowest 17,642 of the 1.76 million default probabilities in the sample) should have the lowest number of defaults and the 100th percentile should have the highest as shown in Exhibit 16.13. Its graph shows these test results for version 5 of the Merton model and Jarrow-Chava model. The Jarrow-Chava model shows a better performance because the actual default experience is more heavily concentrated in the highest default probability percentile buckets. In contrast, percentage of defaults captured by the Merton model in the 99th percentile is nearly 20 percentage points lower than the 42 percent captured by the Jarrow-Chava model. The difference in discrimination between the two models can be measured quantitatively using a Chi-squared test as suggested in van Deventer and Imai (2003).

PERFORMANCE OF CREDIT MODELS VS. NAÏVE MODELS OF RISK

Knowing the absolute level of the accuracy of a credit model is extremely useful, but it is just as important to know the relative accuracy of a credit model versus naïve models of credit risk. A typical “naïve” model would be one that uses only one financial ratio or equity market statistic to predict default. It is well-known among industry experts that many financial ratios outperform the Merton model in the ordinal ranking of riskiness. This result, however, comes as a surprise to many academics and financial services regulators, so we present representative results here.¹⁵




EXHIBIT 16.13 Distribution of Defaults by Percentile Rank of Default Probability for KRIS Version 5 Default Probability Models KDP-jc5 and KDP-ms5

ROC Accuracy Ratios for Merton Model Theoretical Version vs. Selected Naïve Models

As mentioned previously, the pure theory implementation of the Merton model has an ROC accuracy ratio over a one-month horizon of 0.8259. Among the explanatory variables used in the Jarrow-Chava version 5 model, seven variables have a higher stand-alone accuracy ratio than the Merton model. Three of those ratios are summarized here:

0.8803	Three-month standard deviation of stock price returns
0.8687	Excess return over S&P 500 equity index for the last year
0.8522	Percentile rank of firm’s stock price on given date

This observation that the Merton model underperforms common financial ratios has economic and political implications for the practical use of credit models:

Can management of a financial institution or bank regulators approve a model whose performance is inferior to a financial ratio that management and bank regulators would not approve as a legitimate modeling approach on a stand-alone basis?

Of course, the answer should be no. This finding is another factor supporting the use of reduced form models, which by construction outperform any naïve model depending on a single financial ratio.¹⁶

Tests of Credit Models Using Market Data

The first part of this chapter introduced reduced form credit models and a testing regime to show how well such models perform relative to the Merton model and other legacy credit models that we review in more detail in Chapter 18. Our interest in the models and how well they work is more than academic. We are looking for a comprehensive framework for enterprise risk management that creates true risk-adjusted shareholder value. We want to be able to accurately value the Jarrow-Merton put option on the value of the firm's assets that is the best comprehensive measure of integrated credit risk, market risk, liquidity risk, and interest rate risk. If risk is outside of the safety zone discussed in earlier chapters, we want to know the answer to a critical question: What is the hedge? If we cannot answer this question, our credit risk modeling efforts are just an amusement with no practical use. In this section, we trace recent developments in using market data to test credit models and outline directions for future research.

TESTING CREDIT MODELS: THE ANALOGY WITH INTEREST RATES

In Chapters 6 through 13, we examined competing theories of movements in the risk-free term structure of interest rates. In Chapter 14, we reviewed various methods of estimating the parameters of those models. In this section, our task is similar except that the term structure models we are testing is a term structure with credit risk. We devote all of Chapter 17 to that task, but we introduce the topic in this section.

Market Data Test 1: Accuracy in Fitting Observable Yield Curves and Credit Spreads

Continuing the term structure model analogy, one of the easiest and most effective market data tests of a credit model is to see how well the model fits observable yield curve data. All credit models take the existing risk-free yield curve as given, so we can presume we have a smooth and perfectly fitting continuous risk-free yield curve using one of the techniques of Chapter 5. Given this risk-free curve, there is an obvious test of competing credit models: Which model best fits observable credit spreads?

As we discuss in Chapter 17, the market convention for credit spreads is a crude approximation to the exact-day count precision in credit spread estimation that we describe in that chapter. It would be better if we could measure model performance on a set of data that represents market quotations on the basis of credit spreads that we know are "clean." There are many sources of this kind of data, ranging from interest rate swap quotations (where the credit issues are complex) to credit default swaps, a data source of increasing quality concerns because of low trading volume and risks of manipulation. Van Deventer and Imai (2003) present another source—the credit spreads quoted by investment bankers weekly to First Interstate Bancorp, which at the time was the seventh-largest bank holding company in the United States. This data series is given in its entirety in an appendix by van Deventer and Imai.

They then propose a level playing field for testing the ability of the reduced form and Merton models to fit these credit spreads by restricting both models to their two-parameter versions, with these two parameters reestimated for each data point. The

parameters were estimated in a “true to the model” fashion. For the Merton model, for instance, it is assumed that the equity of the firm is an option on the assets of the company as we discuss in Chapter 18. Van Deventer and Imai assume this assertion is true, and solve for the implied value of company assets and their volatility such that the sum of squared error in pricing the value of company equity and the two-year credit spread was minimized. The two-year credit spread was chosen because it was the shortest maturity of the observable credit spreads and therefore closest to the average maturity of the assets held by a bank holding company like First Interstate.

For the Jarrow model, van Deventer and Imai select the simplest version of the model in which the liquidity function discussed previously is a linear function of years to maturity and in which the default intensity $\lambda(t)$ is assumed constant. These assumptions imply that the zero-coupon credit spread is a linear function of years to maturity (Jarrow 1999).

Van Deventer and Imai show that the Jarrow model substantially outperforms the Merton model in its ability to model credit spreads for First Interstate. The Merton model tended to substantially underprice credit spreads for longer maturities, a fact noted by many market participants. The Merton model also implies very short-term credit spreads will be zero for a company that is not extremely close to bankruptcy, while the Jarrow liquidity parameter means the model has enough degrees of freedom to avoid this implication. These problems are discussed extensively in David Lando’s (2004) excellent book on credit modeling

In the tests we have described so far, we are essentially testing the credit model equivalent of the Vasicek term structure model; that is, we are using the pure theory and not “extending” the model to perfectly fit observable data like we do using the extended Vasicek/Hull and White term structure model. We now turn to a test that allows us to examine model performance even for models that fit the observable term structure of credit spreads perfectly.

Market Data Test 2: Tests of Hedging Performance

From a trader’s perspective, a model that cannot fit observable data well is a cause for caution but not a fatal flaw. The well-known “volatility smile” in the Black-Scholes options model does not prevent its use (even though it should), but it affects the way the parameters are estimated. After this tweaking of the model, it provides useful insights on hedging of positions in options. How do the Jarrow and Merton models compare in this regard?

Jarrow and van Deventer (1998) present such a test, again using data from First Interstate Bancorp. They implement basic versions of the Jarrow and Merton models and assume that they are literally true. They then calculate how to hedge a one-week position in First Interstate two-year bonds, staying true to the theory of each model. In the Merton model, this leads to a hedge using U.S. Treasuries to hedge the interest rate risk and a short position in First Interstate common stock to hedge the credit risk. In the Jarrow model, it was assumed that the macro factors driving the default intensities $\lambda(t)$ were interest rates and the S&P 500 index, so U.S. Treasuries and the S&P 500 futures contracts were used to hedge.

Jarrow and van Deventer report that the hedging error using the Jarrow model was on average only 50 percent as large as that of using the Merton model over the entire First Interstate sample.

More importantly, they report the surprising conclusion that the Merton model fails a naïve model test like those we examined earlier in this chapter: a simple duration model produced better hedging results than the Merton model. This surprised Jarrow and van Deventer and many other observers, so Jarrow and van Deventer turned to the next test to understand the reasons for this poor performance.

Market Data Test 3: Consistency of Model Implications with Model Performance

Jarrow and van Deventer investigated the poor hedging performance of the Merton model in detail to understand why the hedging performance of the model was so poor. They ran a linear regression between hedging error and the amount of Merton hedges (in Treasuries and common stock of First Interstate). The results of this standard diagnostic test were again a shock—the results showed that instead of selling First Interstate common stock short to hedge, performance would have been better if instead one had taken a long position in common stock. While the authors do not recommend this as a hedging strategy to anyone, it illustrates what Jarrow and van Deventer found—the movements of stock prices and credit spreads are much more complex than the Merton model postulates. As one can confirm from the first derivatives of the value of risky debt given in Chapter 18, the Merton model says that when the value of company assets rises, stock prices should rise and credit spreads should fall. The opposite should happen when the value of company assets falls, and neither should change if the value of company assets is unchanged.

If the Merton model is literally true, this relationship between stock prices and credit spreads should hold 100 percent of the time. Van Deventer and Imai report the results of Jarrow and van Deventer's examination of the First Interstate credit spread and stock price data—only 40 to 43 percent of the movements in credit spreads (at 2 years, 3, years, 5 years, 7 years, and 10 years) were consistent with the Merton model. This is less than one could achieve with a credit model that consists of using a coin flip to predict directions in credit spreads. This result explains why going short the common stock of First Interstate produced such poor results. Van Deventer and Imai report the results of similar tests involving more than 20,000 observations on credit spreads of bonds issued by companies like Enron, Bank One, Exxon, Lucent Technologies, Merrill Lynch, and others. The results consistently show that only about 50 percent of movements in stock prices and credit spreads are consistent with the Merton model of risky debt. These results have been confirmed by many others and are presented more formally in a recent paper by Jarrow, van Deventer, and Wang (2003).¹⁷

Why does this surprising result come about? And to which sectors of the credit spectrum do the findings apply? First Interstate had credit ratings that ranged from AA to BBB over the period studied, so it was firmly in the investment grade range of credits. Many have speculated that the Merton model would show a higher degree of consistency on lower quality credits. This research is being done intensively at this writing and we look forward to the concrete results of that research.

For the time being, we speculate that the reasons for this phenomenon are as follows:

- Many other factors than the potential loss from default affect credit spreads, as the academic studies outlined in Chapter 17 have found. This is consistent with the Jarrow specification of a general “liquidity” function that can itself be a random function of macro factors
- The Merton assumption of a constant, unchanging amount of debt outstanding, regardless of the value of company assets, may be too strong. Debt covenants and management’s desire for self-preservation may lead a company to over-compensate, liquidating assets to pay down debt when times are bad.
- The Merton model of company assets trading in frictionless efficient markets may miss the changing liquidity of company assets, and this liquidity of company assets may well be affected by macro factors.

What are the implications of the First Interstate findings for valuation of the Jarrow-Merton put option as the best indicator of company risk? In spite of the strong conceptual links with the Merton framework, the Merton model itself is clearly too simple to value the Jarrow-Merton put option with enough accuracy to be useful. (We explore these problems in greater detail in our credit model attic, Chapter 18.)

What are the implications of the First Interstate findings for bond traders? They imply that even a trader with perfect foresight of changes in stock price one week ahead would have lost money more than 50 percent of the time if they used the Merton model to decide whether to go long or short the bonds of the same company. Note that this conclusion applies to the valuation of risky debt using the Merton model—it would undoubtedly be even stronger if one were using Merton default probabilities because they involve much more uncertainty in their estimation (as we shall see in Chapter 18).

What are the implications of the First Interstate findings for the Jarrow model? There are no such implications, because the Jarrow model per se does not specify the directional link between company stock prices and company debt prices. This specification is left to the analyst who uses the Jarrow model and derives the best fitting relationship from the data.

Market Data Test 4: Comparing Performance with Credit Spreads and Credit Default Swap Prices

There is a fourth set of market data tests that one can perform, which is the subject of a very important segment of current research: comparing credit model performance with credit spreads and credit default swap quotations as predictors of default. We address these issues in Chapter 17 in light of recent concerns about manipulation of the LIBOR market and similar concerns about insider trading, market manipulation, and low levels of trading activity in the credit default swap market.

APPENDIX: CONVERTING DEFAULT INTENSITIES TO DISCRETE DEFAULT PROBABILITIES

The reduced form models are attractive because they allow for default or bankruptcy to occur at any instant over a long time frame. Reduced form default probabilities can be converted from:

- Continuous to monthly, quarterly, or annual
- Monthly to continuous or annual
- Annual to monthly or continuous

This section shows how to make these changes in periodicity on the simplifying assumption that the default probability is constant over time. When the default probability is changing over time, and even when it is changing randomly over time, these time conversions can still be made although the formulas are slightly more complicated.

Converting Monthly Default Probabilities to Annual Default Probabilities

In estimating reduced form and hybrid credit models from a historical database of defaults, the analyst fits a logistic regression to a historical default database that tracks bankruptcies and the explanatory variables on a monthly, quarterly, or annual basis. The default probabilities obtained will have the periodicity of the data (i.e., monthly, quarterly, or annual). We illustrate conversion to a different maturity by assuming the estimation is done on monthly data. The probability that is produced by the logistic regression, P , is the probability that a particular company will go bankrupt in that month. By definition, P is a monthly probability of default. To convert P to an annual basis, assuming P is constant, we know that

$$\text{Probability of no bankruptcy in a year} = (1 - P)^{12}$$

Therefore the probability of going bankrupt in the next year is

$$\text{Annual probability of bankruptcy} = 1 - (1 - P)^{12}$$

The latter equation is used to convert monthly logistic regression monthly probabilities to annual default probabilities.

Converting Annual Default Probabilities to Monthly Default Probabilities

In a similar way, we can convert the annual probability of default, say A , to a monthly default probability by solving the equation

$$A = 1 - (1 - P)^{12}$$

for the monthly default probability given that we know the value of A .

$$\text{Monthly probability of default } P = 1 - (1 - A)^{1/12}$$

Converting Continuous Instantaneous Probabilities of Default to an Annual Default Probability or Monthly Default Probability

In the *Kamakura Technical Guide* for reduced form credit models, Robert Jarrow writes briefly about the formula for converting the instantaneous default intensity

$\lambda(t)$ (the probability of default during the instant time t) to a default probability that could be monthly, quarterly, or annual:

For subsequent usage the term structure of yearly default probabilities (under the risk neutral measure) can be computed via

$$Q_t(\tau \leq T) = 1 - Q_t(\tau > T)$$

where

$$Q_t(\tau > T) = E_t \left(e^{-\int_t^T \lambda(u) du} \right)$$

and $Q_t(\cdot)$ is the time t conditional probability.

In this section, we interpret Jarrow's formula. $Q_t(\tau \leq T)$ is the probability at time t that bankruptcy (which occurs at time τ) happens before time T . Likewise, $Q_t(\tau > T)$ is the probability that bankruptcy occurs after time T . Jarrow's formula simply says that the probability that bankruptcy occurs between now (time t) and time T is one minus the probability that bankruptcy occurs after time T . He then gives a formula for the probability that bankruptcy occurs after time T :

$$Q_t(\tau > T) = E_t \left(e^{-\int_t^T \lambda(u) du} \right)$$

E_t is the expected value of the quantity in parentheses as of time t . Because we are using the simplest version of the Jarrow model in this appendix, $\lambda(t)$ is constant so we can simplify the expression for $Q_t(\tau > T)$:

$$Q_t(\tau > T) = e^{-\lambda(T-t)}$$

Converting Continuous Default Probability to an Annual Default Probability

The annual default probability A when λ is constant and $T - t = 1$ (one year) is

$$A = 1 - Q_t(\tau > T) = 1 - e^{-\lambda}$$

Converting Continuous Default Probability to a Monthly Default Probability

The monthly default probability M when λ is constant and $T - t = 1/12$ (one twelfth of a year) is

$$M = 1 - Q_t(\tau > T) = 1 - e^{-\lambda(1/12)}$$

Converting an Annual Default Probability to a Continuous Default Intensity

When the annual default probability A (the KDP), is known, and we want to solve for λ , we solve the equation above for λ as a function of A :

$$\lambda = -\ln(1 - A)$$

Converting a Monthly Default Probability to a Continuous Default Intensity

In a similar way, we can convert a monthly default probability to a continuous default intensity by reversing the formula for M . This is a calculation we would do if we had a monthly default probability from a logistic regression and we wanted to calculate λ from the monthly default probability:

$$\lambda = -12 \ln(1 - M)$$

NOTES

1. Interested readers who would like more technical details on reduced form models should see van Deventer and Imai (2003), Duffie and Singleton [2003], Lando [2004], and Schonbucher (2003). See also the many technical papers in the research section of the Kamakura Corporation website www.kamakuraco.com.
2. See Jarrow and Turnbull for a discussion of the complexities of using a Geske (1979)–style compound options approach to analyze compound credit risks.
3. This assumption is necessary for the rigorous academic framework of the model but not for the model’s practical implementation.
4. Jarrow and Turnbull discuss how to relax this assumption, which we do in the next section.
5. This again means that default risk and interest rates (for the time being) are not correlated. The parameter μ technically represents the Poisson bankruptcy process under the martingale process, the risk-neutral bankruptcy process.
6. In more mathematical terms, $Z(t)$ is standard Brownian motion under a risk-neutral probability distribution $\mathcal{Q}$ with initial value 0 that drives the movements of the market index $M(t)$.
7. The expression for random movements in the short rate is again written under a risk-neutral probability distribution.
8. See Allen and Saunders (2003) in working paper 126 of the Bank for International Settlements for more on the importance of cyclicalities in credit risk modeling.
9. Please see Jarrow, Bennett, Fu, Muxoll, and Zhang, “A General Martingale Approach to Measuring and Valuing the Risk to the FDIC Insurance Funds,” working paper, November 23, 2003 and available on www.kamakuraco.com or on the website of the Federal Deposit Insurance Corporation.
10. The New Basel Capital Accord, Basel Committee on Banking Supervision (May 31, 2001), Section 302, p. 55.
11. For an excellent summary of statistical procedures for evaluating model performance in this regard, see Hosmer and Lemeshow (2000).
12. This calculation can involve a very large number of pairs. The current commercial database at Kamakura Corporation involves the comparison of 1.55 billion pairs of observations, but as of this writing, on a modern personal computer, processing time for the exact calculation is very quick.
13. For a detailed explanation of the explanatory variables used in the KRIS models, please contact e-mail: info@kamakuraco.com.
14. Private conversation with one of the authors, fall 2003.
15. Based on 1.76 million monthly observations consisting of all listed companies in North America for which data was available from 1990 to 2008, Kamakura Risk Information Services database, version 5.0.

16. Recall that the hazard rate estimation procedure determines the best set of explanatory variables from a given set. In the estimation procedure previously discussed, a single financial ratio was a possible outcome, and it was rejected in favor of the multivariable models presented.
17. The treasury department of a top 10 U.S. bank holding company reported, much to its surprise, that their own “new issue” credit spreads showed the same low level of consistency with the Merton model as that reported for First Interstate.

Credit Spread Fitting and Modeling

Chapter 16 provided an extensive introduction to reduced form credit modeling techniques. In this chapter, we combine the reduced form credit modeling techniques with the yield curve smoothing techniques of Chapter 5 and related interest rate analytics in Chapters 6 through 14. As emphasized throughout this book, we need to employ the credit models of our choice as skillfully as possible in order to provide our financial institution with the ability to price credit risky instruments, to calculate their theoretical value in comparison to market prices, and to hedge our exposure to credit risk. The most important step in generating this output is to fit the credit models as accurately as possible to current market data.

If we do this correctly, we can answer these questions:

- Which of the 15 bonds outstanding for Ford Motor Company is the best value at current market prices?
- Which of the 15 bonds should I buy?
- Which should I sell short or sell outright from my portfolio?
- Is there another company in the auto sector whose bonds provide better risk-adjusted value?

Answering these questions is the purpose of this chapter. We continue to have the same overall objective: to accurately measure and hedge the interest rate risk, market risk, liquidity risk, and credit risk of the entire organization using the Jarrow-Merton put option concept as our integrated measure of risk.

INTRODUCTION TO CREDIT SPREAD SMOOTHING

The accuracy of yield curve smoothing techniques has taken on an increased importance in recent years because of the intense research focus among both practitioners and academics on credit risk modeling. In particular, the reduced form modeling approach of Duffie and Singleton (1999) and Jarrow (1999, 2001) has the power to extract default probabilities and the liquidity premium (the excess of credit spread above and beyond expected loss) from bond prices and credit default swap prices.

In practical application, there are two ways to do this estimation. The first method, which is also the most precise, is to use the closed form solution for zero-coupon credit spreads in the respective credit model and to solve for the credit model parameters that minimize the sum of squared pricing error for the observable bonds or credit default swaps (assuming that the observable prices are good numbers, a subject

to which we return later). This form of credit spread fitting includes both the component that contains the potential losses from default and a liquidity premium such as that in the Jarrow model we discussed in Chapter 16. It also allows for the model to be extended to exactly fit observable bond market data in the same manner as the Heath, Jarrow, and Morton (HJM) interest rate models of Chapters 6 through 9 and the Hull and White/extended Vasicek term structure model discussed in Chapter 13.

The second method, which is used commonly in academic studies of credit risk, is to calculate credit spreads on a credit model-independent basis in order to later study which credit models are the most accurate. We discuss both methods in this chapter. We turn to credit model-independent credit spreads first.

THE MARKET CONVENTION FOR CREDIT SPREADS

Credit spreads are frequently quoted to bond-buying financial institutions by investment banks every day. What is a credit spread? From a “common practice” perspective, the credit spreads being discussed daily in the market involve the same kind of simple assumptions such as the yield-to-maturity and duration concepts that we reviewed in Chapters 4 and 12. The following steps are usually taken in the market’s conventional approach to quoting the credit spread on the bonds of ABC Company at a given instant in time relative to a (assumed) risk-free curve, such as the U.S. Treasury curve. That is, to calculate the simple yield to maturity on the bond of ABC Company given its value (price plus accrued interest) and its exact maturity date:

1. Calculate the simple yield to maturity on the “on the run” U.S. Treasury bond with the closest maturity date. This maturity date will almost never be identical to the maturity date of the bond of ABC Company.
2. Calculate the credit spread by subtracting the risk-free yield to maturity from the yield to maturity on the ABC Company bond.

This market convention is simple but very inaccurate for many well-known reasons:

- The yield-to-maturity calculation assumes that zero-coupon yields to each payment date of the given bond are equal for every payment date—it *implicitly assumes the risk-free yield curve is flat*. However, if credit spreads are being calculated for credit risky bonds with two different maturity dates, a different “flat” risk-free curve will be used to calculate the credit spreads for those two bonds because the risk-free yield to maturity will be different for those two bonds.
- The yield-to-maturity calculation usually implicitly assumes *periods of equal length* between payment dates (which is almost never true in the United States for semiannual bond payment dates).
- The maturity dates on the bonds do not match exactly except by accident.
- The payment dates on the bonds do not match exactly except by accident.
- The zero-coupon credit spread is assumed to be flat.
- Differences in coupon levels on the bonds, which can dramatically impact the yield to maturity and, therefore, credit spread, are ignored.
- Often, timing differences in bond price information are ignored. Many academic studies, for example, calculate credit spreads based on bond prices reported

monthly. Since the exact timing of the pricing information is not known, it is difficult to know which date during the month should be used to determine the level of the risk-free yield curve.

- Call options embedded in the bond of ABC Company are often ignored because they cannot be calculated independently of the credit model used. The value of a call option on a bond that is “10 years, noncall five” depends on the probability that the issuer will default in the first five years before the call option can be exercised, for example.

The market convention for fitting credit spreads is very simple, but these problems with the methodology are very serious problems. This kind of inaccuracy is no longer tolerated by traders in fixed income derivatives, such as caps, floors, swaps, swaptions, and so on based on the LIBOR curve. We discuss issues with the LIBOR curve in some detail below. Better analytics for credit spread modeling are now moving into the market for corporate, sovereign, and municipal bonds as well and for very good reason—if you do not use the better technology, you will be “picked off” by traders who do.

A BETTER CONVENTION FOR CREDIT MODEL-INDEPENDENT CREDIT SPREADS

The credit spread calculation that is the market convention has only one virtue—it is simple. We start by proposing a better methodology for analyzing the credit spread of one bond of a risky issuer that avoids most of the problems of the market convention. This methodology continues to assume, however, that the credit spread for each bond of a risky issuer is the same at all payment dates.

We can derive a more precise calculation of credit spread on a credit model-independent basis if we use the yield curve-smoothing technology introduced in Chapter 5. This method has been in commercial use since 1993.¹ Assume there are M payments on the ABC Company bond and that all observable noncall U.S. Treasuries are used to create a smooth, continuous Treasury yield curve using the techniques we described in Chapter 5. Then the continuous credit spread x on the ABC bond can be constructed like this using the following steps:

1. For each of the M payment dates on the ABC Company bond, calculate the continuously compounded zero-coupon bond price and zero-coupon yield from the U.S. Treasury smoothed yield curve. These yields will be to actual payment dates, not scheduled payment dates, because the day count convention associated with the bond will move scheduled payments forward or backward (depending on the convention) if they fall on weekends or holidays.
2. Guess a continuously compounded credit spread of x that is assumed to be the same for each payment date.
3. Calculate the present value of the ABC bond using the M continuously compounded zero-coupon bond yields $y(i) + x$, where $y(i)$ is the zero-coupon bond yield to that payment date on the risk-free curve. Note that $y(i)$ will be different for each payment date but that x is assumed to be constant for all payment dates.
4. Compare the present value calculated in step 3 with the value of the ABC bond (price plus accrued interest) observed in the market.

5. If the theoretical value and observed value are within a tolerance ϵ , then stop and report x as the credit spread. If the difference is outside the tolerance, improve the guess of x using standard methods and go back to step 3.²
6. Spreads calculated in this manner should be confined to noncallable bonds or used with great care in the case of callable bonds.

This method is commonly used by many of the world's largest financial institutions. It can be easily implemented using the solver function in common spreadsheet software.

Deriving the Full Credit Spread of a Risky Issuer

Even the improvement steps listed above for reporting credit spreads ignore the fact that the credit spread is a smooth, continuous curve that takes on the same kind of complex shapes as yield curves do. In order to capture the full credit spread curve, we need to accomplish the following tasks:

- Yield curve smoothing needs to be extended to include inputs like the observable prices of coupon-bearing bonds, not just zero-coupon bond prices.
- Credit spread smoothing analogous to yield curve smoothing has to be implemented.
- Potential problems with the data of risky issuers need to be understood and interpreted with caution. We focus here on issues in the LIBOR market and in the single-name credit default swap market.
- Differences in credit spreads by issuer, maturity date, and observation date need to be explained and understood.

We take on each of these tasks in turn.

Yield Curve Smoothing with Coupon-Bearing Bond Price Input In Chapter 5, we included the maximum smoothness forward rate approach in our comparison of 23 different smoothing techniques, both in terms of smoothness and “tension” or length of the resulting forward and yield curves. In each of our worked examples, we showed how to derive unique forward rate curves and yield curves based on the same set of sample data. This sample data assumed that we had observable zero-coupon yields or zero-coupon bond prices to use as inputs. At most maturities, this will not be the case and the only observable inputs will be coupon-bearing bond prices. In this section, we show how to use coupon-bearing bond prices to derive maximum smoothness forward rates and yields. The same approach can be applied to the 22 other smoothing techniques summarized in Chapter 5.

Revised Inputs to the Smoothing Process In Chapter 5, we used the inputs from Exhibit 17.1 to various yield curve-smoothing approaches.

For the remaining parts of this series on basic building blocks of yield curve smoothing, we assume the following:

- The shortest maturity zero-coupon yield is observable in the overnight market at 4 percent

EXHIBIT 17.1 Actual Input Data

Maturity in Years	Continuously Compounded	
	Zero-Coupon Yield	Zero-Coupon Bond Price
0.000	4.000%	1.000000
0.250	4.750%	0.988195
1.000	4.500%	0.955997
3.000	5.500%	0.847894
5.000	5.250%	0.769126
10.000	6.500%	0.522046

EXHIBIT 17.2 Optimization for Risk-Free Bond Smoothing

Bond Number Bond Name	Additional Inputs for Smoothing		
	1 Bond 1	2 Bond 2	3 Bond 3
Coupon	4.60%	5.40%	6.30%
Years to maturity	1.1	3.25	9.9
Actual NPV	100.30	99.00	99.70

- The 3-month zero-coupon yield is observable at 4.75 percent in the market for short-term instruments, like the U.S. Treasury bill market
- The only other observable instruments are three coupon-bearing bonds with the attributes in Exhibit 17.2.

Note that we are expressing value in terms of observable net present value, which is the sum of price and accrued interest. It is the total dollar amount that the bond buyer pays the bond seller. We ignore the arbitrary accounting division of this number into two pieces (accrued interest and price) because they are not relevant to the smoothing calculation. We assume all three bonds pay interest on a semiannual basis. In doing the smoothing calculation in practice, we would use exact day counts that recognize that semiannual could mean 179 or 183 days or some other number. For purposes of this example, we assume that the two halves of the year have equal length.

Note also that the yield to maturity on these bonds is irrelevant to the smoothing process. The historical yield-to-maturity calculation embeds a number of inconsistent assumptions about forward rates among the three bonds and other observable data, as we discussed earlier in this chapter. To use yield to maturity as an input to the smoothing process is a classic case of “garbage in/garbage out.”

Valuation Using Maximum Smoothness Forward Rates from Chapter 5 The first question we ask ourselves is this: How far off the observable net present values is the net present value we could calculate using Example H Qf1a (where we constrained the forward rate curve to be flat at the 10-year point) in Chapter 5?

To make this present value calculation, we use the coefficients for the five forward rate curve segments from Chapter 5. Recall that we have five forward rate curve segments that are quartic functions of years to maturity:

$$f_i(t) = c_i + d_{i1}t + d_{i2}t^2 + d_{i3}t^3 + d_{i4}t^4$$

The coefficients that we derived for the base case in Example H were the coefficients in Exhibit 17.3.

We can then use these coefficients in this equation to derive the relevant zero-coupon bond yield for each payment date on our three bonds. The yield function in any yield curve segment j is

$$y_j(t) = \frac{1}{t} \left[y^*(t_j)t_j + c_j(t - t_j) + \frac{1}{2}d_{j1}(t^2 - t_j^2) + \frac{1}{3}d_{j2}(t^3 - t_j^3) + \frac{1}{4}d_{j3}(t^4 - t_j^4) + \frac{1}{5}d_{j4}(t^5 - t_j^5) \right]$$

EXHIBIT 17.3 Coefficients

Coefficient Vector	Values
c1	0.04
d11	0.074458207
d12	2.29303E-16
d13	-0.63327238
d14	0.853048721
c2	0.036335741
d21	= 0.133086359
d22	-0.35176891
d23	0.304778043
d24	-0.0850017
c3	0.127299577
d31	-0.23076899
d32	0.194014109
d33	-0.0590773
d34	0.005962133
c4	-0.51565953
d41	0.626509818
d42	-0.23462529
d43	0.036175898
d44	-0.00197563
c5	0.879043488
d51	-0.48925259
d52	0.10010343
d53	-0.0084546
d54	0.000255891

We note that y^* denotes the observable value of y at the left-hand side of the line segment where the maturity is t_j . Within the segment, y is a quintic function of t , divided by t . Using this formula, we lay out the cash flow timing and amounts for each of the bonds, identify which segment is relevant, calculate the zero-coupon yield and the relevant discount factor using the formulas from Part 10 of our blog (Exhibit 17.4).

If we multiply each cash flow by the relevant discount factor, we can get the theoretical bond net present value from the coefficients derived in Example H of Chapter 5 (Exhibit 17.5).

EXHIBIT 17.4 Cash Flow Timing and Amounts

Bond Number	Years to Payment	Cash Flow	Using Example H Qf1a		
			Yield Curve Segment Number	Zero Yield	Discount Factor
1	0.1	2.3	1	4.358165	0.995651
1	0.6	2.3	2	4.860950	0.971256
1	1.1	102.3	3	4.434453	0.952392
2	0.25	2.7	1	4.750000	0.988195
2	0.75	2.7	2	4.729628	0.965150
2	1.25	2.7	3	4.377519	0.946751
2	1.75	2.7	3	4.517091	0.923995
2	2.25	2.7	3	4.940777	0.894789
2	2.75	2.7	3	5.356126	0.863041
2	3.25	102.7	4	5.588093	0.833924
3	0.4	3.15	1	4.912689	0.980541
3	0.9	3.15	2	4.584244	0.959581
3	1.4	3.15	3	4.370700	0.940645
3	1.9	3.15	3	4.629099	0.915804
3	2.4	3.15	3	5.077060	0.885282
3	2.9	3.15	3	5.448956	0.853833
3	3.4	3.15	4	5.614631	0.826217
3	3.9	3.15	4	5.585261	0.804266
3	4.4	3.15	4	5.443462	0.787012
3	4.9	3.15	4	5.279120	0.772072
3	5.4	3.15	5	5.158666	0.756867
3	5.9	3.15	5	5.113988	0.739541
3	6.4	3.15	5	5.150395	0.719193
3	6.9	3.15	5	5.258530	0.695699
3	7.4	3.15	5	5.421918	0.669501
3	7.9	3.15	5	5.621647	0.641395
3	8.4	3.15	5	5.839378	0.612315
3	8.9	3.15	5	6.059341	0.583167
3	9.4	3.15	5	6.269693	0.554687
3	9.9	103.15	5	6.463464	0.527354

EXHIBIT 17.5 Additional Discounts for Smoothing

Bond Number Bond Name	Additional Inputs for Smoothing		
	1 Bond 1	2 Bond 2	3 Bond 3
Coupon	4.60%	5.40%	6.30%
Years to maturity	1.1	3.25	9.9
Actual NPV	100.3	99	99.7
Example H QF1a NPV	101.953547	100.715178	100.694080
Pricing error	1.653547	1.715178	0.994080

EXHIBIT 17.6 Actual Input Data

Maturity in Years	Continuously Compounded	
	Zero-Coupon Yield	Zero-Coupon Bond Price
0.000	4.000%	1.000000
0.250	4.750%	0.988195
1.000	4.500%	0.955997
3.000	5.500%	0.847894
5.000	5.250%	0.769126
10.000	6.500%	0.522046

The table shows that the zero-coupon yields that we used in the smoothing process produced incorrect net present values. As we shall see below, the problem is not the maximum smoothness forward rate technique itself. It is the inputs to the smoothing process at the 1-, 3-, 5-, and 10-year maturity in our input table (Exhibit 17.6).

Given these inputs, our smoothing coefficients and NPVs follow directly as we have shown above and in Chapter 5. We now improve our valuations by changing the zero yields used as inputs to the process.

Iterating on Zero-Coupon Yields to the Smoothing Process We know that the 4 percent yield for a maturity of zero and the 4.75 percent yield for 0.25 year maturity are consistent with observable market data. That is clearly not true for the 1-, 3-, 5-, and 10-year zero yields because (1) there are no observable zero-coupon bond yields at those maturities and (2) bond prices in the market are trading at net present values that are inconsistent with the yields we have been using at those maturities so far. We now pose this question:

What values of zero-coupon bond yields at maturities of 1, 3, 5, and 10 years will minimize the sum of squared pricing errors on our three observable bonds?

EXHIBIT 17.7 Actual Input Data

Maturity in Years	Continuously Compounded	
	Zero-Coupon Yield	Zero-Coupon Bond Price
0.000	4.000%	1.000000
0.250	4.750%	0.988195
1.000	5.912%	0.942591
3.000	6.138%	0.831815
5.000	6.078%	0.737946
10.000	6.384%	0.528157

We can answer this question using a powerful enterprise-wide risk management like Kamakura Risk Manager (see www.kamakuraco.com) or even by using common spreadsheet software's nonlinear optimization (solver) routines. Using the latter approach, we find that this set of inputs eliminates the bond pricing errors (Exhibit 17.7).

The forward rate curve coefficients for maximum smoothness forward rate smoothing that are consistent with these inputs are given in Exhibit 17.8.

The table in Exhibit 17.9 confirms that the sum of the cash flows multiplied by the relevant discount factors from this revised set of coefficients produces bond net present values that exactly match those observable in the marketplace.

This simple example shows that using bond prices as input to the smoothing process is only a small step forward from using zero-coupon bond yields. It goes without saying that there is no need to use an arbitrary functional form, like the Nelson-Siegel approach, that is neither consistent with the observable bond prices nor optimum in terms of smoothness. It continues to mystify us why that technique is ever employed.

For the remainder of this chapter, we use bond prices as inputs. What if the bonds are issued by a firm with credit risk? How should the analysis differ from the case when the issuer of the bonds is assumed to be risk free? We turn to that analysis in the next section.

CREDIT SPREAD SMOOTHING USING YIELD CURVE-SMOOTHING TECHNIQUES

In this section, we show that it is incorrect, sometimes disastrously so, to apply smoothing to a yield curve for an issuer with credit risk, ignoring the risk-free yield curve. We show alternative methods for smoothing the zero-coupon credit spread curve, relative to the risk-free rate, to achieve far superior results. Smoothing a curve for an issuer with credit risk without reference to the risk-free curve is probably the biggest mistake one can make in yield curve smoothing and we show why below.

Setting the Scene: Smoothing Results for the Risk-Free Curve

We start first by assuming that there is a risk-free bond issuer in the currency of interest. Prior to the 2006–2011 credit crisis, this assumption might have been taken

EXHIBIT 17.8 Coefficients

Coefficient Vector	Values
c1	0.040000
d11	0.064352
d12	0.000000
d13	-0.182575
d14	0.216590
c2	0.039117
d21	0.078478
d22	-0.084757
d23	0.043444
d24	-0.009429
c3	0.047700
d31	0.044147
d32	= -0.033261
d33	0.009113
d34	-0.000847
c4	0.126353
d41	-0.060723
d42	0.019174
d43	-0.002539
d44	0.000124
c5	0.052668
d51	-0.001776
d52	0.001490
d53	-0.000181
d54	0.000007

EXHIBIT 17.9 Additional Inputs for Smoothing

Bond Number	1	2	3
Bond Name	Bond 1	Bond 2	Bond 3
Coupon	4.60%	5.40%	6.30%
Years to maturity	1.1	3.25	9.9
Actual NPV	100.30	99.00	99.70
Theoretical NPV	100.30	99.00	99.70
NPV errors	0.00	0.00	0.00

for granted, but given current concerns about sovereign risk we make the assumption here that the government yield curve is truly free of credit risk. We use the data from the previous section for our risk-free curve.

We smoothed the risk-free yield curve for two reasons: (1) So we can derive pricing for other risk-free transactions than those with observable prices, and, (2) so we can derive a high-quality credit spread for ABC Company. The ABC yield curve




EXHIBIT 17.10 Risk-Free Zero Yields and Forward Rates Using Maximum Smoothness Forward Rates

will allow us to make a new financing proposal to ABC Company that is consistent with observable credit risky bonds of ABC Company.

The zero-coupon bond yields and forward rates for the risk-free government bond curve derived in the previous section are graphed in Exhibit 17.10.

We now use this information to derive the best possible credit spreads for ABC Company.

A Naïve Approach: Smoothing ABC Yields by Ignoring the Risk-Free Curve

We start by first taking a naïve approach that is extremely common in the yield curve literature: We smooth the yield curve of ABC Company on the assumption that we don't need to know what the risk-free curve is. This mistake is made daily with respect to the U.S. dollar interest rate swap curve by nearly every major financial institution in the world. They assume that the observable data points are so numerous that the results from this naïve approach will be good, and then they blame the yield curve smoothing method when the results are obviously garbage. In truth, this is not one of Nassim Taleb's black swans—it's an error in finance made by the analyst, and the mistake is due to incorrect assumptions by the analyst rather than a flaw in the smoothing technique.

We show why in this example. We assume that these zero-coupon bonds are outstanding issues of ABC Company with observable net present values (Exhibit 17.11).

EXHIBIT 17.11 Input Data for Risky Yield Curve Creation

Maturity in Years	Risk-Free Zero Yield	Observable Credit Spread	Risky Zero Yield	Zero-Coupon Bond Price
0.100	4.317628	0.450000	4.767628	0.995244
2.750	6.164844	0.490000	6.654844	0.832761
3.250	6.114068	0.500000	6.614068	0.806576
3.400	6.101630	0.510000	6.611630	0.798680
9.400	6.353080	0.620000	6.973080	0.519198
9.900	6.378781	0.630000	7.008781	0.499639

The risky zero yield is derived using continuous compounding from the (assumed) observable zero-coupon bond price. Derivation is similar if the observable data were risky coupon bearing bonds of ABC Company. We know what the risk-free zero-coupon yield is from the smoothing of the risk-free curve in the prior section and in Chapter 5. The observable credit spread is simply the difference between the observable ABC Company zero-coupon yield and the risk-free yield that we derived above on exactly the same maturity date, to the exact day. We now want to smooth the risky yield curve, naïvely assuming that the risk-free curve is irrelevant because we have six observable bonds and we think that's enough to avoid problems.

We make the additional assumption that the credit spread for a very short maturity is also 45 basis points, so the assumed zero maturity yield for ABC Company is the risk-free zero yield of 4.00 percent plus 0.45 percent, a total of 4.45 percent. We apply the maximum smoothness forward rate technique to these zero-coupon bonds, exactly as in Chapter 5. If instead ABC bonds were coupon bearing, we would have used the approach for coupon-bearing bonds like we did at the start of this chapter.

We compare the zero-coupon yield curve that we derive for ABC Company with the zero-coupon yield curve for the government in the graph in Exhibit 17.12.

We can immediately see that we have a problem. ABC Company's fitted yield curve, which ignored the government curve, causes negative credit spreads on the short end of the yield curve, because the ABC yield curve is less than the risk-free zero-coupon yield curve.

When we compare the forward rates from the two curves, we can see this problem more clearly (see Exhibit 17.13).

Taking one more view, we plot the zero-coupon credit spread (which is the difference between the ABC Company zero yields and government yields) and the credit spread forwards (the difference between the ABC Company forward rates and the government forward rates) in the graph in Exhibit 17.14.

The darker line is the zero-coupon credit spread. It is volatile and goes negative in the maturity interval between 0.25 years and 1.25 years. The problem with the credit spread forwards makes these issues even more obvious. The forward credit spreads go negative by approximately 50 basis points. This means that any transactions for ABC Company in this interval will be grossly mispriced.

What does the analyst do now? Most analysts are like the football coach whose team loses because the coach had a bad game plan—the coach blames the




EXHIBIT 17.12 Comparison of Zero-Coupon Bond Yield Curves for Risk-Free Curve and Risky Curve Using Maximum Smoothness Forward Rate Smoothing




EXHIBIT 17.13 Comparison of Zero-Coupon Bond Forward Curves for Risk-Free Curve and Risky Curve Using Maximum Smoothness Forward Rate Smoothing




EXHIBIT 17.14 Credit Spread and Forward Credit Spreads between Risk-Free Curve and Risky Curve Using Maximum Smoothness Forward Rate Smoothing on Risky Yield

quarterback. In this case, most analysts would blame the smoothing method, but in fact the problem is caused by the analyst's decision to ignore the risk-free curve. We show how to fix that now.

FITTING CREDIT SPREADS WITH CUBIC SPLINES

A natural alternative to smoothing the yields of ABC Company is to smooth the zero-coupon credit spreads. One should never commit the unimaginable sin of smoothing the differences between the yield to maturity on an ABC bond with N years to maturity and the simple yield to maturity on a government bond that also has something close to N years to maturity. The mistakes embedded in this simple calculation were outlined previously. The fact that this mistake is made often does not make the magnitude of the error any smaller. We make two choices in this section. First, we again assume that the credit spread at a zero maturity for ABC Company is the same 45 basis point spread that we see on the shortest maturity bond. Second, we choose to use a cubic spline to solve for the ABC credit spread curve and we require that the observable credit spreads and actual credit spreads be equal.

Each credit spread curve segment will have a cubic polynomial that looks like this:

$$c(t) = a + b_1t + b_2t^2 + b_3t^3$$

EXHIBIT 17.15 Credit Spreads

	1	2	3	4	5	6
Risky zero yield						
Maturity in years	0.100	2.750	3.250	3.400	9.400	9.900
Actual zero spread	0.450%	0.490%	0.500%	0.510%	0.620%	0.630%
Theoretical zero spread	0.450%	0.489%	0.504%	0.506%	0.623%	0.628%
Yield errors	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Sum of squared yield errors times 100,000,000						0.46

EXHIBIT 17.16 Actual Input Data

Maturity in Years	Zero-Coupon, On-the-Run
0.000	0.450%
0.250	0.449%
1.000	0.428%
3.000	0.499%
5.000	0.498%
10.000	0.628%

where t is the years to maturity of that instantaneous credit spread. In Chapter 5, we showed how to fit cubic splines to any data. In that chapter, it was zero-coupon yields. In this section, we fit to zero-coupon credit spreads.

We arbitrarily choose to use knot points of 0, 0.25 years, 1, 3, 5, and 10 years. These are the same knot points we used for risk-free curve smoothing. We could have just as easily used the observable yields and maturities in the chart above. In that case, all that credit spread smoothing does is to connect the observable credit spreads: 45 basis points at a zero maturity, 45 basis points at 0.1 years, 49 basis points at 2.75 years, and so on.

We use the risk-free curve knot points and iterate until we have an (almost perfect) match to observable credit spreads (Exhibit 17.15).

The match can be made perfect by changing parameters in the optimization routine. The zero-coupon credit spreads that produce this result are shown in Exhibit 17.16.

The zero maturity is not shaded because it is not part of the iteration. We have assumed that the zero-maturity credit spread is the same 45 basis points as the observable data for 0.1 years to maturity. Once we have the spline parameters as in Chapter 5, we can plot zero-coupon credit spreads and zero-coupon forward credit spreads. The results are plotted later in the chapter.

We now turn to maximum smoothness forward credit spreads.

MAXIMUM SMOOTHNESS FORWARD CREDIT SPREADS

In this section, we apply the maximum smoothness forward rate approach of Chapter 5 to forward credit spreads, not forward rates. The forward credit spreads

EXHIBIT 17.17 Credit Spreads

Risky zero yield	1	2	3	4	5	6
Maturity in years	0.100	2.750	3.250	3.400	9.400	9.900
Actual zero spread	0.450%	0.490%	0.500%	0.510%	0.620%	0.630%
Theoretical zero spread	0.450%	0.489%	0.504%	0.507%	0.621%	0.629%
Yield errors	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Sum of squared yield errors times 100,000,000						0.23

EXHIBIT 17.18 Actual Input Data for On-the-Run, Risky
Zero-Yield Credit Spread Forwards with $f''(0) = 0$

Maturity in Years	On-the-Run Zero-Coupon Yield Spread
0.000	0.450%
0.250	0.449%
1.000	0.440%
3.000	0.497%
5.000	0.532%
10.000	0.631%

are simply the difference between the ABC Company forward rates and the government yield curve forward rates. When we recognize that relationship, the approach in Chapter 5 can be used with extremely minor modifications. We iterate on zero-coupon credit spreads instead of zero-coupon yields.

When we do that iteration, we get an almost perfect fit. A perfect fit just requires tweaking the tolerance in the iteration routine (Exhibit 17.17).

The zero-coupon credit spreads that produce this happy result are given in Exhibit 17.18.

Again, the zero maturity spread of 45 basis points is not part of the iteration as above.

We can plot the zero-coupon credit spreads and forward credit spreads that result (see Exhibit 17.19).

The resulting credit spread curve is very well-behaved and the forward credit spread is smooth with only modest variation. This confirms what we asserted in the naïve example: No problem exists with maximum smoothness forward rate smoothing. The early negative credit spreads resulted from the analyst's error in ignoring risk-free rates.

COMPARING RESULTS

We can compare our three examples by looking first at the credit spreads from risky yield smoothing, cubic credit spread smoothing, and maximum smoothness forward credit spread smoothing (Exhibit 17.20).




EXHIBIT 17.19 Credit Spread and Forward Credit Spreads between Risk-Free Curve and Risky Curve Using Maximum Smoothness Forward Credit Spread Smoothing with $f''(0) = 0$ and $f'(T) = 0$ on Risky Yield




EXHIBIT 17.20 Credit Spreads between Risk-Free Curve and Risky Curve Using Alternative Methods to Smooth Credit Spreads




EXHIBIT 17.21 Forward Credit Spreads between Risk-Free Curve and Risky Curve Using Alternative Methods to Smooth Credit Spreads

Most analysts would assert that the maximum smoothness forward credit approach is superior because the resulting zero-coupon credit spread curve is smoother. The same is true, by definition, for forward credit spreads (see Exhibit 17.21).

These results make it obvious that standard yield curve smoothing techniques can be applied to the zero-coupon credit spread curve. The performance of alternative techniques is essentially identical to the performance of the same techniques applied to zero-coupon yields themselves.

In the next two sections, we discuss how data problems can impact credit spread smoothing using any smoothing technique. We start with the case of the London Interbank Offered Rate and interest rate swap curve.

DATA PROBLEMS WITH RISKY ISSUERS

The Case of LIBOR

There are two problems that immediately confront an analyst who is seeking a smooth credit spread curve between the U.S. Treasury curve and a combination curve composed of one-, three-, and six-month LIBOR rates and the interest rate swap quotations for longer maturities. This graph plots the zero-coupon bond yields that



EXHIBIT 17.22 U.S. Treasury and USD Interest Rate Swap Zero-Coupon Yields Derived from the Federal Reserve H15 Statistical Release Using Maximum Smoothness Forward Rate Smoothing

Sources: Kamakura Corporation; Federal Reserve.

are produced when one applies the maximum smoothness forward rate technique to the U.S. Treasury curve and, as we have outlined above, the same technique to the credit spreads between the LIBOR-swap curve and the Treasury curve. There is a distinct “kink” in the LIBOR-swap zero-coupon bond yields as shown clearly in this graph using data for June 7, 2012, reported by the Federal Reserve in its H15 statistical release (Exhibit 17.22).

Many analysts will blame this kink on the yield curve–smoothing technique, but it is actually the result of a mistake by the analyst. The analyst has made a common but serious mistake by assuming that Eurodollar deposits and interest rate swaps are homogeneous and can be appropriately considered to be part of the same curve. In fact, Eurodollar deposits are much more risky, with a potential loss to the depositor of 100 percent of principal. Interest rate swaps, by contrast, involve no exchange of principal at maturity and even a mark-to-market loss during the life of the swap is required to be collateralized. For this reason, the kink recognizes the “regime switch” from a high credit risk level on the short end of the curve to a lower credit risk regime on the long end of the curve. At 30 years, the interest rate swap yield is below (has a negative spread) the U.S. Treasury yield and has been for more than two years as this book is going to press. Again, this is a data issue and not a smoothing issue.

The second data issue with the LIBOR-swap curve is more serious and more disturbing: corruption and market manipulation. As of press time, three major

lawsuits alleging manipulation of U.S. dollar LIBOR are pending in the U.S. courts and investigations are under way in Japan and Europe. Market participants have been fired and a leading bank is cooperating with the investigation. The authors have analyzed patterns of LIBOR indications compared to third-party estimates of funding costs and default risk and believe that the allegations of LIBOR manipulation are likely to be proven true.

For this reason, using the LIBOR curve as a basis for risk management and using the LIBOR specific “LIBOR market model” is fraught with danger.

DETERMINANTS OF CREDIT SPREAD LEVELS

Let’s accept as reality the risk that the publicly available credit default swap quotes have been manipulated (see Chapter 20 for a detailed discussion of CDS liquidity). What factors have had an impact on the level of CDS spreads quoted? Jarrow, Li, Matz, Mesler, and van Deventer (2007) examine this issue in detail.

Jarrow, Li, Matz, Mesler, and van Deventer (JLMMV) argue that corporate credit spreads, as measured in the credit default swap market, depend not only on expected losses but also on market liquidity, institutional restrictions, and market frictions. The difference is a credit risk premium. This credit risk premium is justified theoretically using a simple economic model, and this credit risk premium is validated empirically using credit default swap spreads.

It is still common to hear traders argue that observable spreads in the credit default swap market are equal to the product of the default probability for the underlying credit times the loss given default. But, this simple trader’s argument obviously ignores supply and demand considerations in the market for credit insurance. Indeed, if CDS spreads only reflect expected loss with no risk premium, then why would institutional investors be willing to be providers of credit insurance in the CDS market? The Jarrow model reviewed in Chapter 16 recognizes this phenomenon explicitly with its provision for a “liquidity premium.”

Consistent with financial theory, and contrary to the simple trader’s argument, researchers have consistently found that bond spreads and credit default swap quotations are higher than historical credit loss experience for those particular credits. These results are usually obtained by comparing credit spreads, derived by the previous smoothing techniques, with either actual or projected default risk and recovery. Empirically, a credit risk premium appears to exist.

Nonetheless, the simple trader’s argument is still common in the industry. Consequently, the purpose of this section is twofold. First, to help dispel the trader’s argument, we explain in simple terms and using standard economic theory why credit spreads should exceed expected losses (the default probability times loss given default). Second, given the existence of a credit risk premium, we then provide a statistical model for estimating credit spreads in the credit default swap market that uses other explanatory variables (in addition to default probabilities and losses given default) to capture credit risk premium. The statistical model performs quite well. The estimation is performed using a large database of credit default swap prices provided by the broker GFI. We accept the GFI data “as is” with no adjustments for possible data manipulation by market participants.

THE CREDIT RISK PREMIUM: THE SUPPLY AND DEMAND FOR CREDIT

The degree of competition among lenders for a particular financial instrument varies dramatically with the attributes of the instrument and the attributes of the lenders. A hedge fund trader recently commented:

Bond prices in the market place anticipate ratings changes with one exception. When a firm is downgraded to “junk,” a very large number of institutions are compelled to sell immediately. This causes a step down in bond prices of large magnitude, unlike any other change in ratings.

This statement reflects market segmentation, generated by heterogeneity in lenders' risk aversion, trading restrictions, and private information. Implicitly this quotation says that the number of potential lenders to a junk credit is much less than the number of potential lenders to an investment grade credit, even if the default risk for these companies are the same. There is less demand for the financial liabilities of the junk credit, and spreads therefore must widen. In the limit, a small and risky business may have only one lender willing to lend it monies. In this case, the fact that the lender is a monopoly supplier (with their own risk aversion, private information, and institutional constraints) has at least as much impact on the level of the spread as does the default probability or loss given default of the borrower.

Consider an example. On the demand side, as a company's true one-year default probability increases, its need for funds increases. It needs more funds because, as the default probability rises, its revenues are declining and its expenses are increasing. Conversely, when the company has strong cash flows, it has very low default risk and it has little need for additional funds.

On the supply side, when the company is a strong credit, the potential supply of funds to the company is very large. More lenders are willing to participate in loans to the company. This is due to the heterogeneity of lenders' risk aversion, private information, institutional constraints, and transaction costs. Indeed, those lenders that are very risk averse, or who are not privy to private information concerning this company's cash flows, are willing to lend because the risk is low and private information is less relevant. Also, any institutional investment restrictions regarding junk lending will not be binding. As the company becomes more risky, the number of lenders drops, due to the desire to avoid risk by an increasing number of lenders and related concerns regarding private information about the borrower's prospects. Institutional constraints (as mentioned above) will also kick in at an appropriate level. The maximum potential supply of funds drops steadily until, in the limit, it reaches zero.

Equilibrium is *not* where the supply of funds as a function of the one-year default probabilities exactly equals the funding needs. This is because there are other considerations to an economic equilibrium in the credit risk market. For example, loss given default is a relevant and missing consideration in the determination of a borrower's risk, as are third-party bankruptcy costs (e.g., legal fees) and transactions costs incurred in the market for borrowed funds. From the borrower's side, a financial strategy that retains the option to seek additional funds at a later date holds a real option of considerable value and this option provides a reason not to borrow

the maximum at any moment in time. These missing considerations make the analysis more complex than this simple example indicates, but the downward-sloping nature of the supply curve for lenders' funds will still hold. That is, as the credit risk of a borrower increases, the volume of available loans declines.

An individual lender, realizing the downward sloping nature of the aggregate supply curve, will decide on the volume of their lending based on their marginal costs and revenues. Marginal cost includes the expected loss component of the credit spread plus the marginal costs of servicing the credit. The expected loss component, equal to the default probability times loss given default, depends on the lender's private information regarding a borrower's default probabilities and potential loss given default.

Marginal revenue represents the lender's risk-adjusted percentage revenue per additional dollar of loan. Due to risk aversion, the risk-adjusted marginal revenue curve will be downward sloping.³ To induce a risk-averse lender to hold more of a particular loan in a portfolio, increasing idiosyncratic risk, the lender must be compensated via additional spread revenue. Alternatively, and independent of risk aversion, if the lender's loan volume is sufficient to affect the market spread (i.e., the lender is not a price taker), then the marginal revenue curve will be downward sloping based on this consideration alone. For either reason, as the loan volume increases, the percentage revenue per dollar of additional lending declines.

Standard economic reasoning then implies that the lender will extend loans until their risk-adjusted marginal revenue equals their marginal cost. Ignoring for the moment market frictions (the difference between marginal costs and the expected loss), we see that, for the marginal dollar lent, the marginal credit spread equals the expected loss ($PD * LGD$). Next, adjusting for market frictions, on the margin, the marginal credit spread equals (marginal costs + $PD * LGD$). But, the observed credit spread reflects the *average spread per dollar loan*, and not the marginal spread. Since the marginal revenue spread curve is downward sloping, the average revenue spread per dollar loan exceeds the marginal revenue spread curve. Thus, the observed credit spread per dollar loan is given by:

$$\text{Credit spread} = (\text{Average marginal revenue spread} + \text{Marginal costs} + PD * LGD)$$

The first component is due to risk aversion and or market liquidities; that is, the fact that the supply curve for loans is downward sloping. The second component is due to market frictions—the marginal costs of servicing loans. The third component is the expected loss itself.

This simple economic reasoning shows that a trader who argues that the CDS quote equals the default probability times the loss given default is really assuming an extreme market situation. The extreme market situation requires that the first and second components of the credit spread (as depicted above) are zero. The first component is zero only if the supply curve for loans is horizontal. That is, there is an infinite supply of funds at the marginal cost of credit—markets are perfectly liquid and there is no credit risk premium. And, the second component is zero only if there are no costs in servicing loans. These are unrealistic conditions, and their absence explains why the simple trader's argument is incorrect.

JLMMV argue that there are two methods available to determine the drivers of credit risk spreads. The first method is to build an economic model, akin to that described previously, and fit the model to observed credit risk spreads. The second is to fit a statistical model. A statistical model is often useful when building an (equilibrium) economic model is too complex. And, furthermore, a statistical model can be thought of as an approximation to the economic model.

As suggested by the discussion in the previous section, building a realistic equilibrium model for credit spreads with lender heterogeneity in risk aversion, private information, and institutional restrictions, is a complex task. Furthermore, such a construction is subject to subjective assumptions regarding lender preferences, endowments, institutional structures as well as the notion of an economic equilibrium. To avoid making these subjective assumptions, which limit the applicability of the model selected, and, in order to obtain a usable model for practice, JLMMV elect to build a statistical model for credit risk spreads instead.

The JLMMV statistical model uses various explanatory variables, selected to capture the relevant components of the credit spread, as discussed in the previous section. In particular, JLMMV seek explanatory variables that will capture the credit risk premium, market liquidity, institutional constraints, and expected losses. Before discussing these explanatory variables, it is important to first discuss the functional form of the explanatory variables fit to the data.

Previous authors (Campbell and Taksler 2003, Collin-Dufresne et al. 2001, Huang et al. 2003, and Elton et al. 2001) fitting a statistical model to credit spreads typically use a linear function to link the credit spread to explanatory variables via ordinary least squares regression. Implicit in the linear regression structure, however, is the possibility that, when the model is used in a predictive fashion, the statistical model may predict negative credit spreads. Negative credit spreads, of course, are inconsistent with any reasonable economic equilibrium as long as the risk-free curve is truly risk free. If one were in Greece in 2012, for example, it would be a serious mistake to consider the credit curve for the Hellenic Republic to be a risk-free curve.

The solution to this problem is simple. One only needs to fit a functional form that precludes negative credit spreads. Fitting a linear function to $\ln[\text{Credit spread}(t)]$ is one such transformation. Another transformation is to use the logistic function,

$$\text{Credit spread}(t) = \frac{1}{1 + e^{-\alpha - \sum_{i=1}^n \beta_i X_i}}$$

where alpha and beta for $i = 1, \dots, n$ are constants and X_i for $i = 1, \dots, n$ are the relevant explanatory variables.

Unlike the natural logarithm, the logistic formula has the virtue that predicted CDS spreads always lie between 0 and 100 percent. And, similar to the use of the natural logarithm, one can estimate the alphas and betas in the logistic formula by using the transformation, $(-\ln[(1 - \text{Credit spread}[t])/\text{Credit spread}[t]])$ to fit a linear function in the explanatory variables via ordinary least squares regression. This linear function can be thought of as Altman's Z-score, for example. Alternatively, one can use a general linear model for the derivation. In this post, we use the logistic regression approach to model credit risk spreads.

The broker GFI supplied the CDS quotes used in the JLMMV estimation. Their database included daily data from January 2, 2004, to November 3, 2005, which includes more than 500,000 credit default swap bid, offered, and traded price observations. This was a particularly benign period in U.S. economic history, a brief period of calm prior to the 2006 to 2011 credit crisis. Bid prices, offered prices, and traded prices were all estimated separately. In the JLMMV database, there were 223,006 observations of bid prices, 203,695 observations of offered prices, and 19,822 observations of traded prices. Traded CDS prices were only one-tenth as numerous as the bid and offered quotations, consistent with the concerns about market manipulation discussed previously. CDS quotes where an up-front fee was charged were excluded from the estimation because conversion of the up-front fee to a spread equivalent requires a joint hypothesis about the term structure and relationship of default probabilities and credit spreads, and that is what we are trying to derive in this post. The JLMMV data set predates the change in ISDA credit default swap convention that made up-front fees the standard, not the exception.

The explanatory variables that JLMMV used to fit CDS prices include the credit rating, estimated default probabilities, and company-specific attributes. Two types of estimated default probabilities were used: one from the Jarrow-Chava reduced form model that we discussed in Chapter 16, and the second from a Merton-type structural model, which we review in Chapter 18. All of these estimates were obtained from version 3.0 of the Kamakura Risk Information Services default probability service. The CDS maturities, agency ratings, and all of the individual macroeconomic factors and company-specific balance sheet ratios that are inputs to the Jarrow-Chava reduced form model were also included as inputs. These macro variables and company-specific factors are listed below. These macro and micro variables are intended to capture the credit risk premium, market liquidities, institutional constraints, and market frictions as previously discussed.

Forty-seven variables were found to be statistically significant in predicting CDS spreads. For bid prices, the explanatory variables included:

- Seven maturities of KRIS version 3.0 Jarrow-Chava default probabilities
- KRIS version 3.0 Merton default probabilities
- Dummy variables for each rating category
- A dummy variable indicating whether or not the company is a Japanese company
- 10 company-specific financial and equity ratios (see Jarrow and Chava 2004 for a list of the relevant variables)
- Dummy variables for each CDS maturity
- Dummy variables for senior debt
- Dummy variables for the restructuring language of the CDS contract
- Selected macroeconomic factors (see Jarrow, Li, Matz, Mesler, and van Deventer 2007 for details)

The best-fitting JLMMV relationship for each of the three series (bid, offer, and trade) explained more than 80 percent of the variation in the transformed CDS quotations. The best-fitting relationship also explains more than 90 percent of the variation in the raw CDS quotations (after reversing the log transformation

explained previously). The t -scores of the explanatory variables ranged from 2 to 236. The five-year maturity Jarrow-Chava default probability had the highest t -score among all the different maturity default probabilities included. The Japanese dummy variable was statistically significant with a t -score equivalent of 133 standard deviations, indicating that CDS spreads on Japanese names are much narrower than the otherwise equivalent non-Japanese name. Company size had a t -score equivalent of 100 standard deviations from zero, implying that CDS spreads are narrower for large companies, everything else constant. The JLMMV statistical model fits CDS market prices quite well, validating the existence of credit risk premium in CDS prices.

JLMMV found that Merton default probabilities explained 12 to 28 percent of the variation in the transformed CDS quotes. Ratings explained 36 to 42 percent of the transformed CDS quotes. The Kamakura Risk Information Services (KRIS) version 3.0 Jarrow-Chava default probabilities and all of their explanatory variables explained 56 to 61 percent of the transformed variables. The superhybrid approach, which includes all factors, explained 81 to 83 percent of the variation in the transformed CDS quotes.

CONCLUSION

Credit spreads are driven by the supply and demand for borrowed funds, and not just the default probability and loss given default. This is true for all borrowers from retail to corporate to sovereign. When fitting a statistical model to corporate sector CDS spreads, hazard rate default probabilities and their inputs dominate ratings, and ratings dominate Merton default probabilities, in their ability to explain movements in CDS quotes. A superhybrid approach provides the best overall explanatory power for CDS quotes. In the next chapter, we compare this modern approach to credit risk assessment to legacy approaches that have been used over the past century.

NOTES

1. This method has been a standard calculation of the Kamakura Risk Manager enterprise-wide risk management system since 1993.
2. Chapter 1 of van Deventer and Imai's *Financial Risk Analytics* (1996) illustrates one method commonly used for this purpose.
3. It might be argued that lenders—banks—are risk neutral and not risk averse. However, it can be shown that risk-neutral lenders, subject to either regulatory capital constraints and/or subject to deadweight losses in the event of their own default, will behave as if they are risk averse (see Jarrow and Purnanandam 2005).

Legacy Approaches to Credit Risk

In this chapter, we do for credit risk models what we did for interest rate risk models in Chapter 13. We go up into the risk management attic and dust off the tools that may have been best practice once, but which are only described that way now by a few. We review legacy credit ratings and the Merton model using the tools of Chapter 16. We do this because the selection of credit risk tools is not a beauty contest. It's all about accuracy, period. Arguments to the contrary are usually about regulatory or corporate politics and the slow speed at which very large organizations respond to and absorb new technology. In that spirit, we explain why legacy credit ratings and the Merton default model, used almost exclusively for public firms, have been firmly outperformed by reduced form credit models. The dominance of the reduced form approach for retail borrowers and nonpublic firms has never been challenged by tools like legacy ratings and the Merton model, where the assumptions we outline below are even less appropriate than they are for public firms.

In this chapter, we return to using “Model Risk Alerts” with respect to statements known to cause problems or known to be false but commonly stated by practitioners.

THE RISE AND FALL OF LEGACY RATINGS

Almost every financial institution has used legacy credit ratings actively in both individual asset selection and in portfolio management at some point in its past. Ratings have a history measured in decades and they will probably be used actively many decades from now despite the fatal flows outlined in this section. That said, rating agencies and experts in internal ratings now face substantial pressure, both commercial and conceptual, from providers of quantitative default probabilities both inside and outside of the organization. In this section, we summarize some of the strengths and weaknesses of a traditional ratings approach for comparison with the quantitative approaches that increasingly dominant risk management and credit derivatives pricing. We present proof consistent with Hilscher and Wilson (2011) that ratings are consistently outperformed from an accuracy point of view by quantitative default models built using logistic regression and the reduced form modeling approach. We also present the even more serious problems with ratings summarized by Senator Carl Levin and the U.S. Senate in the wake of the 2006 to 2011 credit crisis.

RATINGS: WHAT THEY DO AND DON'T DO

For many, many years Moody's Investors Service and Standard & Poor's were essentially the only providers of credit risk assessment on large corporations and sovereign issuers of debt. This has changed recently and dramatically with a plethora of credit rating agencies springing up around the world to challenge the older rating agencies. Competing rating agencies and vendors of default probabilities have arisen for a number of reasons, just some of which are noted here:

- The *granularity* of ratings offered by the major rating agencies is low. The number of ratings grades (AAA, AA, A, BBB, BB, B, CCC, CC and D for Standard & Poor's for example) is small, so companies with ratings are divided into only nine groups by major ratings groups and 21 groups if the pluses and minuses are added to ratings from AA to CCC. Market participants need a much higher degree of precision than the rating alone allows, even giving credit for the pluses and minuses attached to each rating.
- The *annual default rates by rating are very volatile* over time, so there is considerable uncertainty about what the default probability is that is associated with each rating.
- Ratings are an *ordinal measure of credit risk*, rather than a quantitative default probability.
- *Ratings change very infrequently*, so market participants are getting a view of an issuer's credit that changes only every six to 12 months at best. General Electric's rating did not change between 1956 and 2009, a period of 53 years.
- Ratings take the long "through-the-cycle view," so *they show relatively little of the cyclicity* that many analysts believe is the key to credit risk management.
- The *ratings universe is a very small percentage* of the total number of counterparties that a typical financial institution would lend to (retail clients, small business clients, and corporate clients totaling in the millions at some large banks), so an institution relying on ratings as its primary credit risk management tool is forced by definition to take a piecemeal approach to enterprise credit risk management. Each class of counterparties ends up having a different type of rating, some internal and some external.
- Rating agencies are normally paid by the structurer of a complex CDO or the issuer of the securities, so there is substantial *commercial pressure* on the agencies for a "good" rating. The report by Senator Levin's committee summarizes these problems in detail.
- Rating agencies have a substantial "nonrating" consulting business, so they are increasingly being scrutinized as having *potential conflicts of interest* like the major accounting firms before they spun off their consulting arms in the past decade.
- The *accuracy of legacy ratings has been consistently found to be much lower* than default probabilities developed using logistic regression and the reduced form approach.

Each of these issues could fill a book in and of itself. We deal with the most serious issues in this section.

THROUGH THE CYCLE VS. POINT IN TIME, A DISTINCTION WITHOUT A DIFFERENCE

There has been a long debate in the risk management industry about whether legacy credit ratings are intended to be “point-in-time” ratings or whether they are intended to be “through-the-cycle” credit indicators. The reason that such a debate occurs is the vagueness of the maturity and default probabilities associated with a given rating for a specific company at any point in time. When viewed in the context of a modern quantitative default probability, the debate about point in time versus through the cycle is a distinction without a difference. This section explains the reason why.

Even after more than 100 years of experience with traditional credit ratings, practical credit portfolio managers still do not have clear answers to a number of critical questions:

- What is the default probability associated with a specific rating for a specific company over a given time horizon on a specific date?
- What is the default probability associated with the same issuer over different time horizons?
- How do the ratings and default probabilities associated with them change when macroeconomic factors change?
- Put differently, how does the default probability associated with a given rating change as the company moves through various points in the business cycle?

Because the legacy rating agencies have been unable to articulate clear answers to these questions, two other questions are often asked and debated:

- Is a given credit rating a point-in-time rating, an index of risk that prevails today but (by implication) not necessarily tomorrow?
- Is a given credit rating a through-the-cycle rating intended to measure the average level of credit risk for a given company over an unspecified but presumably long period of time that extends through the cycle both for the current business cycle and presumably many others?

A precise answer to these questions only becomes more muddled if one poses it to senior people at the rating agencies. In February 2003 at a well-lubricated dinner in midtown Manhattan, one of the authors posed this question to the head of corporate credit ratings at one of the two leading rating agencies: “What would your agency do with respect to a company’s rating if the company was a long-run risk level of BBB/Baa, but faced 60 days of CCC/Caa risk?” His answer, consistent with the benign credit conditions of the time, was this: “We cannot afford to damage our firm’s reputation by letting a BBB/Baa-rated company default, so we would set the rating at CCC/Caa.” This comment contrasts with more recent statements by other senior rating executives who argue that they have a special obligation to avoid “type 1” error that overly harshly rates companies. The reason for this concern is that derivatives-related margin requirements, unfortunately, have been tied to ratings and the legacy rating agencies fear an error that would trigger a margin call that leads to bankruptcy. If the first quote is accurate in reflecting rating agency thinking, it means

that a rating on any given date reflects the worse of (1) the short-term risk of the company or (2) the long-run risk of the company. From this view, the rating is neither point in time nor through the cycle. A little reflection in a default modeling context makes the muddled semantics much more clear and transparent.

How does a modern quantitative default model differ from ratings? The differences are very great and very attractive from a practice use point of view:

- Each default probability has an explicit maturity.
- Each default probability has an obvious meaning. A default probability of 4.25 percent for a three-year maturity means what it says: there is a 4.25 percent (annualized) probability of default by this company over the three-year period starting today. Cumulative default risk is calculated using the formulas in the appendix of Chapter 16.
- Each company has a full-term structure of default probabilities at maturities from one month to 10 years, updated daily, in a widely used default probability service.

What does “point in time” mean in this context? We illustrate the issues with the default probabilities for Hewlett Packard on June 14, 2012, as displayed by Kamakura Risk Information Services in Exhibit 18.1.

All of the default probabilities for Hewlett Packard (HP) on June 14, 2012, are default probabilities that are point in time for HP. Default probabilities at a different point in time, say June 15, will change if the inputs to the default models are different on June 15 than they were on June 14. What does “through the cycle” mean with respect to the default probabilities for HP on June 14? Through the cycle implies the longest default probability available on the term structure of default probabilities, because this maturity does the best job of extending through as much of the business cycle as possible. For Kamakura Risk Information Services (KRIS) version 5, the longest default probability available is the 10-year default probability.




EXHIBIT 18.1 Hewlett Packard Co.

If the default probability for Hewlett Packard on June 14 is 0.38 percent at 10 years, it means that the through-the-cycle default probability for HP prevailing on June 14, 2012, is a 0.38 percent (annualized) default rate over the 10 years ending June 14, 2022. What could be clearer than that?

To summarize, all of the default probabilities prevailing for HP on June 14, 2012, are the point-in-time default probabilities for HP at all maturities from one month to 10 years. The through-the-cycle default probability for HP on June 14 is the 10-year default probability, because this is the longest maturity available. The 10-year default probability is also obviously a point-in-time default probability because it prevails on June 14, the point in time we care about. On June 15, all of the point-in-time default probabilities for HP will be updated, including the 10-year default probability, which has a dual role as the through-the-cycle default probability.

There is no uncertainty about these concepts: all default probabilities that exist today at all maturities are point-in-time default probabilities for HP, and the longest maturity default probability is also the through-the-cycle default probability.

How can these default probabilities be mapped to ratings, and what rating would be point in time and what rating would be through the cycle? An experienced user of quantitative default probabilities would ask in return, “Why would you want to go from an explicit credit risk assessment with a known maturity and 10,000 grades (from 0 basis points to 10,000 basis points) to a vague credit assessment with no known maturity and only 20 grades?” A common answer is that management is used to ratings and ratings have to be produced, even if they’re much less accurate and much less useful than the default probabilities themselves.

Model Risk Alert

Most of the tens of billions of losses incurred by investors in collateralized debt obligations (which we discuss in Chapter 20) during the 2006 to 2011 credit crisis were explained after the fact by the statement “But it was rated AAA when we bought it.”

STRESS TESTING, LEGACY RATINGS, AND TRANSITION MATRICES

In the wake of the 2006 to 2011 credit crisis, financial institutions, regulators have implemented a wide variety of stress tests that require financial institutions to calculate the market value of assets, liabilities, and capital in different economic scenarios. Many of the regulations also require projections of net income under the accounting standard relevant for that institution. When ratings, instead of default probabilities, are at the heart of the credit risk process, institutions are literally unable to do accurate calculations required by regulators because the rating agencies themselves are unable to articulate the quantitative links between macroeconomic factors, legacy ratings, and probabilities of default. We spend a lot of time on these links in later chapters. These stress tests contrast heavily with the much-criticized reliance on ratings in the Basel II Capital Accords. The Dodd-Frank legislation in the United States in 2010 has hastened the inevitable demise of the rating agencies by requiring U.S. government agencies to remove any rules or requirements that demand the use of legacy ratings.

TRANSITION MATRICES: ANALYZING THE RANDOM CHANGES IN RATINGS FROM ONE LEVEL TO ANOTHER

It goes without saying that the opaqueness of the link between ratings, macro factors, and default risk makes transition matrices an idea useful in concept but useless in practice. The *transition matrix* is the probability that a firm moves from ratings class J to ratings class K over a specific period of time. Most users of the transition matrix concept make the assumption—a *model risk alert*—that the transition probabilities are constant. (Chapter 16 illustrates clearly that this assumption is false in its graph of the number of bankruptcies in the United States from 1990 to 2008.) Default risk of all firms changes in a highly correlated fashion. The very nature of the phrase “business cycle” conveys the meaning that most default probabilities rise when times are bad and fall when times are good. This means that transition probabilities to lower ratings should rise in bad times and fall in good times. Since the very nature of the ratings process is nontransparent, there is no valid basis for calculating these transition probabilities that matches the reduced form approach in accuracy.

MORAL HAZARD IN “SELF-ASSESSMENT” OF RATINGS ACCURACY BY LEGACY RATING AGENCIES

As we noted in the introduction to this chapter, the choice of credit models is not a beauty contest. It is all about accuracy and nothing else (except when corporate and regulatory politics interferes with the facts). One of the challenges facing an analyst is the moral hazard of accuracy “self-assessments” by the legacy rating agencies, a topic to which we now turn.

On March 9, 2009, one of the authors and Robert Jarrow wrote a blog entitled “The Rating Chernobyl,” predicting that rating agency errors during the ongoing credit crisis were so serious that even their own self-assessments would show once and for all that legacy ratings are a hopelessly flawed measure of default risk.¹ That 2009 prediction has not come true for a very simple reason: The rating agencies have left their mistakes out of their self-assessments. This section, nevertheless, illustrates the dangers of relying on rating agency default rates in credit risk management, because the numbers are simply not credible. We also predicted that the ratings errors during the credit crisis were so egregious that even the rating agencies’ own annual self-assessments of ratings performance would show that an analyst of credit risk should not place any serious reliance on legacy ratings.²

In the wake of the financial crisis, rating agencies were forced to make available their ratings for public firms for the years 2007–2009. We use those disclosed ratings from Standard & Poor’s in what follows. Like the rating agency self-assessments, we analyze firms’ subsequent default rates based on their legacy ratings on January 1 of each year. If we use all of the data released by Standard & Poor’s, we have 6,664 observations for the three-year period, using the January 1 rating in each case for each firm (Exhibit 18.2):

How many firms failed during this three-year period among firms with legacy ratings? To answer that question, we first use the KRIS database and then compare

EXHIBIT 18.2 Analysis of 2007–2009 Default Rates by Rating Grade

Ratings Rank	Rating	Number of Observations			
		2007	2008	2009	Total Observations
1	AAA	17	15	12	44
2	AA+	8	8	8	24
3	AA	35	65	48	148
4	AA–	101	67	72	240
5	A+	120	113	104	337
6	A	193	172	172	537
7	A–	238	202	202	642
8	BBB+	272	253	236	761
9	BBB	306	255	273	834
10	BBB–	207	213	225	645
11	BB+	149	132	121	402
12	BB	186	147	138	471
13	BB–	201	196	183	580
14	B+	161	153	138	452
15	B	92	104	101	297
16	B–	52	44	61	157
17	CCC+	18	15	19	52
18	CCC	10	6	11	27
19	CCC–	1	1	2	4
20	CC	2	3	5	10
Grand Total		2,369	2,164	2,131	6,664

Sources: Kamakura Corporation; Standard & Poor's.

with results from Standard & Poor's. The KRIS definition of failure includes the following cases:

- A D or uncured SD rating from a legacy rating agency
- An ISDA event of default
- A delisting for obvious financial reasons, such as failure to file financial statements, failure to pay exchange fees, or failure to maintain the minimum required stock price level

The KRIS database very carefully distinguishes failures from rescues. Bear Stearns, for example, failed. After the failure, Bear Stearns was rescued by JPMorgan Chase in return for massive U.S. government assistance.³

The Kamakura failure count includes all firms that would have failed without government assistance. In many of the government rescues, common shareholders and preferred shareholders would be amused to hear that the rescued firm did not default just because the senior debt holders were bailed out. As the Kamakura Case Studies in Liquidity Risk series makes clear, the government's

definition of “too big to fail” changed within two or three days of the failure of Lehman Brothers, so the prediction of a probability of rescue is a much more difficult task than the prediction of failure. Using the KRIS database of failures in the 2007–2009, we have 86 failed firms during this period. This count of failed firms is being revised upward by Kamakura in KRIS version 6 in light of recent government documents (Exhibit 18.3).

We can calculate the failure rate in each year and the weighted average three-year failure rate by simply dividing the number of failures by the number of observations. We do that and display the results in graphical form from AAA to B–rating (Exhibit 18.4).

The results are consistent with Jarrow and van Deventer’s “The Ratings Chernobyl.” The three-year weighted average failure rate for AAA-rated firms was 4.55 percent, much higher than any other failure rate except B– ratings in the B to AAA range. By contrast, BBB firms failed at only a 0.12 percent rate and BB+ firms failed at only a 0.25 percent rate. In tabular form the results can be summarized as shown in Exhibit 18.5.

By any statistical test over the B to AAA range, legacy ratings were nearly worthless in distinguishing strong firms from weak firms. How does the Standard & Poor’s self-assessment present these results? It doesn’t. See Exhibit 18.6 for what it shows.

Note that the first-year default rate for AAA-rated companies is zero. What happened to FNMA and FHLMC? And what about the 2007–2009 failure rate of 1.19 percent for A+ weighted companies? Why is the A-rated first-year default rate so small in light of that? The reason is simple. A large number of important failed companies have been omitted from the self-assessment, avoiding a very embarrassing self-assessment (Exhibit 18.7).

Which firms did Kamakura have listed as failures in the ratings grades from A to AAA?

Kamakura included these 11 firms:

AAA	Federal National Mortgage Association (FNMA or Fannie Mae)
AAA	Federal Home Loan Mortgage Corporation (FHLMC or Freddie Mac)
AA	American International Group
AA–	Wachovia Corporation
A+	Northern Rock PLC
A+	Merrill Lynch & Co Inc.
A+	Lehman Brothers Holdings Inc.
A+	Ageas SA/NV
A	Bear Stearns Companies Inc.
A–	Washington Mutual Inc.
A–	Anglo Irish Bank Corporation

Remember also that American International Group was rated AAA in January 2005. Which firms were omitted from the Standard & Poor’s list? Nine of the 11 failures tallied by Kamakura were omitted by S&P from their self-assessment. (Some, of course, debate whether Fannie Mae and Freddie Mac were in fact failures.⁴)

EXHIBIT 18.3 Analysis of 2007–2009 Default Rates by Rating Grade

Ratings Rank	Rating	Number of Defaults			Total Defaults
		2007	2008	2009	
1	AAA	0	2	0	2
2	AA+	0	0	0	0
3	AA	0	1	0	1
4	AA–	0	1	0	1
5	A+	1	3	0	4
6	A	0	1	0	1
7	A–	0	1	1	2
8	BBB+	0	0	2	2
9	BBB	0	1	0	1
10	BBB–	0	2	3	5
11	BB+	0	1	0	1
12	BB	1	1	1	3
13	BB–	0	2	0	2
14	B+	0	4	2	6
15	B	0	2	10	12
16	B–	1	3	11	15
17	CCC+	1	1	8	10
18	CCC	2	3	7	12
19	CCC–	0	0	1	1
20	CC	1	1	3	5
	Grand Total	7	30	49	86

Sources: Kamakura Corporation; Standard & Poor's.

When Standard & Poor's was asked by a regional bank in the United States why FNMA and FHLMC had been omitted from the corporate ratings self-assessment, the response from Standard & Poor's did not discuss the nuances of a rescue of a failed firm. Instead, the e-mail reply stated only that FNMA and FHLMC were no longer exchange-listed, privately owned corporations and that they were therefore no longer an object of this study. This justification, obviously, could be used to exclude every firm that went bankrupt from the legacy ratings accuracy self-assessment.

A complete third-party audit of the entire rating agency default history is mandatory before historical rating agency default rates can be relied upon with confidence. Kamakura Risk Information Services maintains an independent and active database of corporate failures (for information, contact info@kamakuraco.com).

COMPARING THE ACCURACY OF RATINGS AND REDUCED FORM DEFAULT PROBABILITIES

In a recent paper, Jens Hilscher of Brandeis University and senior research fellow at Kamakura Corporation and Mungo Wilson of Oxford University (2011) compared




EXHIBIT 18.4 Three-Year Failure Rate by Legacy Rating Grade, 2007–2009

Source: Kamakura Corporation; Standard & Poor's.

the accuracy of legacy credit ratings with modern default probabilities. The authors concluded the following:

This paper investigates the information in corporate credit ratings. We examine the extent to which firms' credit ratings measure raw probability of default as opposed to systematic risk of default, a firm's tendency to default in bad times. We find that credit ratings are dominated as predictors of corporate failure by a simple model based on publicly available financial information ("failure score"), indicating that ratings are poor measures of raw default probability.

We encourage the serious reader to review the full paper to understand the details of this and related conclusions.

The Hilscher and Wilson results are consistent with the Kamakura Risk Information Services Technical Guides (versions 2, 3, 4, and 5), which have been released in sequence beginning in 2002. Version 5 of the Technical Guide (Jarrow, Klein, Mesler, and van Deventer, March 2011) includes ROC accuracy ratio comparisons for legacy ratings and the version 5 Jarrow-Chava default probabilities (discussed in Chapter 16). The testing universe was much smaller than the 1.76 million observations and 2,046 company failures in the full KRIS version 5 sample. The rated universe included only 285,000 observations and 276 company failures. The results in Exhibit 18.8 show that the ROC accuracy ratio for the version 5 Jarrow-Chava

EXHIBIT 18.5 Analysis of 2007–2009 Default Rates by Rating Grade

Ratings Rank	Rating	Default Rate by Year			
		2007	2008	2009	3-Year Default Rate
1	AAA	0.00%	13.33%	0.00%	4.55%
2	AA+	0.00%	0.00%	0.00%	0.00%
3	AA	0.00%	1.54%	0.00%	0.68%
4	AA–	0.00%	1.49%	0.00%	0.42%
5	A+	0.83%	2.65%	0.00%	1.19%
6	A	0.00%	0.58%	0.00%	0.19%
7	A–	0.00%	0.50%	0.50%	0.31%
8	BBB+	0.00%	0.00%	0.85%	0.26%
9	BBB	0.00%	0.39%	0.00%	0.12%
10	BBB–	0.00%	0.94%	1.33%	0.78%
11	BB+	0.00%	0.76%	0.00%	0.25%
12	BB	0.54%	0.68%	0.72%	0.64%
13	BB–	0.00%	1.02%	0.00%	0.34%
14	B+	0.00%	2.61%	1.45%	1.33%
15	B	0.00%	1.92%	9.90%	4.04%
16	B–	1.92%	6.82%	18.03%	9.55%
17	CCC+	5.56%	6.67%	42.11%	19.23%
18	CCC	20.00%	50.00%	63.64%	44.44%
19	CCC–	0.00%	0.00%	50.00%	25.00%
20	CC	50.00%	33.33%	60.00%	50.00%
	Grand Total	0.30%	1.39%	2.30%	1.29%

Sources: Kamakura Corporation; Standard & Poor's.

model is 7 to 10 percentage points more accurate than legacy ratings at every single time horizon studied. Note that the accuracy reported for month N is the accuracy of predicting failure in month N of those companies who survived the period from period 1 through period $N - 1$.

PROBLEMS WITH LEGACY RATINGS IN THE 2006 TO 2011 CREDIT CRISIS

One of the reasons for the moral hazard in accuracy self-assessment by the rating agencies is the conflicting pressures on management of the agencies themselves. Management is put in the situation where the revenue growth demanded by shareholders of the firm in the short run was in conflict with the need for accuracy in ratings to preserve the “ratings franchise of the firm,” in the words of our rating agency friend mentioned previously. This conflict was examined in detail by the U.S. Senate in the wake of the credit crisis in a very important research report led by Senator Carl Levin (Exhibit 18.9).

EXHIBIT 18.6 Global Corporate Average Cumulative Default Rates (1981–2010)

Rating	Time Horizon (years)														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
AAA	0.00 (0.00)	0.03 (0.20)	0.14 (0.39)	0.26 (0.47)	0.38 (0.59)	0.50 (0.69)	0.56 (0.75)	0.66 (0.82)	0.72 (0.83)	0.79 (0.83)	0.83 (0.83)	0.87 (0.84)	0.91 (0.84)	1.00 (0.91)	1.09 (0.99)
AA	0.02 (0.08)	0.07 (0.12)	0.15 (0.19)	0.26 (0.28)	0.37 (0.36)	0.49 (0.47)	0.58 (0.56)	0.67 (0.63)	0.74 (0.68)	0.82 (0.72)	0.90 (0.72)	0.97 (0.73)	1.04 (0.70)	1.10 (0.70)	1.15 (0.71)
A	0.08 (0.12)	0.19 (0.21)	0.33 (0.27)	0.50 (0.36)	0.68 (0.44)	0.89 (0.48)	1.15 (0.54)	1.37 (0.59)	1.60 (0.67)	1.84 (0.77)	2.05 (0.86)	2.23 (0.91)	2.40 (0.88)	2.55 (0.87)	2.77 (0.83)
BBB	0.25 (0.27)	0.70 (0.60)	1.19 (0.88)	1.80 (1.09)	2.43 (1.32)	3.05 (1.48)	3.59 (1.59)	4.14 (1.61)	4.68 (1.64)	5.22 (1.57)	5.78 (1.40)	6.24 (1.33)	6.72 (1.19)	7.21 (1.04)	7.71 (1.00)
BB	0.95 (1.05)	2.83 (2.32)	5.03 (3.39)	7.14 (4.08)	9.04 (4.64)	10.87 (4.87)	12.48 (4.78)	13.97 (4.61)	15.35 (4.43)	16.54 (4.24)	17.52 (4.43)	18.39 (4.50)	19.14 (4.56)	19.78 (4.51)	20.52 (4.63)
B	4.70 (3.31)	10.40 (5.69)	15.22 (6.93)	18.98 (7.70)	21.76 (8.10)	23.99 (7.87)	25.82 (7.33)	27.32 (6.84)	28.64 (6.27)	29.94 (5.97)	31.09 (5.58)	32.02 (5.35)	32.89 (4.55)	33.70 (3.91)	34.54 (3.89)
CCC/C	27.39 (12.69)	36.79 (13.97)	42.12 (13.61)	45.21 (14.09)	47.64 (14.05)	48.72 (12.98)	49.72 (12.70)	50.61 (12.10)	51.88 (11.65)	52.88 (10.47)	53.71 (10.75)	54.64 (11.42)	55.67 (12.06)	56.55 (10.38)	56.55 (9.61)
Investment grade	0.13 (0.12)	0.34 (0.28)	0.59 (0.41)	0.89 (0.52)	1.21 (0.61)	1.53 (0.66)	1.83 (0.69)	2.12 (0.72)	2.39 (0.79)	2.68 (0.84)	2.94 (0.86)	3.16 (0.85)	3.37 (0.78)	3.59 (0.68)	3.83 (0.65)
Speculative grade	4.36 (2.80)	8.53 (4.55)	12.17 (5.61)	15.13 (6.18)	17.48 (6.41)	19.45 (6.12)	21.13 (5.61)	22.59 (4.98)	23.93 (4.34)	25.16 (4.07)	26.21 (3.94)	27.10 (3.95)	27.93 (3.72)	28.66 (3.53)	29.40 (3.57)
All rated	1.61 (1.06)	3.19 (1.85)	4.60 (2.40)	5.80 (2.73)	6.79 (2.88)	7.64 (2.82)	8.38 (2.68)	9.02 (2.54)	9.62 (2.43)	10.18 (2.39)	10.67 (2.28)	11.08 (2.15)	11.47 (1.91)	11.82 (1.90)	12.20 (2.08)

Numbers in parentheses are standard deviations. Sources: Standard & Poor's Global Fixed Income Research and Standard & Poor's CreditPro®.

EXHIBIT 18.7 Investment Grade Defaults in the Five-Year 2006 Static Pool

Company	Country	Industry	Default Date	Next-to- Last Rating	Date of Next-to- Last Rating	First Rating	Date of First Rating	Year of Default
Aiful Corp.	Japan	Financial institutions	9/24/2009	CC	9/18/2009	BBB	10/6/2003	2009
Ambac Assurance Corp.	U.S.	Insurance	11/18/2009	CC	7/28/2009	AAA	12/31/1980	2009
Ambac Financial Group, Inc.	U.S.	Insurance	11/2/2010	CC	6/9/2008	AA+	7/30/1991	2010
BluePoint Re Limited	Bermuda	Insurance	8/14/2008	A	6/9/2008	AA	10/25/2004	2008
Caesars Entertainment Corp.	U.S.	Leisure time/media	12/24/2008	CC	11/18/2008	A	12/31/1980	2008
CIT Group, Inc.	U.S.	Financial institutions	8/17/2009	CC	7/16/2009	AA	12/31/1980	2009
Clear Channel Communications Inc.	U.S.	Leisure time/media	12/23/2008	CC	12/5/2008	BBB-	9/26/1997	2008
Colonial BancGroup Inc.	U.S.	Financial institutions	8/17/2009	CC	7/30/2009	BBB-	1/17/1997	2009
Colonial Bank	U.S.	Financial institutions	8/17/2009	CCC-	7/30/2009	BBB	1/21/1997	2009
Commonwealth Land Title Insurance Co.	U.S.	Insurance	12/4/2008	BB-	11/24/2008	A-	6/25/1997	2008
Controladora Comercial Mexicana, S. A. B. de C.V.	Mexico	Consumer/service sector	10/9/2008	CC	10/8/2008	BB+	3/31/1998	2008
Downey Financial Corp.	U.S.	Financial institutions	11/24/2008	CCC-	11/21/2008	BBB-	6/7/1999	2008
Downey S&L Assn	U.S.	Financial institutions	11/24/2008	CCC	11/21/2008	A+	12/31/1980	2008
Energy Future Holdings Corp.	U.S.	Energy and natural resources	11/16/2009	CC	10/5/2009	BBB	10/3/1997	2009

(Continued)

EXHIBIT 18.7 (Continued)

Company	Country	Industry	Default Date	Next-to- Last Rating	Date of Next-to- Last Rating	First Rating	Date of First Rating	Year of Default
FGIC Corp	U.S.	Insurance	8/3/2010	CC	9/24/2008	AA	1/5/2004	2010
General Growth Properties, Inc.	U.S.	Real estate	3/17/2009	CC	12/24/2008	BBB-	6/2/1998	2009
IndyMac Bank, FSB	U.S.	Financial institutions	7/14/2008	B-	7/9/2008	BBB-	9/4/1998	2008
LandAmerica Financial Group Inc.	U.S.	Insurance	11/26/2008	B-	11/24/2008	BBB-	11/19/2004	2008
Lehman Brothers Holdings Inc.	U.S.	Financial institutions	9/16/2008	A	6/2/2008	AA-	1/1/1985	2008
Lehman Brothers Inc.	U.S.	Financial institutions	9/23/2008	BB-	9/15/2008	AA	10/5/1984	2008
Mashantucket Western Pequot Tribe	U.S.	Leisure time/media	11/16/2009	CCC	8/26/2009	BBB-	9/16/1999	2009
McClatchy Co. (The)	U.S.	Leisure time/media	6/29/2009	CC	5/22/2009	BBB-	2/8/2000	2009
Residential Capital, LLC	U.S.	Financial institutions	6/4/2008	CC	5/2/2008	BBB-	6/9/2005	2008
Sabre Holdings Corp.	U.S.	Transportation	6/16/2009	B	3/30/2009	A-	2/7/2000	2009
Scottish Annuity & Life Insurance Co. (Cayman) Ltd.	Cayman Islands	Insurance	1/30/2009	CC	1/9/2009	A-	12/18/2001	2009
Takefuji Corp.	Japan	Financial institutions	12/15/2009	CC	11/17/2009	A-	2/10/1999	2009
Technicolor S.A.	France	Consumer/service sector	5/7/2009	CC	1/29/2009	BBB+	7/24/2002	2009
Tribune Co.	U.S.	Leisure time/media	12/9/2008	CCC	11/11/2008	AA	3/1/1983	2008
Washington Mutual Bank	U.S.	Financial institutions	9/26/2008	BBB-	9/15/2008	B+	1/24/1989	2008
Washington Mutual, Inc.	U.S.	Financial institutions	9/26/2008	CCC	9/15/2008	BBB	7/17/1995	2008
YRC Worldwide Inc.	U.S.	Transportation	1/4/2010	CC	11/2/2009	BBB-	11/19/2003	2010




EXHIBIT 18.8 Kamakura Corporation, ROC Accuracy Ratio in Month N Conditional on Surviving Months 1 through $N - 1$ for KRIS Jarrow-Chava Version 5 vs. Legacy Ratings, January 1990 to December 2008

We list the conclusions of the Levin Report with respect to the rating agencies in this excerpt, which begins on page 245:

A. Subcommittee Investigation and Findings of Fact

For more than one year, the Subcommittee conducted an in-depth investigation of the role of credit rating agencies in the financial crisis, using as case histories Moody's and S&P. The Subcommittee subpoenaed and reviewed hundreds of thousands of documents from both companies including reports, analyses, memoranda, correspondence, and email, as well as documents from a number of financial institutions that obtained ratings for RMBS and CDO securities. The Subcommittee also collected and reviewed documents from the SEC and reports produced by academics and government agencies on credit rating issues. In addition, the Subcommittee conducted nearly two dozen interviews with current and former Moody's and S&P executives, managers, and analysts, and consulted with credit rating experts from the SEC, Federal Reserve, academia, and industry. On April 23, 2010, the Subcommittee held a hearing and released 100 hearing exhibits. In connection with the hearing, the Subcommittee released a joint memorandum from Chairman Levin and Ranking Member Coburn summarizing the investigation into the credit rating agencies and the problems with the credit ratings assigned to RMBS and CDO securities.

United States Senate

PERMANENT SUBCOMMITTEE ON INVESTIGATIONS

Committee on Homeland Security and Governmental Affairs

Carl Levin, Chairman

Tom Coburn, Ranking Minority Member

WALL STREET AND THE FINANCIAL CRISIS: Anatomy of a Financial Collapse

**MAJORITY AND MINORITY
STAFF REPORT**

**PERMANENT SUBCOMMITTEE
ON INVESTIGATIONS**

UNITED STATES SENATE



April 13, 2011

EXHIBIT 18.9 U.S. Senate Permanent Subcommittee on Investigations Majority and Minority Staff Report

The memorandum contained joint findings regarding the role of the credit rating agencies in the Moody's and S&P case histories, which this Report reaffirms. The findings of fact are as follows.

1. **Inaccurate Rating Models.** From 2004 to 2007, Moody's and S&P used credit rating models with data that was inadequate to predict how high-risk residential mortgages, such as subprime, interest only, and option adjustable rate mortgages, would perform.

2. **Competitive Pressures.** Competitive pressures, including the drive for market share and the need to accommodate investment bankers bringing in business, affected the credit ratings issued by Moody's and S&P.
3. **Failure to Re-evaluate.** By 2006, Moody's and S&P knew their ratings of RMBS and CDOs were inaccurate, revised their rating models to produce more accurate ratings, but then failed to use the revised model to reevaluate existing RMBS and CDO securities, delaying thousands of rating downgrades and allowing those securities to carry inflated ratings that could mislead investors.
4. **Failure to Factor in Fraud, Laxity, or Housing Bubble.** From 2004 to 2007, Moody's and S&P knew of increased credit risks due to mortgage fraud, lax underwriting standards, and unsustainable housing price appreciation, but failed adequately to incorporate those factors into their credit rating models.
5. **Inadequate Resources.** Despite record profits from 2004 to 2007, Moody's and S&P failed to assign sufficient resources to adequately rate new products and test the accuracy of existing ratings.
6. **Mass Downgrades Shocked Market.** Mass downgrades by Moody's and S&P, including downgrades of hundreds of subprime RMBS over a few days in July 2007, downgrades by Moody's of CDOs in October 2007, and actions taken (including downgrading and placing securities on credit watch with negative implications) by S&P on over 6,300 RMBS and 1,900 CDOs on one day in January 2008, shocked the financial markets, helped cause the collapse of the subprime secondary market, triggered sales of assets that had lost investment grade status, and damaged holdings of financial firms worldwide, contributing to the financial crisis.
7. **Failed Ratings.** Moody's and S&P each rated more than 10,000 RMBS securities from 2006 to 2007, downgraded a substantial number within a year, and, by 2010, had downgraded many AAA ratings to junk status.
8. **Statutory Bar.** The SEC is barred by statute from conducting needed oversight into the substance, procedures, and methodologies of the credit rating models.
9. **Legal Pressure for AAA Ratings.** Legal requirements that some regulated entities, such as banks, broker-dealers, insurance companies, pension funds, and others, hold assets with AAA or investment grade credit ratings, created pressure on credit rating agencies to issue inflated ratings making assets eligible for purchase by those entities.

These "findings of fact" summarize briefly the reasons why legacy ratings have been called into question generally, even though many of the specific comments above refer to ratings of mortgage-backed securities and collateralized debt obligations, which we discuss later in this chapter.

THE JARROW-MERTON PUT OPTION AND LEGACY RATINGS

We end this section by noting the big picture objectives of Chapters 1 and 2. We want to measure market risk, interest rate risk, liquidity risk, and credit risk on an integrated basis and we want to be able to do something about it if we are uncomfortable with the level of risk we have. We want to be able to compute the value of the Jarrow-Merton put option (say in one year at the par value of the liabilities of the financial

institution) on the value of company assets that is the best coherent estimate of the risk the institution has taken on.

In simpler language, we want to know “what’s the hedge” if our risk is too large. The use of credit ratings leaves us short of this goal line. The link between ratings and pricing and, valuation and hedging is very tenuous in most financial institutions. We now turn to another approach that once seemed promising but which has been exposed as inadequate in the light of the 2006 to 2011 credit crisis: the Merton model of risky debt.

THE MERTON MODEL OF RISKY DEBT

In Geneva, in December 2002, Robert Merton was speaking to a large conference audience of 400 risk managers when a question arose from the back of the room. “Professor Merton,” the questioner started, “what else do you recommend to financial institutions using commercial versions of your credit model?” After a long pause, Merton said, “Well, the first thing you have to remember is that the model is 28 years old.” This is both the strength and the weakness of the Merton model. As it ages it is better understood. At the same time, more extensive credit model testing like that in Chapter 16 have made it apparent that the model simply has not performed well from an accuracy point of view. This view is shared by Merton himself, who was one of a panel of four Nobel Prize winners who awarded the Harry Markowitz Award for best paper of 2011 to Campbell, Szilagyi, and Hilscher (2011) for their paper in the *Journal of Investment Management*, which concluded the following about their reduced form credit model: “Our model is . . . between 49 percent and 94 percent more accurate than [a Merton distance to default model].”

Many authors have examined this topic and reached the same conclusions that we reached in Chapter 16 about the superior performance of reduced form models versus the Merton model. The Chapter 16 results were originally reported in Kamakura Risk Information Services Technical Guides 2, 3, 4, and 5 released between 2002 and 2010. Most of the following papers have been in circulation in banking and academic circles for more than a decade:

- S. Bharath and T. Shumway. “Forecasting Default with the Merton Distance to Default Model.” *Review of Financial Studies* 21 (2008): 1339–1369.
- J. Y. Campbell, J. Hilscher, and J. Szilagyi. “In Search of Distress Risk.” *Journal of Finance* 63 (2008): 2899–2939.
- S. Chava and R. A. Jarrow. “Bankruptcy Prediction with Industry Effects.” *Review of Finance* 8 (2004): 537–569.
- J. Hilscher and M. Wilson. “Credit Ratings and Credit Risk.” Unpublished paper, Brandeis University and Oxford University, 2011.
- T. Shumway. “Forecasting Bankruptcy More Accurately: A Simple Hazard Model.” *Journal of Business* 74 (2001): 101–124.

Shumway and Bharath (2008) conducted an extensive test of the Merton approach. They tested two hypotheses. Hypothesis 1 was that the Merton model is a “sufficient statistic” for the probability of default; that is, a variable so powerful that in a logistic regression like the formula in the previous section no other explanatory

variables add explanatory power. Hypothesis 2 was the hypothesis that the Merton model adds explanatory power even if common reduced form model explanatory variables are present. They specifically tested modifications of the Merton structure partially disclosed by commercial vendors of the Merton model. The Shumway and Bharath (2008) conclusions, based on all publicly traded firms in the United States (except financial firms) using quarterly data from 1980 to 2003 are as follows:

We conclude that the . . . Merton model does not produce a sufficient statistic for the probability of default

Models 6 and 7 include a number of other covariates: the firm's returns over the past year, the log of the firm's debt, the inverse of the firm's equity volatility, and the firm's ratio of net income to total assets. Each of these predictors is statistically significant, making our rejection of hypothesis one quite robust. Interestingly, with all of these predictors included in the hazard model, the . . . Merton probability is no longer statistically significant, implying that we can reject hypothesis two

Looking at CDS implied default probability regressions and bond yield spread regressions, the . . . Merton probability does not appear to be a significant predictor of either quantity when our naïve probability, agency ratings and other explanatory variables are accounted for.

The Campbell, Hilscher, and Szilagyi (2008) paper was done independently on a separate monthly data set and reaches similar conclusions, all consistent with the results in Chapter 16.

Any rational analyst who has read these papers and replicated their tests on a large data set (more than 500,000 monthly observations is recommended) will reach the same conclusions and would issue a Model Risk Alert because the Merton model alone is simply not an accurate tool for credit risk assessment.

Why then include the Merton model in this chapter? First of all, this is a chapter on credit risk management antiques, and the Merton model is an antique that held the affections of one of the authors of this book for more than a decade. That author was ultimately convinced by this comment by John Maynard Keynes (quoted by Malabre) that he should turn to other models: "When the facts change, I change my mind. What do you do, sir?"

That being said, the intuition of the Merton model is attractive and sometimes useful in thinking about credit risk issues, even though the default probabilities derived from the model are much less accurate than the reduced form approach. For that reason, we summarize the attributes of the model and its derivation here.

THE INTUITION OF THE MERTON MODEL

The intuition of the Merton model is dramatically appealing, particularly to one of the authors, a former treasurer at a major bank holding company who used the model's insights in thinking about financial strategy. Merton (1974) was a pioneer in options research at MIT at the time that Black and Scholes (1973) had just published their now well-known options model. One of the reasons for the enduring popularity of the Merton credit model is the even greater popularity of the Black-Scholes model.

The Black-Scholes model is simple enough that it can be modeled in popular spreadsheet software. Because of this simplicity, Hull (1993) and Jarrow and Turnbull (1996) note with chagrin that traders in interest rate caps and floors (a Model Risk Alert that we issue in later chapters) frequently use the related Black commodity options model, even though it implicitly assumes interest rates are constant. In the same way, extensive use of the Merton model has persisted even though it is now clear that other modeling techniques for credit risk are much more effective.

Merton's great insight was to recognize that the equity in a company can be considered an option on the assets of the firm. More formally, he postulated that the assumptions that Black and Scholes applied in their stock options model can be applied to the assets of corporations. Merton (1974) assumes that the assets of corporations are completely liquid and traded in frictionless markets.

Model Risk Alerts

Cash is liquid, but most other assets are not, invalidating many of the model's conclusions. Merton assumed that interest rates are (Model Risk Alert) constant, and that the firm has only (Model Risk Alert) one zero-coupon debt issue outstanding. He assumed, like Black and Scholes, a (Model Risk Alert) one-period world where management choices were made only at (Model Risk Alert) time zero and the consequences of those decisions would become apparent at the end of the single period.

Given these assumptions, what happens at the end of the single period at time T ? If the value of company assets $V(T)$ is greater than the amount of debt that has to be paid off, K , then the excess value $V(T) - K$ goes to equity holders in exactly the same way as if the company were liquidated in the perfectly liquid markets Merton assumes. If the value of company assets is less than K , then equity holders get nothing. More important, from a Basel II and III perspective, debt holders have a loss given default of $K - V(T)$.

Model Risk Alert

Loss given default is built into the Merton model, an insight many market participants have overlooked.

The payoffs on the common stock of the firm are identical to that of a European call option at current time t with these characteristics:

- The option has maturity $T - t$, the same maturity as the debt. The original maturity of the option was $T - 0$.
- The option has a strike price equal to the amount of debt K .
- The option's value is determined by the current value of company assets $V(t)$ and their volatility σ .

We label this call option, the value of the company's equity, $C(V(t), t, T, K)$. Then Merton invokes an argument of Nobel prize winners Modigliani and Miller (1958). Merton assumes that the value of company assets equals the value of its equity plus the value of its debt. That means we can write the value of the company's debt

$$D(V(t), t, T, K) = V(t) - C(V(t), t, T, K)$$

That, in essence, is the heart of the Merton model intuition. Alternative “versions” of the Merton model are relatively minor variations around this theme. Many “versions” in fact are just alternative methods of estimating the Merton parameters.

In the next few sections, we summarize some of the variations on the Merton model. After that, we turn to practical use of the Merton model.

THE BASIC MERTON MODEL

In order to be precise about the Merton model and related default probabilities, we need to be more precise about the assumptions used by Black, Scholes, and Merton.⁵

- r Interest rates are constant at a continuously compounded rate r .
- σ The value of company assets has a constant instantaneous variance of the return on company assets, $\sigma(V(t), t)^2 = \sigma$.
- K K is the amount of zero-coupon debt payable at maturity, time T .

Black, Scholes, and Merton assume that the value of company assets follows a lognormal distribution and that there are no dividends on these assets and no other cash flows from them during the period. Formally, we can write the evolution of the value of company assets over time as

$$\frac{dV(t)}{V(t)} = \alpha dt + \sigma dW(t)$$

$W(t)$ is normally distributed with mean zero and variance t —it is the “random number generator” that produces shocks to the value of company assets. The magnitude of these shocks is scaled by the volatility of company assets, σ . The term α is the drift in the value of company assets. We discuss its value later in this chapter, but we do not need to spend time on that now. From this formula, we know that $V(t)$, the value of company assets, is lognormally distributed. Taking advantage of these facts, the value of company assets at a future point in time can be written like this:

$$V(t) = V(0)e^{[(\alpha - \frac{1}{2}\sigma^2)t + \sigma W(t)]}$$

We can use this equation to get the value of company assets and its distribution at maturity T . $W(t)$, as mentioned above, is normally distributed with mean zero and variance t (standard deviation of $\sqrt{t}$). This allows us to get the well-known “distance to default” in the Merton model. We can do this by first rewriting the formula for the value of company assets to get the instantaneous return on company assets. The return on company assets is normally distributed because W is normally distributed:

$$y(t) = \frac{1}{t} \left[\left(\alpha - \frac{1}{2} \sigma^2 \right) t + \sigma W(t) \right]$$

The distance from default is the number of standard deviations that this return is (when $t = T$, the date of maturity) from the instantaneous return just needed to “break even” versus the value of debt that must be paid, K . We can write this breakeven return as $y(K, T)$ or

$$y(K, T) = \frac{1}{T} \ln \left[\frac{K}{V(0)} \right]$$

The standard deviation of the return on company assets $y(t)$ when $t = T$ is $\sigma/\sqrt{T}$ so we can write the distance to default $z(T)$ as

$$z(T) = \frac{y(T) - y(K, T)}{\frac{\sigma}{\sqrt{T}}}$$

Later in this chapter, we discuss how this distance to default or the Merton default probabilities themselves can be mapped to actual historical default rates. In order to use the Merton default probabilities, we need to evaluate α . We turn to that in the next section.

Making use of the derivation so far, Merton shows that the value of company equity $C(V(0), 0, T, K)$ is equal to the value of a Black-Scholes call option:

$$C[V(0), 0, T, K] = V(0)N(d_1) - KP(0, T)N(d_2)$$

$N()$ is the cumulative normal distribution function that can be found in any popular spreadsheet package. The variable $d_1 = \ln[V(0)/KP(0, T)]$, and $d_2 = d_1 - \sigma\sqrt{(T-0)}$. The expression $P(0, T)$ is the zero-coupon bond price of \$1 paid at time T , discounting at the constant risk-free interest rate r . Using the Modigliani and Miller argument that the value of company assets must equal the sum of the value of this call option, the value of the debt D gives us

$$V(0) = C[V(0), 0, T, K] + D[V(0), 0, T, K]$$

Rearranging this gives the Merton formula for risky debt. The second expression takes advantage of the symmetrical nature of the normal distribution.

$$\begin{aligned} D[V(0), 0, T, K] &= V(0) - V(0)N(d_1) + KP(0, T)N(d_2) \\ &= V(0)N(-d_1) + KP(0, T)N(d_2) \end{aligned}$$

As we have stated throughout this book, our objective is to come up with an integrated approach to market risk, interest rate risk, liquidity risk, and credit risk that provides pricing, hedging, and valuation. Ultimately, we want to be able to value the Jarrow-Merton put option on the assets of our company or business unit that

provides the best measure of risk. Therefore, it is very important to note that all of the Black-Scholes hedging concepts can be applied to the Merton model of risky debt, in theory, but these hedges have been proven to be wildly inaccurate in practice.

The delta of risky debt shows the change in the value of risky debt that results from a change in the value of company assets:

$$\text{Delta} = \frac{\partial D}{\partial V} = 1 - N(d_1) > 0$$

The value of risky debt rises when the value of company assets V rises. The gamma of risky debt is the second derivative of the Merton risky debt formula with respect to the value of company assets:

$$\text{Gamma} = \frac{\partial^2 D}{\partial V^2} = \frac{-f(d_1)}{V\sigma\sqrt{T}} < 0$$

Similarly, we can derive derivatives for changes in time, volatility, and the level of interest rates, in a fashion parallel to that of the Black-Scholes model:

$$\text{Theta} = \frac{\partial D}{\partial t} = \frac{Vf(d_1)\sigma}{2\sqrt{T}} + rKP(0, T)N(d_2) > 0$$

$$\text{Vega} = \frac{\partial D}{\partial \sigma} = -V\sqrt{T} f(d_1) < 0$$

$$\text{Rho} = \frac{\partial D}{\partial r} = -TKP(0, t)N(d_2) \leq 0$$

Note that we can hedge our position in risky debt by taking the appropriate positions in the assets of the company, just like a trader of stock options can hedge his position in options with a replicating position in the underlying common stock. Similarly, in the Merton model, both the value of the debt and the value of the common stock are functions of the value of company assets. This means that a position in the debt of the company can be hedged with the appropriate “short” position in the common stock of the company.

Model Risk Alert

Note that when the value of common stock approaches zero this hedge will break down in practice, for good sound theoretical reasons. The value of debt can keep declining as the value of company assets declines, but the stock price cannot go below zero.

Finally, note that we can do pricing, hedging, and valuation in the Merton model without any discussion whatsoever of default probabilities and loss given default. They are totally unnecessary, since the valuation formulas take into account all possible values of company assets, the probability of default at each asset level, and the loss given default at each asset level. This is obvious from a theoretical point of

view but (Model Risk Alert) it is often forgotten by practitioners who follow the steps of this process:

1. Estimate the Merton model of risky debt
2. Derive the Merton default probability
3. Apply this default probability to calculate the value of a debt instrument

If our objective is valuation and hedging, and if the Merton model is true (it is not), steps 2 and 3 are unnecessary given the sound theoretical structure that Merton has laid out. Before turning to the subject of default probabilities, we emphasize this point by discussing valuation of a bond with many payments.

VALUING MULTIPAYMENT BONDS WITH THE MERTON MODEL OF RISKY DEBT

How do we apply the single-period Merton model to value a bond with many payments, say a 10-year bond with semiannual payments? To do the valuation, we need 20 zero-coupon bond prices to apply the present value techniques of Chapter 4. Remember, if we have a smoothed yield curve using the techniques in Chapter 5, we do not need the Merton model for valuation because we can pull the 20 zero-coupon bond prices we need directly from the yield curve with no need to resort to the Merton theory. When we say “use the Merton model” for valuation, what we are really saying is this—how do we fit the parameters of the model to coupon-paying bonds with observable prices? We will return to this issue below when we discuss parameter fitting. Let’s say we have a three-year bond with an observable price, a five-year bond with an observable price, and a 10-year bond with an observable price. How can we fit the Merton model to this data? If all the bonds have the same payment dates (say June 15 and December 15) then we can guess the value of debt to be paid off K , the value of company assets $V(0)$, and the volatility of company assets σ . We are immediately confronted with a conflict with model assumptions because we will see from the smoothed risk-free yield curve that the value of the continuously compounded zero yield r will not be constant; there will be 20 different values, one for each payment date.

At this point, users of the Merton model get to the quick and dirty nature of application that is consistent with using the constant interest rate Black model to value interest rate caps! Given the 20 values of the risk-free rate, in our example we have three pricing equations for the bonds and three variables to solve for, K , $V(0)$, and σ . At the same time, theory says we should also be fitting the model to the stock price, but we do not know which of the 20 maturities we are using is most relevant to the stock price.

Let’s squint our eyes, hold our nose, and get this done. As a practical matter, the authors would fit the model to the shortest maturity observable debt issue, the longest observable debt issue, and stock price. For any pricing error that arises, we can “extend” the model just like we extended the Vasicek term structure model to get the Hull and White model. We can use the techniques of Chapter 14 to find the best-fitting Merton parameters just like we did for term structure model parameters.

When we do this, what should we extend to fit the observable term structure for say Hewlett Packard using the Merton model? Users can extend K , the value of debt to be paid off, to be different at each maturity. Another option is to extend the volatility of company assets σ . A final option is to allow for drift in the value of company assets. All of these are options. Van Deventer and Imai (2003) show using data from First Interstate Bancorp new issue spreads that the pricing error can be large when using the Merton model without some form of extension.

ESTIMATING THE PROBABILITY OF DEFAULT IN THE MERTON MODEL

As shown previously, we do not need the Merton model default probabilities for valuation, pricing, and hedging in the Merton framework. The model implicitly considers all possible default probabilities and the associated loss given default for each possible ending value of company assets $V(T)$. As an index of risk, however, it was once hoped that Merton default probabilities would be an accurate measure of default risk. As we saw in Chapter 16, this hope did not materialize.

One of the curious features of the Black-Scholes options model in the Merton context is that the drift term in the return on company assets α does not affect the valuation formulas just like the drift in the return on a common stock does not affect its option pricing. We do need to know α when it comes to the probability of default on an empirical or statistical basis (as opposed to the risk-neutral default probabilities that are implicitly embedded in the valuation formulas).

Jarrow and van Deventer (1998) derive the empirical default probabilities using Merton's intertemporal capital asset pricing model. That model recognizes that the value of company assets of all companies are not independent. They will be priced in a market that recognizes the benefits of diversification, a market that knows lending to a suntan lotion company and an umbrella maker diversifies away weather risk. These benefits of correlation in the returns on company assets affect the expected return on the assets. Jarrow and van Deventer note that, when interest rates are constant, the expected return on the assets of company i in Merton's (1973) equilibrium asset pricing model must be

$$\alpha_i = r + \frac{\sigma_i \rho_{iM}}{\sigma_M} (\alpha_M - r)$$

Where r is the continuously compounded risk-free rate, the subscript M refers to the market portfolio, or equivalently, the portfolio consisting of all assets of all companies in the economy, and ρ_{iM} denotes the correlation between the return on firm i 's asset value and the return on the market portfolio of all companies' assets.

Given this formula for α , the probability of default in the Merton model becomes

$$\text{Probability(Default)} = \text{Probability}(V_i(T) < K) = N\left(\frac{\ln(K/V_i(0)) - \mu_i T}{\sigma_i \sqrt{T}}\right)$$

The variable μ is defined as

$$\mu_i = -\frac{1}{2}\sigma_i^2 + (1 - b_i)r + b_i a_M$$

where

$$b_i = \frac{\sigma_i \rho_{iM}}{\sigma_M}$$

is the i th firm's beta, which measures the degree to which the return on the assets of the i th company are correlated with the return on the assets of all companies.

When interest rates are not constant, α takes on a more complex form given by Merton.

IMPLYING THE VALUE OF COMPANY ASSETS AND THEIR RETURN VOLATILITY σ

The market value of company assets in the Merton model is not observable, nor is its volatility. If the asset volatility is known, we can solve implicitly for the value of company assets $V(0)$ by solving the equation relating the total value of the company's equity (stock price S times number of shares outstanding M) to the Black-Scholes call option formula:

$$MS = C[V(0), 0, T, K] = V(0)N(d_1) - KP(0, T)N(d_2)$$

The volatility of the return on company assets is not observable, but we can observe the stock price history of the firm and calculate the volatility of the return on the stock σ_C in typical Black-Scholes fashion.

Model Risk Alert

For highly distressed companies, the stock price may only be observable intermittently, it can remain near zero for an extended period of time, and it can bounce randomly between bid and offered levels. All of these facts lead to inaccurate volatility assessment and inaccurate default probabilities.

Model Risk Alert

Let's assume we are true believers in the Merton model and assume all of these real-world complications do not exist. In the appendix to this chapter, we use Ito's lemma in the same way we did earlier in the book, when we introduced term structure models of interest rate movements. This allows us to link the observable volatility of equity returns σ_C to the volatility of the returns on company assets σ :

$$\frac{\sigma_C C}{C_V V} = \sigma$$

Together with the previous call option formula, we have two equations in two unknowns $V(0)$ and σ . We solve for both of them and create a time series of each. From these time series on the value of company assets for all companies, we can calculate the betas necessary to insert in the previous default probability formula. The expected return on the assets of all companies is usually estimated as a multiple of the risk-free rate in such a way that model performance (as we discussed in Chapter 16) is optimized.

MAPPING THE THEORETICAL MERTON DEFAULT PROBABILITIES TO ACTUAL DEFAULTS

Falkenstein and Boral (2000) and Lando (2004), in his excellent book, note the theoretical values of Merton default probabilities produced in this manner typically result in large bunches of companies with very low default probabilities and large bunches of companies with very high default probabilities. In both cases, the Merton default probabilities are generally much higher than the actual levels of default. The best way to resolve this problem is to use the theoretical Merton model as an explanatory variable in the reduced form framework, using logistic regression. This is one of the models (“KDP-ms”) offered by KRIS. As long as this route is taken, however, as Chapter 16 makes clear, it is better to simply abandon the Merton default probability completely because it adds little or no incremental explanatory power to a well-specified reduced form model.

As we note in Chapter 16, this mapping of Merton default probabilities does not change the ROC accuracy ratio, the standard statistical measure of model performance for 0/1 events like bankruptcy. The mapping process does improve the Falkenstein and Boral test we employed in Chapter 16.

THE MERTON MODEL WHEN INTEREST RATES ARE RANDOM

As mentioned in the introduction to this chapter, one of the earliest concerns expressed about the accuracy of the Merton model was its reliance on the assumption that interest rates are constant, exactly as in the Black-Scholes model. Shimko, Tejima, and van Deventer (1993) solved this problem by applying Merton’s option model with random interest rates to the risky debt problem. It turns out that the random interest rates version of Black-Scholes can be used with many of the single-factor, term-structure models that we reviewed in Chapter 13, which is one of the reasons for their importance in this volume.

THE MERTON MODEL WITH EARLY DEFAULT

Another variation on the Merton problem is focused on another objection to the basic structure of the Merton model: default never occurs as a surprise. It either occurs at the end of the period or it does not, but it never occurs prior to the end of the period. Analysts have adopted the down-and-out option in order to trigger

early default in the Merton model. In this form of the model, equity is a down-and-out call option on the assets of the firm. Equity becomes worthless if the value of company assets falls to a certain barrier level, often the amount of risky debt that must be repaid at the end of the period K . Adopting this formulation, the total value of company equity again equals an option on the assets of the firm, but this option is a down-and-out option. Lando (2004) discusses the variations on this approach extensively. Unfortunately, these variations have little or no ability to improve the accuracy of the default probabilities implied by the model by the standards of Chapter 16.

LOSS GIVEN DEFAULT IN THE MERTON MODEL

As mentioned earlier in this chapter, analysts are divided between those who seek to use the Merton model for valuation, pricing, and hedging, and those who use it to estimate default probabilities and then in turn try to use the default probabilities and a separate measure of loss given default for pricing, valuation, and hedging outside the Merton model structure. The former group is truer to the model, because the Merton model incorporates all potential probabilities of default and loss given default in its valuation formulas. For this reason, there is no need to specify default probabilities and loss given default for valuation in the model—they are endogenous, just like the Black-Scholes model considers all potential rates of return on common stock in the valuation of an option on the common stock.

Loss given default is the expectation (with respect to the value of company assets (V)) for the amount of loss $K - V(T)$, which will be incurred given that the event of default has occurred. The derivation of loss given default in the model is left as an exercise for the reader.

COPULAS AND CORRELATION BETWEEN THE EVENTS OF DEFAULT OF TWO COMPANIES

The Merton model is normally fitted to market data in a univariate manner, one company at a time. As a result, there is no structure in the risky debt model itself, which describes how the events of default of two companies would be correlated. Because this correlation is not derived from the model's structure, many analysts *assume* a correlation structure.⁶ This contrasts sharply with the implementation of the reduced form models in Chapter 16, where the correlation between the events of default of all companies is jointly derived when the model is fitted.

Back to the Merton Case

Let's assume that we need to estimate the correlation that causes two companies to have defaults and default probabilities that are not independent. If we are true to the model, we could make the argument that the returns on the value of company assets for both companies has correlation σ :

$$\begin{aligned}\frac{dV_a}{V_a} &= \mu_a dt + \sigma_a dZ_a \\ \frac{dV_b}{V_b} &= \mu_b dt + \sigma_b dZ_b \\ (dZ_a dZ_b) &= \rho dt\end{aligned}$$

The correlation of the returns on the assets of both companies can be estimated using the formulas above from the historical implied series on the value of company assets. This is the proper specification for correlation in the copula method of triggering correlated events of default. (We return to this issue in Chapter 20, and the results of our investigation in that chapter are not for the squeamish.)

There are a number of other ways of doing the same thing. If the analyst has an observable probability that both companies will default between time t and T , then the value of the correlation coefficient ρ can be implied from the Merton model default probability formulas and the joint multivariate normal distribution of the returns on company assets for the two firms. There is no reason to assume that the pairwise correlation between all pairs of companies is the same—in fact, this is a serious error in analyzing collateralized debt obligation values (as we discuss in Chapter 20).

There are two more elegant alternatives that the authors prefer. First choice is the reduced form models, which we examine in much greater detail in the following chapters, based on the introduction in Chapter 16. A second choice, which stays true to the Merton theory, is to specify the macroeconomic factors driving the returns on company assets and to allow them to drive the correlations in the events of default. This approach is much less accurate.

PROBLEMS WITH THE MERTON MODEL: SUMMING UP

Van Deventer and Imai (2003) extensively summarize the strengths and concerns with each of the models they review, including the Merton and Shimko, Tejima, van Deventer random interest rates version of the Merton model. In their view, the key factors that contribute to the model performance issues in the Merton model can be summarized as follows:

- The model implicitly assumes a highly volatile capital structure by positing that the dollar amount of debt stays constant while the value of company assets varies randomly.
- The model (in traditional form) assumes interest rates are constant, but they are not.
- The model assumes that company assets are traded in a perfectly liquid market so the Black-Scholes replication arguments can apply. In reality, most corporate assets are very illiquid, which makes the application of the Black-Scholes model inappropriate in both theory and practice.

- The model does not allow for early default—default only occurs at the end of the period.
- The model’s functional form is very rigid and does not allow the flexibility to incorporate other explanatory variables in default probability prediction.

Finally, from a portfolio management point of view, one prefers a model where the correlation in defaults can be derived, not assumed. As we have emphasized throughout this volume, we are seeking pricing, valuation, and hedging at both the transaction level and the portfolio level. The Merton model leaves one wondering “how do I hedge?” Even more serious, we showed in Chapter 16 that many naïve models of default have a greater accuracy in the ordinal ranking of companies by riskiness than the Merton model. This result alone has led a very large number of financial institutions and regulators to abandon their experimentation with the Merton idea. It is simply too much to hope that the value of debt and equity of a firm are driven by only one random factor, the value of company assets. We proved in Chapter 16 that the model simply does not work, consistent with the results of the many authors listed previously.

APPENDIX

In the Merton model of risky debt, the volatility of returns on company assets (not the volatility of the value of assets) σ is a critical parameter, but it is not observable.

We can observe the historical volatility of returns on the company’s common stock. (Note, this is *not* the same as the volatility of the stock price or standard deviation of the stock price.) By linking this observable historical volatility of common stock returns by formula to the historical volatility of company asset returns, we can derive the value of the volatility of company asset returns σ .

Assumptions

The value of company assets V follows the same stochastic process as specified in Black-Scholes:

$$\frac{dV}{V} = \alpha dt + \sigma dZ$$

We will use Ito’s lemma to expand the formula for the value of equity, the Black-Scholes call option $C[V(0), 0, T, K]$, and then link its volatility to the volatility of company assets σ .

Using Ito’s Lemma to Expand Changes in the Value of Company Equity

Merton’s assumptions lead to the conclusion that the value of company assets is a call option on the assets of the firm $C[V(0), 0, T, K]$. Using Ito’s lemma as we did in the

chapter on term structure models, we can write the stochastic formula for changes in stock price as a function of V :

$$dC = C_V dV + \frac{1}{2} C_{VV} (dV)^2 dt + C_t dt$$

Substituting for dV gives

$$dC = C_V (V\alpha dt + V\sigma dZ) + \frac{1}{2} C_{VV} (dV)^2 dt + C_t dt$$

We can then rearrange, separating the drift from the stochastic terms:

$$dC = [C_V V\alpha dt + \frac{1}{2} C_{VV} (dV)^2 dt + C_t dt] + C_V V\sigma dZ$$

The equation above is for the change in stock price C . The volatility of stock price (not return on the stock price) is $C_V V\sigma$. We can get the expression for the volatility of the return on the stock price by dividing both sides by stock price C :

$$\frac{dC}{C} = \frac{[C_V V\alpha dt + \frac{1}{2} C_{VV} (dV)^2 dt + C_t dt]}{C} + \frac{C_V V\sigma}{C} dZ$$

Therefore, the market convention calculation for the volatility of the return on common stock σ_C (which is observable given the observable stock price history) is

$$\sigma_C = \frac{C_V V\sigma}{C}$$

We solve this equation for σ to express the volatility of the return on the assets of the firm as a function of the volatility of the return of the firm's common stock:

$$\frac{\sigma_C C}{C_V V} = \sigma$$

This formula can be used for σ in the Merton model of risky debt.

Model Risk Alert

While this approach is fine in theory, it does not work well in practice. Volatilities estimated in this way are distorted by a number of common phenomena: stock prices near constant as they near zero, intermittent trading, trading at bid and offered prices, and so on.

NOTES

1. A copy of “The Ratings Chernobyl” can be found at www.kamakuraco.com/Blog/tabid/231/EntryId/11/The-Ratings-Chernobyl.aspx. The blog was originally redistributed by *RiskCenter.com* and by the Global Association of Risk Professionals.
2. For an example of such a recent self-assessment, please see this example from Standard & Poor’s released in 2011 at www.standardandpoors.com/ratings/articles/en/us/?assetID=1245302234237. The other two major rating agencies release similar self-assessments. In this section, in the interests of brevity, we use the S&P self-assessment. An analysis of the other two rating agencies would yield similar results.
3. This event is documented in Donald R. van Deventer, “Case Studies in Liquidity Risk: JPMorgan Chase, Bear Stearns and Washington Mutual,” Kamakura blog, www.kamakuraco.com, July 8, 2011.
4. For the the credit default swap auction protocols triggered by Fannie Mae and Freddie Mac/s “credit event” and announced by ISDA, see www.dechert.com/library/FS_FRE_10-08_ISDA_Credit_Default_Swap_Auction_Protocols.pdf.
5. For readers who would like a more detailed explanation of options concepts, the authors recommend Ingersoll (1987), Hull (1993), and Jarrow and Turnbull (1996).
6. Private conversation with Robert Jarrow, February 4, 2004.

Valuing Credit Risky Bonds

In this part of the book, we turn to a very concrete set of tasks. Our ultimate objective is to calculate the value of the Jarrow-Merton put option on the net assets (value of assets less value of liabilities) of the firm, which Jarrow and Merton propose as the best measure of total risk. *Total risk* combines interest rate risk, market risk, liquidity risk, credit risk, foreign exchange risk, and so on. In order to value this option, we need to be able to value all of the instruments on the financial institution's balance sheet. We derive the value of the option on all assets less all liabilities by performing this analysis on each underlying instrument, one by one. There is no reason to summarize or stratify the transactions in order to do this calculation. Unless the instruments are literally identical (such as a demand deposit or time deposit with identical coupon and maturity date), any summarization introduces an error whose magnitude is unknown without a comparable transaction level calculation. The cost of the hardware to achieve a target calculation time is a lower cost than the cost of human labor to oversee the data summarization and quality control on that process for almost every imaginable risk calculation. If you are hearing otherwise, someone is misinformed.

Now, in addition to valuation, we need to know how to simulate the value that each financial instrument may take on at multiple points in the future and in multiple scenarios. These scenarios may be deterministic (such as the regulatory stress tests widely imposed in 2012) or stochastic, produced by a Monte Carlo simulation.

We go through this process in a systematic pattern. First, we start with straight bonds, bonds with credit risk, and bonds with no credit risk. In the next chapter, we extend this analysis to credit derivatives and the formerly trendy collateralized debt instruments. In Chapter 21, we begin to get to the “option” portion of this exercise by studying options on bonds. We follow with valuation techniques for a string of related financial instruments thereafter, bringing it all together in Chapters 36 to 41.

We turn now to the valuation of bonds with and without credit risk. We then turn to simulation of future values.

THE PRESENT VALUE FORMULA

In Chapter 4, we reviewed the basic case where the value of a financial instrument was the simple sum of the present value of each cash flow. The expression $P(t)$ is the present value of a dollar received at time t from the perspective of time zero. The expression $C(t)$ is the cash flow to be paid at time t .

$$\text{Present value} = \sum_{i=1}^n P(t_i)C(t_i)$$

Do we need to say any more than this about bond valuation? Surprisingly, the answer is yes.

VALUING BONDS WITH NO CREDIT RISK

Fortunately, in the case of bonds with a default probability of zero, valuation is as simple as the formula above. In a no-arbitrage economy with frictionless trading, the present value formula above applies. As we noted in Chapter 4, “present value” does not mean price. Rather, it means price plus accrued interest. Where do we get the values of the discount factors $P(t)$? We get them from the yield curve-smoothing techniques discussed in Chapter 5.

Given these P values and the terms of the bond, its valuation is very simple. What about simulating future values? Chapter 6 through 9 on the Heath, Jarrow, and Morton (HJM) approach showed us how this should be done.

SIMULATING THE FUTURE VALUES OF BONDS WITH NO CREDIT RISK

The Jarrow-Merton put option that is our ultimate valuation objective can be thought of as a European option, valued on only one date, or it can be thought of as an American option with many possible exercise dates. The European option analogy is a way of thinking about risk that is consistent with traditional value at risk concepts—risk matters on only one date, we can ignore intervening cash flows, and analyze the dickens out of the single date we choose as important. Common choices are 10-day horizons, one-month horizons, and one-year horizons. The longer the horizon, the more tenuous this way of thinking becomes. Most analysts, including the authors, believe that a multiperiod approach to the Jarrow-Merton put option and risk management in general is essential. How do we calculate the variation in the value of a bond with no credit risk over time?

We use the HJM approaches of Chapters 6 through 9 and the parameter-fitting techniques of Chapter 14. In Chapter 3, we presented the statistical evidence that shows why it is necessary to simulate future yield curve movements using multiple risk factors, not a single risk factor like a parallel shift or a change in the short rate that drives the rest of the “single factor” risk-free yield curve. The steps in simulating the risk-free term structure can be summarized as follows, for N scenarios covering M time steps and K risk factors:

1. Capture today’s continuous zero-coupon yield curve using a sophisticated yield curve smoothing technique like the maximum smoothness forward rate technique. We use the yield curve inputs from Chapter 9 from now on.
2. Choose a term structure model and calculate its parameters. For the rest of this book, we rely primarily on the three-factor HJM risk-free yield curve of Chapter 9.

3. Choose a time step to step forward in time, delta t , to arrive at time step 1. For expository purposes, we use the same one-year time step that we employed in Chapter 9.
4. Calculate the random value for scenario 1 for the K risk factors at time step 1. In Chapter 9, the random factors were the change in the year-one spot rate, the change in the one-year forward rate maturing at time $T = 10$, and a third factor that is “everything else.” The volatilities of these factors are given in the volatility tables from Chapter 9, and their mean values differ at each forward point in time, conditional on the state or level of interest rates in that scenario. As in Chapter 9, we do a table lookup to get the right values. Alternatively, we could have specified a mathematical function of the level of rates.
5. Re-smooth the yield curve at time step 1. We apply the yield curve-smoothing techniques of Chapter 5 to smooth the yield curve between the values of the yields whose simulated values we know (the HJM values at each annual time step) so that we can value cash flows that do not fall exactly on a time step.
6. Move to the next time step and repeat steps 4 and 5 until you have completed the simulation to M time steps, completing scenario j (and of course for the first scenario $j = 1$).
7. Repeat steps 4 to 6 for a total of N scenarios.

At each point in time, in each scenario, we can calculate the value of the risk-free bond using the formula above. The number of cash flows and the time before they arrive will of course change as time passes.

We now show how to value a bullet bond, an amortizing loan, and a floating-rate loan using the HJM three-factor interest rate model from Chapter 9.

CURRENT AND FUTURE VALUES OF FIXED INCOME INSTRUMENTS: HJM BACKGROUND AND A STRAIGHT BOND EXAMPLE

We use the risk-free yield curve and the bushy tree from the three-factor HJM model in Chapter 9 for the examples in this chapter. We create a bushy tree for four periods and three risk factors, so there is the same number of nodes in the tree as in Chapter 9, that is, 4 in period 1; 16 in period 2; and 64 in period 3, which we can use to evaluate cash flows in period 4. The zero-coupon bond prices prevailing at time zero are given in Exhibit 19.1, taken directly from Chapter 9.

The probability of each node in the tree is given in Exhibit 19.2.

EXHIBIT 19.1 Key Assumptions

	Period of Maturity			
	1	2	3	4
Input Zero-Coupon Bond Prices	0.9970057865	0.9841101497	0.9618922376	0.9308550992

The discount factor for \$1 received at each state is given in Exhibit 19.3, as calculated in Chapter 9.

These discount factors are consistent with the term structure of zero-coupon bond prices derived in Chapter 9. The probability-weighted discount factors for each state are given in Exhibit 19.4. For each cash flow pattern we want to value, we simply multiply the cash flow at each period in each state times these probability-weighted discount factors to get the valuation at time zero.

EXHIBIT 19.2 Probability of Each State

Row Number	Current Time				
	0	1	2	3	4
1	100.0000%	12.5000%	1.5625%	0.1953%	0.1953%
2		12.5000%	1.5625%	0.1953%	0.1953%
3		25.0000%	3.1250%	0.3906%	0.3906%
4		50.0000%	6.2500%	0.7813%	0.7813%
5			1.5625%	0.1953%	0.1953%
6			1.5625%	0.1953%	0.1953%
7			3.1250%	0.3906%	0.3906%
8			6.2500%	0.7813%	0.7813%
9			3.1250%	0.3906%	0.3906%
10			3.1250%	0.3906%	0.3906%
11			6.2500%	0.7813%	0.7813%
12			12.5000%	1.5625%	1.5625%
13			6.2500%	0.7813%	0.7813%
14			6.2500%	0.7813%	0.7813%
15			12.5000%	1.5625%	1.5625%
16			25.0000%	3.1250%	3.1250%
17				0.1953%	0.1953%
18				0.1953%	0.1953%
19				0.3906%	0.3906%
20				0.7813%	0.7813%
21				0.1953%	0.1953%
22				0.1953%	0.1953%
23				0.3906%	0.3906%
24				0.7813%	0.7813%
25				0.3906%	0.3906%
26				0.3906%	0.3906%
27				0.7813%	0.7813%
28				1.5625%	1.5625%
29				0.7813%	0.7813%
30				0.7813%	0.7813%
31				1.5625%	1.5625%
32				3.1250%	3.1250%
33				0.3906%	0.3906%
34				0.3906%	0.3906%

EXHIBIT 19.2 (Continued)

Row Number	Current Time				
	0	1	2	3	4
35				0.7813%	0.7813%
36				1.5625%	1.5625%
37				0.3906%	0.3906%
38				0.3906%	0.3906%
39				0.7813%	0.7813%
40				1.5625%	1.5625%
41				0.7813%	0.7813%
42				0.7813%	0.7813%
43				1.5625%	1.5625%
44				3.1250%	3.1250%
45				1.5625%	1.5625%
46				1.5625%	1.5625%
47				3.1250%	3.1250%
48				6.2500%	6.2500%
49				0.7813%	0.7813%
50				0.7813%	0.7813%
51				1.5625%	1.5625%
52				3.1250%	3.1250%
53				0.7813%	0.7813%
54				0.7813%	0.7813%
55				1.5625%	1.5625%
56				3.1250%	3.1250%
57				1.5625%	1.5625%
58				1.5625%	1.5625%
59				3.1250%	3.1250%
60				6.2500%	6.2500%
61				3.1250%	3.1250%
62				3.1250%	3.1250%
63				6.2500%	6.2500%
64				12.5000%	12.5000%
Total	100.0000%	100.0000%	100.0000%	100.0000%	100.0000%

EXHIBIT 19.3 Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1.0000	0.9970	0.9759	0.9326	0.8789
2		0.9970	0.9759	0.9326	0.8834

(Continued)

EXHIBIT 19.3 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
3		0.9970	0.9759	0.9326	0.8868
4		0.9970	0.9759	0.9326	0.8632
5			0.9902	0.9623	0.9333
6			0.9902	0.9623	0.9380
7			0.9902	0.9623	0.9431
8			0.9902	0.9623	0.9247
9			0.9862	0.9511	0.9039
10			0.9862	0.9511	0.9326
11			0.9862	0.9511	0.9218
12			0.9862	0.9511	0.9015
13			0.9836	0.9301	0.8823
14			0.9836	0.9301	0.8868
15			0.9836	0.9301	0.8901
16			0.9836	0.9301	0.8665
17				0.9684	0.9298
18				0.9684	0.9593
19				0.9684	0.9482
20				0.9684	0.9273
21				0.9843	0.9631
22				0.9843	0.9789
23				0.9843	0.9768
24				0.9843	0.9726
25				0.9822	0.9571
26				0.9822	0.9859
27				0.9822	0.9796
28				0.9822	0.9613
29				0.9780	0.9609
30				0.9780	0.9657
31				0.9780	0.9710
32				0.9780	0.9521
33				0.9676	0.9460
34				0.9676	0.9507
35				0.9676	0.9560
36				0.9676	0.9373
37				0.9724	0.9507
38				0.9724	0.9555
39				0.9724	0.9608
40				0.9724	0.9420
41				0.9778	0.9435
42				0.9778	0.9719
43				0.9778	0.9657
44				0.9778	0.9477
45				0.9587	0.9173

EXHIBIT 19.3 (Continued)

Row Number	Current Time				
	0	1	2	3	4
46				0.9587	0.9465
47				0.9587	0.9355
48				0.9587	0.9149
49				0.9624	0.9298
50				0.9624	0.9594
51				0.9624	0.9482
52				0.9624	0.9273
53				0.9672	0.9428
54				0.9672	0.9475
55				0.9672	0.9528
56				0.9672	0.9341
57				0.9725	0.9480
58				0.9725	0.9528
59				0.9725	0.9580
60				0.9725	0.9393
61				0.9535	0.9124
62				0.9535	0.9323
63				0.9535	0.9282
64				0.9535	0.9079

EXHIBIT 19.4 Probability of Weighted Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1.0000	0.1246	0.0152	0.0018	0.0017
2		0.1246	0.0152	0.0018	0.0017
3		0.2493	0.0305	0.0036	0.0035
4		0.4985	0.0610	0.0073	0.0067
5			0.0155	0.0019	0.0018
6			0.0155	0.0019	0.0018
7			0.0309	0.0038	0.0037
8			0.0619	0.0075	0.0072
9			0.0308	0.0037	0.0035
10			0.0308	0.0037	0.0036
11			0.0616	0.0074	0.0072
12			0.1233	0.0149	0.0141
13			0.0615	0.0073	0.0069
14			0.0615	0.0073	0.0069
15			0.1230	0.0145	0.0139

(Continued)

EXHIBIT 19.4 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
16			0.2459	0.0291	0.0271
17				0.0019	0.0018
18				0.0019	0.0019
19				0.0038	0.0037
20				0.0076	0.0072
21				0.0019	0.0019
22				0.0019	0.0019
23				0.0038	0.0038
24				0.0077	0.0076
25				0.0038	0.0037
26				0.0038	0.0039
27				0.0077	0.0077
28				0.0153	0.0150
29				0.0076	0.0075
30				0.0076	0.0075
31				0.0153	0.0152
32				0.0306	0.0298
33				0.0038	0.0037
34				0.0038	0.0037
35				0.0076	0.0075
36				0.0151	0.0146
37				0.0038	0.0037
38				0.0038	0.0037
39				0.0076	0.0075
40				0.0152	0.0147
41				0.0076	0.0074
42				0.0076	0.0076
43				0.0153	0.0151
44				0.0306	0.0296
45				0.0150	0.0143
46				0.0150	0.0148
47				0.0300	0.0292
48				0.0599	0.0572
49				0.0075	0.0073
50				0.0075	0.0075
51				0.0150	0.0148
52				0.0301	0.0290
53				0.0076	0.0074
54				0.0076	0.0074
55				0.0151	0.0149
56				0.0302	0.0292
57				0.0152	0.0148
58				0.0152	0.0149

EXHIBIT 19.4 (Continued)

Row Number	Current Time				
	0	1	2	3	4
59				0.0304	0.0299
60				0.0608	0.0587
61				0.0298	0.0285
62				0.0298	0.0291
63				0.0596	0.0580
64				0.1192	0.1135

VALUATION OF A STRAIGHT BOND WITH A BULLET PRINCIPAL PAYMENT AT MATURITY

We saw in Chapter 9 that a bond paying an annual coupon of \$3 and principal of \$100 at time $T = 4$ could be valued directly using the present value formula above. Alternatively, we can do a Monte Carlo simulation or bushy tree using the HJM no-arbitrage framework and get the identical time 0 value. Exhibit 19.5 shows that the valuation can also be calculated using the bushy tree by inserting the correct cash flow in the cash flow table and multiplying each cash flow by the probability-weighted discount factors in Exhibit 19.4. When we do this, we get the same value as in Chapter 9: 104.7071.

We can calculate the future values of this bond, at each node of the bushy tree in each state, by using the term structure of zero-coupon bond prices that prevails on the bushy tree, shown in Chapter 9. We now turn to an amortizing structure.

VALUING AN AMORTIZING LOAN

An amortizing bond or loan pays the same dollar amount in each period, say A . If there is no default risk or prepayment risk, a loan paying this constant amount A per period for n periods can be valued directly using a special case of the present value formula above:

$$\text{Present value} = A \sum_{i=1}^n P(t_i)$$

Keeping with our example from Chapter 9, we can use the zero-coupon bond prices to value an amortizing loan that pays \$50 at times 1, 2, 3, and 4. Alternatively, we can get exactly the same valuation by inserting \$50 for every state at all four times 1, 2, 3, and 4, producing the valuation of \$193.6932 as in Exhibit 19.6.

We now turn to a floating-rate structure.

EXHIBIT 19.5 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.00	3.00	3.00	3.00	103.00
2		3.00	3.00	3.00	103.00
3		3.00	3.00	3.00	103.00
4		3.00	3.00	3.00	103.00
5			3.00	3.00	103.00
6			3.00	3.00	103.00
7			3.00	3.00	103.00
8			3.00	3.00	103.00
9			3.00	3.00	103.00
10			3.00	3.00	103.00
11			3.00	3.00	103.00
12			3.00	3.00	103.00
13			3.00	3.00	103.00
14			3.00	3.00	103.00
15			3.00	3.00	103.00
16			3.00	3.00	103.00
17				3.00	103.00
18				3.00	103.00
19				3.00	103.00
20				3.00	103.00
21				3.00	103.00
22				3.00	103.00
23				3.00	103.00
24				3.00	103.00
25				3.00	103.00
26				3.00	103.00
27				3.00	103.00
28				3.00	103.00
29				3.00	103.00
30				3.00	103.00
31				3.00	103.00
32				3.00	103.00
33				3.00	103.00
34				3.00	103.00
35				3.00	103.00
36				3.00	103.00
37				3.00	103.00
38				3.00	103.00
39				3.00	103.00
40				3.00	103.00
41				3.00	103.00
42				3.00	103.00
43				3.00	103.00
44				3.00	103.00

EXHIBIT 19.5 (Continued)

Row Number	Current Time				
	0	1	2	3	4
45				3.00	103.00
46				3.00	103.00
47				3.00	103.00
48				3.00	103.00
49				3.00	103.00
50				3.00	103.00
51				3.00	103.00
52				3.00	103.00
53				3.00	103.00
54				3.00	103.00
55				3.00	103.00
56				3.00	103.00
57				3.00	103.00
58				3.00	103.00
59				3.00	103.00
60				3.00	103.00
61				3.00	103.00
62				3.00	103.00
63				3.00	103.00
64				3.00	103.00
Risk-Neutral Value =					104.7070997370

EXHIBIT 19.6 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.00	50.00	50.00	50.00	50.00
2		50.00	50.00	50.00	50.00
3		50.00	50.00	50.00	50.00
4		50.00	50.00	50.00	50.00
5			50.00	50.00	50.00
6			50.00	50.00	50.00
7			50.00	50.00	50.00
8			50.00	50.00	50.00
9			50.00	50.00	50.00
10			50.00	50.00	50.00
11			50.00	50.00	50.00
12			50.00	50.00	50.00
13			50.00	50.00	50.00
14			50.00	50.00	50.00

(Continued)

EXHIBIT 19.6 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
15			50.00	50.00	50.00
16			50.00	50.00	50.00
17				50.00	50.00
18				50.00	50.00
19				50.00	50.00
20				50.00	50.00
21				50.00	50.00
22				50.00	50.00
23				50.00	50.00
24				50.00	50.00
25				50.00	50.00
26				50.00	50.00
27				50.00	50.00
28				50.00	50.00
29				50.00	50.00
30				50.00	50.00
31				50.00	50.00
32				50.00	50.00
33				50.00	50.00
34				50.00	50.00
35				50.00	50.00
36				50.00	50.00
37				50.00	50.00
38				50.00	50.00
39				50.00	50.00
40				50.00	50.00
41				50.00	50.00
42				50.00	50.00
43				50.00	50.00
44				50.00	50.00
45				50.00	50.00
46				50.00	50.00
47				50.00	50.00
48				50.00	50.00
49				50.00	50.00
50				50.00	50.00
51				50.00	50.00
52				50.00	50.00
53				50.00	50.00
54				50.00	50.00
55				50.00	50.00
56				50.00	50.00
57				50.00	50.00
58				50.00	50.00

EXHIBIT 19.6 (Continued)

Row Number	Current Time				
	0	1	2	3	4
59				50.00	50.00
60				50.00	50.00
61				50.00	50.00
62				50.00	50.00
63				50.00	50.00
64				50.00	50.00
Risk-Neutral Value =					193.6931636452

VALUING RISK-FREE, FLOATING-RATE LOANS

In this section, we show how to value a floating-rate loan whose payments are tied to the one-period spot rate of interest that is consistent with the bushy tree from Chapter 9. The spot rate of interest at each point in time from time 0 to time 4 is given in Exhibit 19.7.

Let's assume we want to value a loan with \$1,000 in principal that matures at time $T = 4$. Let us also assume that the interest rate on the loan equals the spot rate of interest times principal plus \$50, paid in arrears. The reader can confirm that the cash flow table is consistent with Exhibit 19.8 and that the present value of the loan is 1,193.6932.

Note that the value of \$50 received at times 1, 2, 3, and 4 was valued at \$193.6932 in Exhibit 19.6, so we expect that the present value of a risk-free loan that pays the one-period spot rate plus zero premium is exactly \$1000.0000. We insert those cash flows into the cash flow table, multiply by the probability-weighted discount factors, and indeed get exactly \$1,000. This should come as no surprise, because on each branch of the bushy tree, we are discounting \$1,000 $(1 + \text{spot rate})$ by $(1 + \text{spot rate})$, consistent with no-arbitrage valuation in Chapter 9.

This result is a practical application of this formula from Chapter 4 in the context of a three-factor HJM bushy tree:

$$\begin{aligned} & \text{Value of floating-rate bond} \\ &= P(t_1) \left[\left(1 + \frac{\text{Index}}{m} \right) \text{Principal} \right] + \text{Spread} \left[\sum_{i=1}^n P(t_i) \right] \end{aligned}$$

This seems simple enough—what about the credit risky case?

VALUING BONDS WITH CREDIT RISK

Is the process of valuing bonds with credit risk any different from valuing bonds without credit risk? As we saw in Chapter 17, we need to do more work to fit credit

EXHIBIT 19.7 Spot Rate Process

State	Row Number	0	1	2	3
S-1, S-1, S-1	1	0.3003%	2.1653%	4.6388%	6.1106%
S-1, S-1, S-2	2		0.6911%	1.4142%	5.5667%
S-1, S-1, S-3	3		1.0972%	2.6089%	5.1710%
S-1, S-1, S-4	4		1.3611%	4.9179%	8.0375%
S-1, S-2, S-1	5			2.2496%	3.1041%
S-1, S-2, S-2	6			0.5984%	2.5913%
S-1, S-2, S-3	7			0.8121%	2.0291%
S-1, S-2, S-4	8			1.2459%	4.0628%
S-1, S-3, S-1	9			1.9217%	5.2210%
S-1, S-3, S-2	10			1.4149%	1.9784%
S-1, S-3, S-3	11			0.8591%	3.1798%
S-1, S-3, S-4	12			2.8695%	5.5017%
S-1, S-4, S-1	13			2.2088%	5.4252%
S-1, S-4, S-2	14			1.7005%	4.8848%
S-1, S-4, S-3	15			1.1432%	4.4917%
S-1, S-4, S-4	16			3.1593%	7.3396%
S-2, S-1, S-1	17				4.1517%
S-2, S-1, S-2	18				0.9421%
S-2, S-1, S-3	19				2.1313%
S-2, S-1, S-4	20				4.4295%
S-2, S-2, S-1	21				2.2036%
S-2, S-2, S-2	22				0.5531%
S-2, S-2, S-3	23				0.7667%
S-2, S-2, S-4	24				1.2004%
S-2, S-3, S-1	25				2.6231%
S-2, S-3, S-2	26				−0.3734%
S-2, S-3, S-3	27				0.2593%
S-2, S-3, S-4	28				2.1679%
S-2, S-4, S-1	29				1.7764%
S-2, S-4, S-2	30				1.2703%
S-2, S-4, S-3	31				0.7153%
S-2, S-4, S-4	32				2.7228%
S-3, S-1, S-1	33				2.2804%
S-3, S-1, S-2	34				1.7718%
S-3, S-1, S-3	35				1.2141%
S-3, S-1, S-4	36				3.2315%
S-3, S-2, S-1	37				2.2804%
S-3, S-2, S-2	38				1.7718%
S-3, S-2, S-3	39				1.2141%
S-3, S-2, S-4	40				3.2315%
S-3, S-3, S-1	41				3.6354%
S-3, S-3, S-2	42				0.6093%
S-3, S-3, S-3	43				1.2483%
S-3, S-3, S-4	44				3.1757%

EXHIBIT 19.7 (Continued)

State	Row Number	0	1	2	3
S-3, S-4, S-1	45				4.5070%
S-3, S-4, S-2	46				1.2864%
S-3, S-4, S-3	47				2.4797%
S-3, S-4, S-4	48				4.7858%
S-4, S-1, S-1	49				3.5013%
S-4, S-1, S-2	50				0.3117%
S-4, S-1, S-3	51				1.4934%
S-4, S-1, S-4	52				3.7773%
S-4, S-2, S-1	53				2.5831%
S-4, S-2, S-2	54				2.0730%
S-4, S-2, S-3	55				1.5136%
S-4, S-2, S-4	56				3.5370%
S-4, S-3, S-1	57				2.5831%
S-4, S-3, S-2	58				2.0730%
S-4, S-3, S-3	59				1.5136%
S-4, S-3, S-4	60				3.5370%
S-4, S-4, S-1	61				4.5061%
S-4, S-4, S-2	62				2.2686%
S-4, S-4, S-3	63				2.7252%
S-4, S-4, S-4	64				5.0230%

EXHIBIT 19.8 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	53.0032	71.6527	96.3884	1,111.1062
2		53.0032	71.6527	96.3884	1,105.6672
3		53.0032	71.6527	96.3884	1,101.7103
4		53.0032	71.6527	96.3884	1,130.3747
5			56.9105	64.1419	1,081.0407
6			56.9105	64.1419	1,075.9132
7			56.9105	64.1419	1,070.2913
8			56.9105	64.1419	1,090.6283
9			60.9723	76.0890	1,102.2100
10			60.9723	76.0890	1,069.7841
11			60.9723	76.0890	1,081.7977
12			60.9723	76.0890	1,105.0166
13			63.6107	99.1794	1,104.2520
14			63.6107	99.1794	1,098.8481

(Continued)

EXHIBIT 19.8 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
15			63.6107	99.1794	1,094.9168
16			63.6107	99.1794	1,123.3960
17				72.4959	1,091.5175
18				72.4959	1,059.4211
19				72.4959	1,071.3126
20				72.4959	1,094.2955
21				55.9835	1,072.0359
22				55.9835	1,055.5310
23				55.9835	1,057.6673
24				55.9835	1,062.0037
25				58.1208	1,076.2311
26				58.1208	1,046.2660
27				58.1208	1,052.5932
28				58.1208	1,071.6787
29				62.4592	1,067.7642
30				62.4592	1,062.7027
31				62.4592	1,057.1531
32				62.4592	1,077.2283
33				69.2173	1,072.8044
34				69.2173	1,067.7178
35				69.2173	1,062.1408
36				69.2173	1,082.3153
37				64.1486	1,072.8044
38				64.1486	1,067.7178
39				64.1486	1,062.1408
40				64.1486	1,082.3153
41				58.5911	1,086.3541
42				58.5911	1,056.0934
43				58.5911	1,062.4830
44				58.5911	1,081.7567
45				78.6949	1,095.0704
46				78.6949	1,062.8645
47				78.6949	1,074.7966
48				78.6949	1,097.8579
49				72.0883	1,085.0128
50				72.0883	1,053.1169
51				72.0883	1,064.9341
52				72.0883	1,087.7735
53				67.0053	1,075.8313
54				67.0053	1,070.7297
55				67.0053	1,065.1361
56				67.0053	1,085.3704
57				61.4322	1,075.8313

EXHIBIT 19.8 (Continued)

Row Number	Current Time				
	0	1	2	3	4
58				61.4322	1,070.7297
59				61.4322	1,065.1361
60				61.4322	1,085.3704
61				81.5926	1,095.0608
62				81.5926	1,072.6860
63				81.5926	1,077.2519
64				81.5926	1,100.2298
Risk-Neutral Value =					1,193.6932

EXHIBIT 19.9 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0032	21.6527	46.3884	1,061.1062
2		3.0032	21.6527	46.3884	1,055.6672
3		3.0032	21.6527	46.3884	1,051.7103
4		3.0032	21.6527	46.3884	1,080.3747
5			6.9105	14.1419	1,031.0407
6			6.9105	14.1419	1,025.9132
7			6.9105	14.1419	1,020.2913
8			6.9105	14.1419	1,040.6283
9			10.9723	26.0890	1,052.2100
10			10.9723	26.0890	1,019.7841
11			10.9723	26.0890	1,031.7977
12			10.9723	26.0890	1,055.0166
13			13.6107	49.1794	1,054.2520
14			13.6107	49.1794	1,048.8481
15			13.6107	49.1794	1,044.9168
16			13.6107	49.1794	1,073.3960
17				22.4959	1,041.5175
18				22.4959	1,009.4211
19				22.4959	1,021.3126
20				22.4959	1,044.2955
21				5.9835	1,022.0359
22				5.9835	1,005.5310
23				5.9835	1,007.6673
24				5.9835	1,012.0037

(Continued)

EXHIBIT 19.9 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
25				8.1208	1,026.2311
26				8.1208	996.2660
27				8.1208	1,002.5932
28				8.1208	1,021.6787
29				12.4592	1,017.7642
30				12.4592	1,012.7027
31				12.4592	1,007.1531
32				12.4592	1,027.2283
33				19.2173	1,022.8044
34				19.2173	1,017.7178
35				19.2173	1,012.1408
36				19.2173	1,032.3153
37				14.1486	1,022.8044
38				14.1486	1,017.7178
39				14.1486	1,012.1408
40				14.1486	1,032.3153
41				8.5911	1,036.3541
42				8.5911	1,006.0934
43				8.5911	1,012.4830
44				8.5911	1,031.7567
45				28.6949	1,045.0704
46				28.6949	1,012.8645
47				28.6949	1,024.7966
48				28.6949	1,047.8579
49				22.0883	1,035.0128
50				22.0883	1,003.1169
51				22.0883	1,014.9341
52				22.0883	1,037.7735
53				17.0053	1,025.8313
54				17.0053	1,020.7297
55				17.0053	1,015.1361
56				17.0053	1,035.3704
57				11.4322	1,025.8313
58				11.4322	1,020.7297
59				11.4322	1,015.1361
60				11.4322	1,035.3704
61				31.5926	1,045.0608
62				31.5926	1,022.6860
63				31.5926	1,027.2519
64				31.5926	1,050.2298
Risk-Neutral Value =					1,000.0000

risky bonds to observable market data and we need to be able to separate out the credit loss–related portion of the credit spread from the liquidity portion as we discussed then. What if we do that? Having done that and having calculated the value of a risky zero-coupon bond with maturity t , $P(t)$, can we use the formulas for value previously?

$$\text{Present value} = \sum_{i=1}^n P(t_i)C(t_i)$$

Robert Jarrow (2004) provides the answer. The necessary and sufficient conditions for using the present value formula directly, even when the issuer can default, can be summarized as follows:

- The present value formula can be used correctly if the recovery rate δ is constant and the same for all bonds with the same seniority, the assumption used by Jarrow and Turnbull (1995).
- The present value formula can be used if the recovery rate is random and depends only on time and seniority.
- If the default process is a Cox process (like the reduced form models in Chapter 16), then a sufficient condition to use the present value formula is for recovery to be a fraction of the bond's price an instant before default.

In reduced form models, under these assumptions then, the present value formula can be used in a straightforward way. In the Merton model and its variations, things are not so easy, and that is why it has been relegated to our review of legacy credit models, models past their prime, in Chapter 18. Use of the Merton model for enterprise risk management is not accurate or practical and will not be covered in the rest of this book except in passing when reviewing the problems of using the copula method of portfolio simulation for valuation of collateralized debt obligations (see Chapter 20).

Under the Jarrow (2004) conditions, the valuation of the previous structures (a bullet bond, an amortizing loan, and a floating-rate loan) is exactly the same for credit risky bonds as it is for risk-free bonds. Only the values in the bushy tree would be different, because the zero-coupon bond inputs are different for a credit risky borrower.

What happens when the Jarrow conditions are not met? We deal with that special case in our chapter on American fixed income options, Chapter 27.

SIMULATING THE FUTURE VALUES OF BONDS WITH CREDIT RISK

Simulating the value of risky bonds involves a straightforward number of steps using the reduced form models. For the reduced form model, the steps are as follows:

1. Simulate the risk-free term structure as given earlier.
2. Choose a formula and the risk drivers for the liquidity component of credit spread as discussed in Chapter 17.

3. Choose a formula and the risk drivers for the default intensity process in the Jarrow model as discussed in Chapter 16.
4. Simulate the random values of the drivers of liquidity risk and the default intensity for K risk factors and M time periods over N scenarios, consistent with the risk-free term structure.
5. Calculate the default intensity and the liquidity component at each of the M time steps and N scenarios. Note that these will be changing randomly over time because interest rates and other factors impact each of them. This is essential to capturing cyclicalities in bond prices and defaults.
6. Apply the Jarrow model as we have done in Chapters 16 and 17 to get the zero-coupon bond prices for each maturity relevant to the bond.
7. Calculate the bond's value at each of the M time steps and N scenarios, conditional on the company not defaulting in that period or in a prior date.

This is only slightly more complicated than the risk-free bond simulation. It relies on the identical interest rate scenarios and simulation process. Note that if the simulation is done for many companies with common risk factors driving default, we will see correlated default behavior that we have discussed throughout this book.

VALUING THE JARROW-MERTON PUT OPTION

If we have a portfolio of straight bonds, either fixed rate or floating rate, on the asset side of our financial institution, we are almost ready to value the Jarrow-Merton put option on the assets of the firm. We defer that until we reach Chapter 21, but the goal is near at hand. We move first to the discussion of credit default swaps and collateralized debt obligations, because they help set a more general portfolio context for the problem.

Credit Derivatives and Collateralized Debt Obligations

The credit crisis of 2006 to 2011 represented one of the greatest episodes of financial hysteria seen since the Nikkei stock price index (which traded at 8,734.62 on June 25, 2012) reached an intraday peak of 38,957.44 on December 29, 1989. The introduction to this book summarizes much of this mass hysteria, and it is very difficult to resist the temptation to make this a “What were they thinking?” chapter when recounting the tens of billions of dollars lost in the credit markets during this period. As best we can, we will try to reemphasize some common risk management fallacies that we raised in the introduction to this book. We reprise those fallacies in this chapter with specific reference to credit default swaps (CDS) and collateralized debt obligations (CDOs) and how Wall Street responded to market participants who clung to these fallacies as if they were true:

- “If it hasn’t happened to me yet, it won’t happen to me, even if it’s happened to someone else.”

Wall Street response: If everyone thinks home prices won’t go down, sell them 100 percent loan-to-value mortgage loans until they learn otherwise.

- “Silo risk management allows my firm to choose the ‘best of breed’ risk model for our silo.”

Wall Street response: If the credit silo can’t model the value of securities whose default is tied to macroeconomic factors such as home prices, sell them mortgage-related CDOs.

- “I don’t care what’s wrong with the model. Everyone else is using it.”

Wall Street response: After convincing many to use the flawed copula method, Wall Street stuffed investors with CDOs until they finally figured out the model didn’t work.

- “I don’t care what’s wrong with the assumptions. Everyone else is using them.”

Wall Street response: If everyone is using historical value at risk during a period of rising home prices, work fast to stuff them with mortgage-backed securities after home prices begin to fall.

- “Mathematical models are superior to computer simulations.”

Wall Street response: Try to stuff the firms relying on “F9 Model Monkeys” (see the upcoming Felix Salmon 2012 reference) before they realize that only the kind of sophisticated simulation we describe next can accurately measure the risk.

- “Big North American and European banks are more sophisticated than other banks around the world and we want to manage risk like they do.”

Wall Street has long used the sales strategy “real men buy [insert the name of an overpriced derivative]” with great success. Once buy-side firms stopped buying CDOs, traders used this strategy on their own management teams, loading up CDO structuring firms like Citigroup and Merrill Lynch with unsalable CDO tranches in the hopes that bonuses would be paid before their strategy was discovered.

- “Goldman says they do it this way and that must be right.”

Naïve investors have long regarded their investment bankers as someone with whom they need “a relationship.” Too many look at a Wall Street salesperson wearing a suit and envision that person as a saintly combination of wise finance professor and Mother Teresa. Sadly, the investment bankers have proven in the 2006 to 2011 credit crisis that their view of a client relationship is the same type of relationship that they have with a cheeseburger at McDonald’s (where they rarely have to eat).

Too many journalistic classics have emerged from this human comedy to summarize them all here.¹

With this real-world background out of the way, we turn to the theory and practice of the credit default swap and collateralized debt obligation markets. As we often do during this volume we highlight flawed approaches with the Model Risk Alert label.

CREDIT DEFAULT SWAPS: THEORY

The focus of this chapter is on pricing, valuation, hedging, and simulation of CDS and CDOs with the same ultimate goal of this volume: to measure and hedge risk in a fully integrated way, as captured by the Jarrow-Merton put option that is discussed in Chapter 1. We don’t focus on the nitty-gritty details of transaction terms and structures, which are covered in great depth by a number of other others including Tavakoli (2003), Ranson (2003), Meissner (2008), and Raynes and Rutledge (2003). Our focus here is on analyzing credit derivatives in the context of a highly accurate enterprise-wide risk management framework with a complete integration of credit risk, market risk, liquidity risk, and interest rate risk.

Credit default swaps have been highly touted as a credit risk management tool because of the difficulty in hedging single-name credit risk by shorting the common stock of the issuer, which we discussed in Chapter 18, or the publicly traded bonds of the issuer. Clearly, if shorting the bonds were easy to do, one would either sell the debt issue being hedged outright or sell it short. Short sales of risky debt are rarely an option because of the illiquidity of the bond market (which is explicitly incorporated in the Jarrow model of risky debt discussed in Chapter 16). For that reason, the development of the credit default swap market was initially greeted with enthusiasm by market participants over the past decade. This enthusiasm has faded, however, as we explain in the next section.

Two credit default swap structures are common in the marketplace. In the first structure, which we call the traditional structure, the credit protection provider pays

the purchaser of credit protection the “recovery amount” associated with the notional principal of the credit default swap with a specified delay after a specified time frame from the carefully defined event of default. This recovery amount is determined by an auction process overseen by ISDA.org. In the second structure, the digital default swap, the provider of credit protection essentially pays \$1 upon the occurrence of an event of default. By comparing the pricing on the two structures for the same issuer and the same maturity, one can imply the recovery rate and loss given default implicit in market prices. The latter structure is quite rare, but it is of great academic interest because it ignores the thorny issues of recovery.

The conceptual benefits of these derivatives structures to financial institutions are powerful, but we summarize the increasingly dramatic problems in the next section: lack of liquidity, market concentration, and potential market manipulation like that mentioned in Chapter 17 with respect to the LIBOR market.

Initially, the market for credit default swaps called for a periodic payment in advance (the insurance premium) by the purchaser of credit protection. Up-front insurance payments by the purchaser of credit protection were restricted to troubled reference credits. After the credit default swap market “big bang” in 2009, up-front payments with more modest periodic payments became the norm for credits of all quality ranges.

Jarrow (1999) provides a valuation framework for credit default swaps that explicitly allows for the fact that there may be illiquidity in the bond markets, which is, in return, reflected in risky bond prices. We reviewed the essentials of the Jarrow model in Chapter 16. The hazard rate, or “default intensity” in the Jarrow model is given by a linear combination of three terms:

$$\lambda(t) = \lambda_0 + \lambda_1 r(t) + \lambda_2 Z(t)$$

Recall that the first term, λ_0 , can be made time dependent, extending the credit model exactly like we extended term structure models in Chapters 6 through 9 and Chapter 13. The $r(t)$ term is the short rate of interest in a sophisticated term structure model. The term $Z(t)$ is the “shock” term with mean zero and standard deviation of one which creates random movements in the macroeconomic factor(s) (such as home prices in the case of the 2006 to 2011 credit crisis), which drives default for that particular company.² Movements in this macro factor are generally written in this form:

$$dM(t) = M(t)[r(t)dt + \sigma_m dZ(t)]$$

We start our valuation review with the simplest case: a digital default swap with a constant continuous default intensity. Jarrow shows that the value of a digital default swap when the default intensity λ_0 is constant can be written

$$I_t(T)1_{\{\tau > t\}} = \lambda_0 \int_t^T v(t, s; i) ds = \lambda_0 \int_t^T p(t, s) e^{-\lambda_0(s-t)} ds$$

In this expression, the expression $v(t, s, i)$ refers to the time t zero-coupon bond of the risky issuer with maturity s and seniority i . The expression $p(t, s)$ is a risk-free

zero-coupon bond as of time with maturity s . The integrals in this expression recognize that default can occur any time between time t and the maturity of the credit protection at time T . Jarrow shows how this expression changes as interest rates and macro factors are allowed to drive the default intensities. The first of the two equivalent expressions makes it clear that the value of the digital default swap is the continuous time equivalent of the promise to pay \$1 at each point in time that default could occur, weighted by the probability that default can occur at that time. When the default swap is the more conventional type, where the amount of recovery δ comes into play, the formula needs to be modified. The nature of the modification depends on the assumptions about recovery: both timing and sensitivity of recovery to the macro-economic factors that drive default. For a detailed discussion, see Jarrow (1999) and Chapter 16. In the Jarrow-Turnbull model, we discussed the valuation of a zero-coupon bond in which the holder received either \$1 at maturity or a recovery δ . The Jarrow-Turnbull model can be thought of as a “credit default swap” zero-coupon bond where the credit protection provider pays \$0 (instead of \$1) if there is no default and δ if there is a default. The value of this simple credit default swap is a slight variation on the Jarrow-Turnbull defaultable zero-coupon bond valuation formula from Chapter 16:

$$v(t, T) = [e^{-\lambda\mu(T-t)} + (1 - e^{-\lambda\mu(T-t)})\delta]P(t, T)$$

The first term in brackets is the probability of no default during the term of the zero-coupon bond. The second term in brackets is probability of default during the term of the bond. If there is no payoff in the event of no default, the value of a promise to pay a credit insurance payment of δ at the scheduled maturity date (instead of the date of the event of default) if default occurs during the life of the bond is

$$v(t, T) = [(1 - e^{-\lambda\mu(T-t)})\delta]P(t, T)$$

The Jarrow formulation discussed earlier allows for payment of \$1, not δ , and the credit payment is made at the time of default, not the scheduled maturity date.

In modeling the valuation, loss distribution, and cost of the Jarrow-Merton put option of a credit portfolio that consists of many underlying borrowers or “reference names,” market participants frequently confuse the meaning of the term *correlation*. Correlation can mean three things, and in two of the three cases the term has a model-specific meaning. Jarrow and van Deventer (2005) review the links between these correlation concepts:

- *Correlation in the value of company assets* in the Merton framework, which is definitely not the same as the other two meanings of the term correlation. Note that the correlation in the value of company assets between firms A and B can be 100 percent, but the correlation in the events of default will not be 100 percent except when the two firms A and B have identical default probabilities.
- *Correlation in default intensities* due to dependence on common macro factors like interest rates and the S&P 500 in a U.S. context, the $r(t)$ and $Z(t)$ terms given in the Jarrow formulas above. The correlation between the default intensities of companies A and B in the reduced form modeling context depends on the magnitude of the parameters λ_0 , λ_1 , and λ_2 for each of the two companies A and B.

- *Correlation in the events of default, not the default probabilities themselves.* This is the correlation between the two vectors for Companies A and B, which have a zero in a period when the company did not default and a one when the company does default. For two solvent companies, obviously, these historical time series vectors have nothing but zeros so the historical correlation in the events of default cannot be implied by the historical data on events of default for Companies A and B.

We will try to make these distinctions clear in what follows.

How does one model total credit risk in a portfolio of companies in a reduced form model? Jarrow (1999) shows that the present value of a contract, which pays \$1 if both firms 1 and 2 default prior to time T has the following formula, where the subscripts for 1 and 2 denote the firm:

$$V_t 1_{\{\tau_1 > t, \tau_2 > t\}} = p(t, T) - E_t \left(e^{-\int_t^T [r(u) + \lambda_1(u)] du} \right) - E_t \left(e^{-\int_t^T [r(u) + \lambda_2(u)] du} \right) + E_t \left(e^{-\int_t^T [r(u) + \lambda_1(u) + \lambda_2(u)] du} \right)$$

The expression E_t denotes the risk-neutral expected value of the expressions in brackets as of time t . Jarrow goes on to show how closed form solutions can be obtained to value basket swaps in the reduced form model context. We can arrive at the same conclusion by simulating the macro factors driving the default intensities in the reduced form models, simulating default/no default in a multiperiod way, and taking the correct risk-adjusted present value of the cash flows that result. We outline this procedure below in the context of collateralized debt obligations.

The Jarrow formula above shows us how to model counterparty credit risk. Assume that Company A is buying digital default swap protection on Company C with Company B as the insurance provider. We know that this transaction pays:

- 0 if Company A defaults before Company B or Company C (since Company A defaults on the credit insurance premiums).
- 0 if Company B defaults before Company C.
- \$1 if Company C defaults (1) before Companies A and B but (2) before the credit protection contract expires.

We can use this knowledge to simulate the counterparty credit risk using either the Merton or Jarrow reduced form model framework discussed above. Implicit in the formulas above are a very key assumption—that markets are perfectly liquid and that no volume of trading in credit default swaps or (Model Risk Alert) risky debt will change their price. Recent events involving JPMorgan Chase illustrate the tenuous nature of this assumption. We turn to the implications in the next section.

CREDIT DEFAULT SWAPS: PRACTICE³

The May 2012 sight of the mighty JPMorgan Chase, wallowing in a credit default swap position from which it cannot exit, raises many questions about the valuation

formulas from the previous section. We know from Chapter 17 that the credit spread, both in the bond and CDS markets, reflects the intersection of supply and demand. None of the complications discussed in Chapter 17 is represented in the previous digital default swap formula, which assumed perfectly liquid markets with no impact on the price of credit protection from trade size. One of the key real-world questions is a particularly simple one: Is the arena for credit default swap trading a marketplace or a mud pit? In this section, we sadly conclude that the answer is “mud pit” and outline some of the very many steps that are necessary for credit default swap arena to become a marketplace where the assumptions of the previous section have a better chance of being accurate. Without these reforms, the CDS mud pit is likely to suffer the same fate as collateralized debt obligation-total irrelevance, as we discuss next.

JPMorgan Chase’s May 8, 2012, announcement of \$2 billion in losses from a trade that the firm could not easily exit turned the spotlight on the arena for credit default swaps. How can a sophisticated market participant like JPMorgan Chase end up suffering such discomfort and illiquidity in the U.S. financial markets? We explain that in this section. We focus on single-name credit default swaps, not the credit default swap index market in which JPMorgan Chase has suffered its losses. The reason for our focus is also simple: The CDS index market is built on the foundation of the single-name CDS market. If the single-name CDS market is a mud pit, so is the index market.

What do we mean when we ask if the CDS market is a mud pit? We are in essence asking the following questions:

- Is there a transparent and open transmission of traded prices, bids, and offers in real time by reference name like there is for common stocks and options?
- Is there reliable reporting of trade volume in real time like there is for common stocks and options?
- Is there sufficient competition in the market to prevent market manipulation by a concentrated set of dealers, like there is for common stocks?
- If the market is highly concentrated, is there sufficient regulatory scrutiny to prevent rampant market manipulation?
- Is there sufficient regulatory scrutiny to prevent CDS trading on inside information?

Sadly, the answer to all of these questions is no. The unfortunate answer to another key question is yes—that question is “Has there been tacit cooperation among market participants and data vendors to preserve the status quo in the CDS mud pit?”

We briefly deal with each of these questions in turn:

- Is there a transparent and open transmission of traded prices, bids, and offers in real time by reference name like there is for common stocks and options?

In spite of changes in reporting by the Depository Trust & Clearing Corporation (DTCC) in recent years, the single-name CDS bids, offers, and traded prices are tightly controlled by a small cadre of dealers, market data vendors, and DTCC itself. The data is not freely available on a website like Yahoo! or Google in the way that common stock prices or even bond prices are. While one broker who distributes CDS data does time stamp bids, offers and trades in its

data service, they are the exception, not the rule, and they see a relatively small market share of total activity. The primary product offered by the largest CDS data vendor does not distinguish between bids, offers, and traded prices. Moreover, the data service typically lists CDS spreads for 2,000 corporates and sovereigns even though there have never been more than 1,000 reference names traded in a given week since DTCC began reporting trade volume (but not bids, offers, or traded prices) since the week ended July 16, 2010.

- Is there reliable reporting of trade volume in real time like there is for common stocks and options?

No. DTCC has this information, but it makes trade volume available on a weekly basis with a two or three business day lag, starting with the week ended July 16, 2010. Through January, 2012, the terms of use agreement for that volume data made it illegal for the authors to tell readers how many single-name CDS were traded in JPMorgan Chase because of this phrase:

You agree to treat any Report containing data on a specific entity, rather than aggregate position or transaction activity or other aggregate data, as confidential, or, if you are a regulator or governmental entity, in accordance with any statutory confidentiality requirements applicable to you.

Fortunately, since January, 2012, this phrase has been removed.

- Is there sufficient competition in the market to prevent market manipulation by a concentrated set of dealers, like there is for common stocks?

No. In a recent blog, using the standard measures of market concentration by the Department of Justice and volume data from the Office of the Comptroller of the Currency, we showed that the market for credit default swaps has been “highly concentrated” since its inception.⁴ The European Commission announced on April 29, 2011, that it was launching two antitrust investigations into the CDS market.⁵

- If the market is highly concentrated, is there sufficient regulatory scrutiny to prevent rampant market manipulation?

No. In fact, a courageous journalist recently described how to manipulate the credit default swap market in this post from three years ago.⁶

- Is there sufficient regulatory scrutiny to prevent CDS trading on inside information?

No. The SEC launched the first action alleging insider trading in 2009, but it was ultimately unsuccessful.⁷ The fact that insider trading is a serious problem was highlighted by this announcement from the UK’s Financial Services Authority in 2011.⁸ That leads us to perhaps the most saddening question of those posed above.

- Has there been tacit cooperation among market participants and data vendors to preserve the status quo in the CDS mud pit?

Yes. Perhaps the most egregious form of cooperation is the effort to preserve the impression that there is active trading in a large number of reference names when in fact there is not.⁹

Breathless reporters or rating agencies claim “Dell’s CDS widen 42%” when, in fact, there were only 9.6 gross trades per day and 1.75 nondealer trades per day in

Dell during the week ended May 25, 2012, according to DTCC.¹⁰ Reporters need a story, and the CDS mud pit provides material. Rating agencies need a product that is not a rating, and the CDS mud pit provides one. CDS data vendors could not sell CDS data if it were well known that (1) 81.68 percent of trades in the DTCC trade warehouse were trades between dealers or (2) more than 50 percent of the reference names for which data were reported by the vendor had *no* trades in a given week, on average.

Existing beneficiaries of the CDS mud pit have no incentive to change things and lots of incentives not to change things. These changes must be forced by laws and regulations. The necessary changes should emulate the information and data flow of the market for common stocks and options. These markets, while not yet perfect, are 1 million times more transparent and competitive than the CDS mud pit:

- Distribute data through a large number of data vendors as the New York Stock Exchange does
- Display bid levels and bid volumes in real time
- Display offered levels and offered volumes in real time
- Display traded prices and traded volumes in real time
- Require disclosure of all entities with large open positions, both long and short, as is required in the market for common stock
- Aggressively enforce insider trading rules
- Force trading onto legitimate exchanges. Industry resistance to this is long standing, but that is irrelevant. The investors' interest is of greater import.
- Use the Department of Justice whenever necessary to establish competition in the marketplace

Until then, whenever one sees a credit default swap quotation, it should be regarded with an extremely high degree of suspicion. With this sobering information as background, we now turn to the rise and fall of the market for collateralized debt obligations.

COLLATERALIZED DEBT OBLIGATIONS: THEORY

Collateralized debt obligations are generalized versions of first to default swaps, second to default swaps, third to default swaps, and so on. In a typical CDO structure, there may be 100 to 150 underlying reference names with a (roughly) equal amount of (actual or notional) principal associated with each of the reference names. Synthetic CDOs have single-name credit default swaps as the underlying reference instruments. Cash-flow CDOs have financial instruments such as bonds, mortgage-backed securities, asset-backed securities, commercial mortgage-backed securities, or other instruments as the underlying collateral. If the underlying collateral is in the form of bonds issued by the 100 reference names, the proceeds of the issuance of senior, subordinated, and equity tranches of the CDO are used to buy these bonds. If the underlying "collateral" is in the form of credit default swaps (a synthetic CDO), then the proceeds of the issuance of the tranches are used to buy a bond of a single issuer, the proceeds of which are paid to the CDO tranche holders over time depending on how many of the reference names default. The fact that all of

the cash in the synthetic structure is invested in one reference name is a huge factor in differentiating between the so-called “cash flow CDOs” (the bond structure) and the synthetic CDO structure. This concentration of risk must not be ignored in CDO analysis.

The value and cash flow generated by the underlying reference names in the CDO is identical in its nature to what we discussed in Chapter 19, when we discussed simulating future values of credit risky bonds, and to the first part of this chapter. The only difference is that the cash flow is parceled out according to the priorities of the tranches. If there is only one tranche (i.e., the equity tranche), CDO analysis is no different from credit risk portfolio analysis. If there are two tranches, the first cash flows go to the senior tranche as specified in the CDO indenture (the waterfall) and the remaining cash flows (if any) go to the equity tranche, and so on. During the heyday of the CDO market, waterfall structures with 5, 10, or 15 layers or tranches were common.

The simulation of which cash flows go to whom is at the heart of CDO analysis, as is the simulation of correlated defaults. The investment banking community and the rating agencies have aggressively advocated a method of analysis that dramatically understates the worst-case losses from a CDO portfolio, and users of this analysis should recognize that (if they take the analysis seriously) they will grossly overpay for CDO tranches. (See the last section of this chapter, on the copula approach, for the details of this legacy calculation.) Felix Salmon (2012) labeled the diehard users of this approach “F9 Model Monkeys” in honor of the key in common spreadsheet software that would initiate the lumbering spreadsheet calculation of the copula method.

The correct procedure for valuing credit default swaps (traditional or digital), first to default swaps, and CDOs is to follow a slight variation on the process we outlined in a corporate bond context in Chapter 19. For the reduced form model, the steps are as follows:

1. Simulate the risk-free term structure of interest rates as given in Chapters 6 through 9 and summarized in Chapter 10. This should be a simulation with 5 to 10 risk factors driving the term structure of interest rates.
2. Choose a formula and the risk drivers for the liquidity component of credit spread as discussed in Chapter 17.
3. Choose a formula and the risk drivers for the default intensity process in the Jarrow model as discussed in Chapter 16. The macroeconomic factors driving correlated risk can be estimated using historical default data and historical macroeconomic factors using logistic regression or from observable market prices of bonds. Twenty-six macroeconomic factors are embedded in the instructions for the Comprehensive Capital Analysis and Review 2012 of the Federal Reserve. They provide a good starting point for any serious analyst of risk:
 - a. BBB Corporate Yield
 - b. CPI Inflation Rate
 - c. Developing Asia Inflation Rate
 - d. Euro Area Inflation Rate
 - e. Japan Inflation Rate
 - f. UK Inflation Rate

- g. Commercial Real Estate Price Index
 - h. Dow Jones Total Stock Market Index
 - i. Developing Asia Bilateral Dollar Exchange Rate
 - j. Euro Area Bilateral Dollar Exchange Rate
 - k. Japan Bilateral Dollar Exchange Rate
 - l. UK Bilateral Dollar Exchange Rate
 - m. House Price Index
 - n. Mortgage Rate
 - o. Nominal Disposable Income Growth
 - p. Nominal GDP Growth
 - q. Real Disposable Income Growth
 - r. Real GDP Growth
 - s. Developing Asia Real GDP Growth
 - t. Euro Area Real GDP Growth
 - u. Japan Real GDP Growth
 - v. UK Real GDP Growth
 - w. Ten-Year Treasury Yield
 - x. Three-Month Treasury Yield
 - y. Unemployment Rate
 - z. Market Volatility Index—VIX
4. Simulate the random values of the drivers of liquidity risk and the default intensity for M time periods over N scenarios, consistent with the risk-free term structure.

Model Risk Alert

One should take care not to make common risk-management assumptions (i.e., that the macro factors have normally distributed returns and are independent from period to period) when those common assumptions are not true.

5. Calculate the default intensity and the liquidity component at each of the M time steps and N scenarios. Note again, as in Chapter 19, that these will be changing randomly over time because interest rates and other macro factors are changing at each time step of each scenario. The simulation of these macro factors is essential to capturing cyclicalities in bond prices and defaults that create correlated defaults and which have a major impact on the equity tranches and junior tranches of CDOs.
6. Simulate the default/no default of each reference name at each of the M time steps and N scenarios.
7. Calculate the payments made on the derivative structure given the number and amount of defaults to that date in each scenario using the “waterfall” described in the indenture to that structured product.
8. Apply the Jarrow model as we have done in Chapters 16 and 17 to get the zero-coupon bond prices for the risk-free curve and for each reference name in the CDO collateral pool or other portfolio. This calculation is very relevant to tranche value. It was fairly common for naïve investors to argue a CDO tranche was still worth par “because no principal or interest payment had been missed,” even though the underlying reference collateral was worth far less than par.

9. Calculate the derivative's value at each of the M time steps and N scenarios, conditional on the number of defaults among the reference names that have occurred to that point in the simulation.

This is slightly more complicated but not that much more complicated than the credit risky bond simulation discussed in Chapter 19 and the risk-free yield curve simulation discussed in Chapters 6 through 10. It relies on the identical interest rate scenarios and simulation process. Note that if the simulation is done for many companies with common risk factors driving default, we will see correlated default behavior that we have discussed throughout this book.

COLLATERALIZED DEBT OBLIGATIONS: A WORKED EXAMPLE OF REDUCED FORM SIMULATION

In this section, we review how to simulate the correlated default risk of the reference names underlying a credit portfolio or a collateralized debt obligation. The process is identical for retail borrowers, small businesses, public firms, and sovereign credits. Only the driving risk factors and coefficients differ. For simplicity, we focus on two hypothetical corporations, ABC Company and XYZ Company. We assume we have completed an extensive historical analysis that proves ABC Company's one-year default probability is sensitive to the two-year change in home prices and the real growth rate in gross domestic product, as shown in Exhibit 20.1.

Similarly, XYZ Company is sensitive to a separate set of macroeconomic factors. A combination of historical analysis and business judgment establishes that its one-year default probability has the sensitivity to the two-year change in home prices and the 10-year yield on government bonds as shown in Exhibit 20.2.

In a reduced form modeling context, theoretical default probabilities are usually modeled in discrete blocks of time, and the default rate in each period is normally given as a logistic function of macroeconomic factors and other risk factors that move randomly. We outlined this procedure in the prior section and in Chapter 16. For simulation purposes, default probabilities are normally generated for discrete daily, weekly, monthly, quarterly, or annual time periods. For simulation purposes, as we discussed in Chapter 16, the logistic formula is an attractive choice for modeling because:

EXHIBIT 20.1 ABC Company One-Year Default Probability and Response to Macro Factors

2-Year Home Price Return	Real Growth Rate of Gross Domestic Product				
	−4.00%	−2.00%	0.00%	2.00%	4.00%
40.00%	2.46%	2.32%	2.19%	2.06%	1.95%
20.00%	3.63%	3.42%	3.23%	3.05%	2.87%
10.00%	4.39%	4.15%	3.92%	3.70%	3.49%
0.00%	5.32%	5.02%	4.74%	4.48%	4.23%
−10.00%	6.42%	6.07%	5.73%	5.42%	5.12%
−20.00%	7.73%	7.31%	6.91%	6.54%	6.18%
−40.00%	11.11%	10.53%	9.98%	9.45%	8.95%

EXHIBIT 20.2 XYZ Company One-Year Default Probability and Response to Macro Factors

2-Year Home Price Return	10 Year Government Yields				
	1.00%	3.00%	5.00%	7.00%	9.00%
40.00%	1.03%	1.09%	1.15%	1.22%	1.30%
20.00%	2.25%	2.39%	2.53%	2.69%	2.85%
10.00%	3.32%	3.52%	3.73%	3.95%	4.19%
0.00%	4.88%	5.17%	5.47%	5.79%	6.12%
−10.00%	7.11%	7.52%	7.94%	8.39%	8.87%
−20.00%	10.25%	10.81%	11.41%	12.03%	12.68%
−40.00%	20.26%	21.25%	22.27%	23.33%	24.42%

- It is the maximum likelihood estimator for 0/1 problems like default/no default.
- It will never produce simulated default probabilities outside of the range from 0 to 100 percent.

The logistic function takes the form

$$P[t] = \frac{1}{1 + e^{-\alpha - \sum_{i=1}^n B_i X_i}}$$

where $P[t]$ is the unannualized default probability over some discrete interval. The variables X_i are the explanatory variables and the alphas and the betas are the best-fitting coefficients that produce the maximum likelihood estimates of default probabilities.

For ABC Company, the coefficients that produce the table of default probabilities above are consistent with this logistic function for one-year time intervals:

$$P_{ABC}[t] = \frac{1}{1 + e^{-(-3) - (-3X_1 - 2X_2)}}$$

X_1 is the annual growth rate in real gross domestic product and X_2 is the two-year change in home prices, both expressed as a decimal. These coefficients can be derived either from a historical database, like the sample we showed for Lehman Brothers in Chapter 16, or they can be derived from the numbers in Exhibit 20.1 using the solver function in common spreadsheet software.

For XYZ Company, the coefficients that produce the default probabilities above as a function of 10-year government yields and the two-year change in home prices are given here:

$$P_{XYZ}[t] = \frac{1}{1 + e^{-(-3) - (3X_2 - 4X_3)}}$$

X_2 , as before, is the two-year change in home prices and X_3 is the 10-year government yield, both expressed as a decimal.




EXHIBIT 20.3 Sensitivity of ABC Company and XYZ Company One-Year Default Probabilities to Two-Year Returns on Home Prices

We can simulate these default probabilities forward for M time steps in N scenarios by generating N scenarios for the three macro factors that drive defaults for these two firms:

- The growth in real gross domestic product
- The two-year change in home prices
- The level of 10-year government yields

ABC Company and XYZ Company have default probabilities that we know are correlated because they have a common dependence on the two-year change in home prices as given by the previous tables. Assuming the current real GDP growth is 0 percent and that 10-year government yields are 5.00 percent, a stress test of default probabilities for both firms with respect to the 2-year change in home prices is as shown in Exhibit 20.3.

ABC Company and XYZ Company may also have other sources of implicit correlation if the growth in real domestic product (which directly affects only ABC Company) and the 10-year government yield (which directly affects only XYZ Company) are correlated.

Simulating these macro factors forward in an accurate way requires a careful analyst to specify:

- The probability distribution for each factor's random movements

Model Risk Alert

Often these distributions are not normally distributed, although that assumption is common.

- Whether or not the distribution of the factor in period J is in fact independent of its value in period $J - 1$

Model Risk Alert

Often, they are *not* independent, although that independence assumption is common.

- The correlations between the random factors themselves

In our N -period simulations over M periods, we can value a portfolio of securities for ABC Company and XYZ Company if we know how spreads move, conditional on the level of default probabilities for each firm and the three previous macro factors. We discussed this issue in Chapter 17.

We can value the total collateral in a credit portfolio or collateralized debt obligation by adding the value of the securities for ABC, XYZ, and all other reference names included in the collateral pool. With the default/no default data and valuation data at each of the M time periods in the N scenarios, we can accurately allocate cash flow and value to each tranche of a CDO and calculate its value. In an efficient and competitive market (which the CDO market is not), the sum of the value of all tranches (less the structuring fee of the sponsor of the CDO) should equal the sum of the value of the reference collateral. Felix Salmon's "F9 Model Monkeys" sought to approximate the correct answer in spreadsheet software using the copula approach that we discuss in an upcoming section. There is no need to settle for that kind of extremely crude approximation when this straightforward procedure gives an exact answer conditional on the macro factors and parameters chosen.

This simple example shows that using reduced form default models as part of a CDO analysis or as part of a comprehensive enterprise-wide risk management simulation is powerful and straightforward. It allows the analyst to produce government-mandated stress tests with respect to each macro factor that is important in driving defaults. The macro factor links are explicitly derived from either historical data or current market prices and their history. Any simulation approach that uses ratings, which we discussed in Chapter 18, as part of the simulation will not have the same degree of accuracy or transparency with respect to the macro factors that drive risk. The reason is that the rating agencies have been unable to articulate even the time horizon over which ratings apply, let alone the exact formulas by which macro factors affect ratings. With the reduced form approach, the linkages are clear and based on best practice econometric methodology.

COLLATERALIZED DEBT OBLIGATIONS: PRACTICE

We now turn from best practice risk analytics to the horrifying reality of what took place in the CDO market during the credit crisis of 2006 to 2011. The spectacular rise and fall of the collateralized debt obligation market is summarized in the press release "Global CDO Issuance" of the Securities Industry and Financial Markets Association:

Global CDO Issuance Billions of U.S. Dollars	
2004	157.4
2005	251.3
2006	520.6
2007	481.6
2008	61.9
2009	4.3
2010	8.0

The 2010 issuance total is for a partial year. It is obvious from this issuance history that there has been a complete rejection of the CDO instrument by the investing public, which experienced tens of billions of dollars of losses that were partially chronicled in the introduction to this book. The reasons for the market's demise have been richly described by journalists and government studies. Perhaps the most comprehensive U.S. government review of the abuses that led to the market's demise is the Levin report mentioned earlier in this book.¹¹

The authors strongly recommend that this report be reviewed in full, at least three times, by any investor willing to consider the purchase of any security that looks or smells like a CDO. At more than 600 pages in length, the Levin report is the definitive narrative of the greed, corruption, and misbehavior by “F9 Model Monkeys” that caused the biggest financial collapse of a generation.

THE COPULA METHOD OF CDO VALUATION: A POSTMORTEM

We close this chapter with a eulogy for the copula approach to CDO simulation. The copula approach to valuation of collateralized debt obligations has been blamed for much of the credit crisis. The recent articles in *Wired*, *Mother Jones*, *Wall Street Journal*, and so on listed earlier in this chapter are three very entertaining examples of this genre. Rather than placing blame, other commentators emphasize the positive, saying “There’s plenty of blame to go around, let’s not waste time assigning blame, let’s move forward.” This section takes a different view than either group. We argue “Formulas don’t cause losses, people cause losses.” There are lots of lessons to be learned from the copula formula’s uses and abuses.

This list of lessons is not definitive, but it highlights many of the reasons why and how financial engineers and management went astray in CDO valuation:

Lesson 1: Senior Management Knew or Should Have Known the Copula Approach Was Flawed, but They Did Nothing

Robert Rubin famously commented that he did not have “enough experience” on Wall Street to understand the liquidity puts of CDO tranches that were at the heart of Citigroup’s CDO related losses. But we are sure Mr. Rubin reads the *Wall Street Journal*, which on August 12, 2005, ran an excellent story on page 1 by Mark

Whitehouse, “Slices of Risk: How a Formula Ignited Market That Burned Some Big Investors.” The story described David Li’s role in bringing the use of copulas into finance, and it then went on at length on how the technique had failed in the May 2005 period, causing losses of hundreds of millions of dollars. Did the heads of the audit committees of the Boards of Directors of Citigroup, Merrill Lynch, and UBS call in the CEO and chief risk officers and grill them on the use of the copula method in their rapidly growing CDO businesses? Did Charles Prince or Stanley O’Neal read that article and put on the brakes? No, they did nothing. People cause losses, not formulas, even when one of the world’s most important business newspapers warned us all in a page 1 story.

Lesson 2: The Copula Approach Assumed Away All but One Risk Factor, but Analysts Used It Anyway

The August 12, 2005 story in the *Wall Street Journal* was very specific about why the copula approach caused such losses. It implicitly assumes that there is only one common risk factor driving the returns on company assets of the reference names underlying a CDO tranche. Other than this common risk factor, all other risk is assumed to be idiosyncratic and uncorrelated with the risks of other companies. We know this is a gross oversimplification; in fact, a March 23, 2009, press release on www.kamakuraco.com recounts how 40 macro risk factors are used for modeling correlated default on the Kamakura Risk Information Services (KRIS) Web-based credit portfolio management tool, KRIS-cpm. For CDO market participants, however, the assumption of one common risk factor was hugely attractive. It made the CDO analysis “Excel friendly” and it made the math tractable to the smartest 2 percent of Wall Street, Felix Salmon’s “F9 Model Monkeys.” The copula approach allows the kind of mathematical exposition that is attractive to academic journals and that one can find in the very interesting compilation of articles in the book *The Definitive Guide to CDOs* (edited by Gunter Meissner, Risk Publications, 2008). The cost to this is that the results from this highly simplified abstraction to reality are just not accurate. The *Wall Street Journal* put it plainly. The implication of one risk factor is that you should be able to go long one tranche of a CDO and hedge it with the proper short position in another tranche of the same CDO. This trade caused enormous losses because the implication of the copula model was plain wrong.

There was another anecdote in the story that seriously called into question the Merton model of risky debt (see Chapter 18), which is closely linked with the copula approach. The 1974 classic by Robert Merton assumes that only (Model Risk Alert) one risk factor, the value of the assets of a company, drives the prices of the firm’s debt and equity. If this is true, one should be able to buy the debt of a company and hedge it with a short position in the firm’s common stock. Unfortunately, as the *Wall Street Journal* recounted, this implication is simply not true and traders in GM and Ford lost tons of money in May 2005 using the Mertonian trading strategy. Multiple factors drive bond and common stock prices, but people did not want to hear it. In fact, Jarrow and van Deventer (1998) were very proud to receive hate mail from a prominent junk bond analyst when they made exactly this point in a 1998 publication.¹²

For a detailed study of the lack of accuracy of the highly stylized Merton model of risky debt, see Campbell, Hilscher, and Szilagyi (2008), a key reference we highlight in Chapter 18.

Lesson 3: Lots of Correlations Matter in CDO Valuation, but Copula Users Assumed Only One Correlation Matters

If there are N reference names in a CDO collateral pool, there are $N(N - 1)/2$ pairs of companies in the pool. The correlations between the default probabilities of each pair of companies are of course different. The same is true for the correlation in the returns on the assets of each pair of companies. The two concepts of correlation are mathematically linked (for the formula, see Jarrow and van Deventer 2005). Copula analysts assume (Model Risk Alert) that the pairwise correlations in asset returns are the same for every one of the $N(N - 1)/2$ pairs of companies, but this is simply not true. The fact that it is not true was well known to traders—the proof was simple. The correlation implied by observable market prices was not the same across tranches of a CDO on the same reference collateral. The assumption was simply wrong, like the assumption in the Black-Scholes model that implied volatility is constant and should be the same for stock options with any maturity and strike price.

Why did traders, analysts, and risk managers use such a simple assumption that they knew to be wrong? For many, the reason was because it was the easy thing to do. For the very smartest people on Wall Street, who knew how flawed the copula approach was, they pushed the technique as an easy way to induce investors to buy CDO tranches at the wrong price. The smart ones used a much more realistic valuation technique and they have the profitability to prove it!

Lesson 4: Default Probabilities Vary Randomly, but Copula Users Assumed Away This Randomness

There are two closely related assumptions about the default probabilities used in a copula analysis. The most common assumption is that they are (Model Risk Alert) constant. A rarer but only slightly better assumption is to assume that (Model Risk Alert) they drift over time (nonrandomly) in a way that matches an observable term structure of default probabilities, like those in the CDS market. Unfortunately, these assumptions are also dramatically untrue. As the 2006–2011 credit crisis showed, default probabilities rise and fall randomly over the business cycle, driven by common macro factors like home prices and interest rates. The upshot of assuming away this randomness is dramatic errors in valuation, particularly from underestimating the fat tails of losses when the economy turns bad. This was especially devastating for “super senior” tranches of CDOs.

Lesson 5: Default Can Happen at Any Time, but Copula Users Assumed That Away

A surprisingly large number of copula users employed a (Model Risk Alert) single-period simulation, including many third-party vendors as explained by Michel Araten of JP Morgan Chase in a presentation at the ICBI Risk Conference in December in Geneva a few years ago. In a single-period simulation, defaults can happen only at time zero or at the end of the single period, not in between. For a CDO of bank trust

preferred securities, to name one example, maturities can be more than 30 years, so this assumption about timing of defaults is extreme. The right approach is obvious—to simulate default/no default on a multiperiod basis, typically monthly. CDO market participants, however, generally did not do this. This explains the widely held but completely false belief that an increase in the single correlation in a copula model will cause the price of the equity tranche to rise (see Jarrow and van Deventer 2008). That conclusion about equity tranches is due to the assumption of a single modeling period. A multiperiod simulation was essential for accuracy, but most analysts and some vendors clung to a single-period model.

Lesson 6: CDO Analysts Can Be Slaves to Fashion, but At High Costs

Gunter Meissner's fascinating compilation of articles in *The Definitive Guide to CDOs* (Risk Publications, 2008) is definitive proof of the trendiness of the copula approach. Even though most of the chapters in the book were penned as the current credit crisis began to unfold, the overwhelming majority of the chapters endorse or extend the copula method. A reviewer on Amazon.com cited only two chapters in the book (Chapter 4 on hedging, and Chapter 16 by Jarrow and van Deventer entitled "CDO Valuation: Fact and Fiction") that were critical of the copula approach. Even as events were revealing (again) the flaws of the copula approach that the *Wall Street Journal* highlighted in 2005, many were still singing its praises. This is certainly not the fault of David Li, who first used the copula approach in finance. (It is the fault of all of us who came after Li but failed to maintain our scientific objectivity about what works and what does not, and why.)

A multiperiod simulation of CDO values with default probabilities driven up and down by macroeconomic factors is an exercise in computer science that's hard to do without sophisticated third-party risk management software. It's not fashionable, because there's no closed form solution, so you will not see many academic publications on its implications. Nonetheless, it is the only method that works. That's perhaps the most important thing we have learned from the attempt to use the copula approach for credit portfolio management.

VALUING THE JARROW–MERTON PUT OPTION

We have analyzed the reference names in a credit portfolio or in a collateralized debt obligation just as if they were loans in a bank lending portfolio or private placements on the books of a pension fund or insurance company. The next step is to integrate them with the liabilities of the institution so that we can evaluate the Jarrow-Merton put option as an integrated measure of firm risk. We now turn to that task in the remainder of this book.

NOTES

1. The authors recommend a few classics, however, both because of their timely warning and their "plain English" warning of what can and did go wrong in the CDS and CDO markets. We list them here in chronological order:

- Mark Whitehouse, “Slices of Risk: How a Formula Ignited Market That Burned Some Big Investors,” *Wall Street Journal*, August 12, 2005.
 - Felix Salmon, “Recipe for Disaster: The Formula that Killed Wall Street,” *Wired Magazine*, February 23, 2009.
 - Kevin Drum, “The Gaussian Copula,” *Mother Jones*, February 24, 2009.
 - Sam Jones, “Of Couples and Copulas,” *Financial Times*, April 24, 2009.
 - Lisa Pollack, “On the Clairvoyance of Sovereign CDS,” *Financial Times*, October 10, 2011.
 - Robert A. Jarrow, “Problems with Using CDS to Infer Default Probabilities,” Cornell University and Kamakura Corporation memorandum, January 25, 2012, www.kamakuraco.com.
 - Donald MacKenzie and Taylor Spears, “‘The Formula That Killed Wall Street’? The Gaussian Copula and the Material Cultures of Modeling,” University of Edinburgh, June, 2012.
 - Felix Salmon, “The Dangerous Gaussian Copula Function,” Reuters, June 21, 2012.
2. In more mathematical terms, $Z(t)$ is standard Brownian motion under a risk-neutral probability distribution Q with initial value 0 that drives the movements of the market index $M(t)$.
 3. An earlier version of this section appeared in the *BankThink* column of www.AmericanBanker.com on June 11, 2012, under the title “Credit Default Swaps: Market or Mud Pit?” The authors wish to thank the editors for many helpful comments.
 4. For the full text, see this link www.kamakuraco.com/Blog/tabid/231/EntryId/371/The-Credit-Default-Swap-Market-and-Anti-Trust-Considerations.aspx.
 5. European Union, “Antitrust: Commission Probes Credit Default Swaps Market,” *Rapid*, <http://europa.eu/rapid/pressReleasesAction.do?reference=IP/11/509&format=HTML&aged=0&language=EN&guiLanguage=en>.
 6. Mark Mitchell, “Did The Markit Group, A Black-Box Company Partially Owned by Goldman Sachs and JP Morgan, Devastate Markets?” *Marketrap*, www.marketrap.com/article/view_article/91172/did-the-markit-group-a-black-box-company-partially-owned-by-goldman-sachs-and-jp-morgan-devastate-markets.
 7. U.S. Securities and Exchange Commission, “SEC Charges Hedge Fund Manager and Bond Salesman in First Insider Trading Case Involving Credit Default Swaps,” www.sec.gov/news/press/2009/2009-102.htm.
 8. Ben Moshinsky, “FSA Said to Review Credit-Default Swaps in Market Abuse Probe,” *Bloomberg*, www.bloomberg.com/news/2011-03-25/fsa-said-to-review-credit-default-swaps-in-market-abuse-probe.html.
 9. For a review of trading volume reported by the DTCC, see the following links:

All Reference Names:

www.kamakuraco.com/Blog/tabid/231/EntryId/365/CDS-Trading-Volume-for-1-090-Reference-Names.aspx

U.S. Banks:

www.kamakuraco.com/Blog/tabid/231/EntryId/366/Credit-Default-Swaps-and-Deposit-Insurance.aspx

Sovereigns:

www.kamakuraco.com/Blog/tabid/231/EntryId/368/Sovereign-Credit-Default-Swap-Trading-Volume.aspx

Sub-Sovereigns and Municipals:

www.kamakuraco.com/Blog/tabid/231/EntryId/367/Municipal-Credit-Default-Swap-Trading-Volume.aspx

10. Fitch Solutions, “Fitch Solutions: Dell’s CDS Widen 42%; Hewlett Packard Out 95% Ahead of Earnings,” *Yahoo! Finance*, <http://finance.yahoo.com/news/fitch-solutions-dells-cds-widen-120600373.html>.
11. See Committee on Homeland Security and Governmental Affairs, Carl Levin, Chairman, Tom Coburn, Ranking Minority Member, *Wall Street and the Financial Crisis: Anatomy of a Financial Collapse, Majority and Minority Staff Report, Permanent Subcommittee on Investigations*, United States Senate, April 13, 2011 www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&ved=0CCIQFjAA&url=http%3A%2F%2Fwww.hsgac.senate.gov%2Fdownload%2F%3Fid%3D273533f4-23be-438b-a5ba-05efe2b22f71&ei=J750UKKIC-T7iwL894C4Bw&usg=AFQjCNH4e7J5tvhekmXhnMV64RvGo3NG2g
12. Robert A. Jarrow and Donald R. van Deventer, “Integrating Interest Rate Risk and Credit Risk in Asset and Liability Management,” *DefaultRisk.com*, www.defaultrisk.com/pp_model_06.htm.

PART

Four

Risk Management Applications: Instrument by Instrument

European Options on Bonds

In this chapter, we use the three-factor Heath, Jarrow, and Morton (HJM) yield curve model to value European options on corporate bonds. As we noted in Chapter 19, Jarrow (2004) has summarized the conditions under which a risky bond issuer's (called ABC Corporation in this chapter) coupon-bearing bond issues can be analyzed as a portfolio of zero-coupon bonds using the formulas of Chapter 16 and Chapter 19. We explained in Chapter 17 that we can create today's yield curve for ABC Corporation in one of two ways. In method one, we estimate the term structure of default risk using logistic regression for ABC and then apply the techniques of Chapter 17 to estimate what the intersection of supply and demand for credit would be for ABC corporation on that day; this would be the yield curve that we would use to price ABC securities and derivatives on them. The second method is to collect data on existing ABC bond prices and apply the yield curve smoothing techniques of Chapter 5 to extract a smooth, continuous yield curve.

As an example, we looked at the net present value ("dirty price," which is price plus accrued interest) on 77 bonds for a prestigious U.S. bond issuer on August 2, 2004. We smoothed the U.S. Treasury curve using the maximum smoothness forward rate approach of Chapter 5. We then solved for the maximum smoothness forward credit spread curve that minimized the sum of squared pricing errors on the bonds of this firm, which we call ABC Corporation. The zero-coupon yield curves for the risk-free curve (lower line) and the ABC Corporation curve (upper line) are shown in Exhibit 21.1.

We could then apply an HJM interest rate model to either generate Monte Carlo simulations for both the risk-free and ABC yield curves (this is the most common and most accurate technique) or to generate an HJM bushy tree. The bushy tree approach is the preferred approach for American options. We use the three-factor HJM bushy tree from Chapter 9 for expository purposes. We assume, too, in this chapter that the bushy tree in Chapter 9 is the ABC Corporation bushy tree. We could derive it in exactly the same manner that we did in Chapter 9, but we leave that to the reader.

We will use the following information from Chapter 9:

1. The zero-coupon bonds used as input
2. The volatility assumptions for the three factors
3. The resulting zero-coupon bond prices, zero-coupon bond yields, and unpounded spot rates of interest




EXHIBIT 21.1 U.S. Treasury and ABC Corporation Zero-Coupon Yields Derived Using Maximum Smoothness Forward Rate Smoothing, August 2, 2004

Sources: Kamakura Corporation; Federal Reserve.

4. The probability of occurrence of the 4 nodes at time step 1, the 16 nodes at time step 2, and the 64 nodes that are relevant for time steps 3 and 4
5. The present-value factors for each of those nodes
6. Most importantly, the probability-weighted discount factors for each node

For all of the examples in this chapter, we price derivatives on an underlying reference bond of ABC Corporation. We choose as our reference instrument the 3 percent bond paying \$3 of interest at times 1, 2, 3, and 4 and principal of \$100 at time 4. We could price this bond in one of two ways. We could multiply the cash flows by the zero-coupon bond prices that we assumed were visible in the market. We show them again in Exhibit 21.2.

Alternatively, we could populate the cash flow table cash flows received at each point of time in each state, as shown in Exhibit 21.3.

We then multiply these cash flows by the matching probability-weighted discount factor from the table in Exhibit 21.4.

Regardless of the method chosen, we get the same values because we have been careful to construct the bushy tree so that the input zero-coupon bond prices are correctly valued using the bushy tree: 104.7071 is the value for the 3 percent coupon-bearing bond. The table we will use the most in this chapter is the term structure of zero-coupon bond prices that will prevail at time 0 (1 state), time 1 (4 states), and especially time 2 (16 states), because we will be analyzing European options with an exercise time of 2 years. We see in Exhibit 21.5 that there are 16 combinations of zero-coupon bonds that mature at time $T = 3$ and $T = 4$.

EXHIBIT 21.2 Key Assumptions

	Period of Maturity			
	1	2	3	4
Input Zero-Coupon				
Bond Prices	0.9970057865	0.9841101497	0.9618922376	0.9308550992

EXHIBIT 21.3 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		3.0000	3.0000	3.0000	103.0000
3		3.0000	3.0000	3.0000	103.0000
4		3.0000	3.0000	3.0000	103.0000
5			3.0000	3.0000	103.0000
6			3.0000	3.0000	103.0000
7			3.0000	3.0000	103.0000
8			3.0000	3.0000	103.0000
9			3.0000	3.0000	103.0000
10			3.0000	3.0000	103.0000
11			3.0000	3.0000	103.0000
12			3.0000	3.0000	103.0000
13			3.0000	3.0000	103.0000
14			3.0000	3.0000	103.0000
15			3.0000	3.0000	103.0000
16			3.0000	3.0000	103.0000
17				3.0000	103.0000
18				3.0000	103.0000
19				3.0000	103.0000
20				3.0000	103.0000
21				3.0000	103.0000
22				3.0000	103.0000
23				3.0000	103.0000
24				3.0000	103.0000
25				3.0000	103.0000
26				3.0000	103.0000
27				3.0000	103.0000
28				3.0000	103.0000
29				3.0000	103.0000
30				3.0000	103.0000
31				3.0000	103.0000
32				3.0000	103.0000
33				3.0000	103.0000

(Continued)

EXHIBIT 21.3 (Continued)

Row Number	Current Time				
	0	1	2	3	4
34				3.0000	103.0000
35				3.0000	103.0000
36				3.0000	103.0000
37				3.0000	103.0000
38				3.0000	103.0000
39				3.0000	103.0000
40				3.0000	103.0000
41				3.0000	103.0000
42				3.0000	103.0000
43				3.0000	103.0000
44				3.0000	103.0000
45				3.0000	103.0000
46				3.0000	103.0000
47				3.0000	103.0000
48				3.0000	103.0000
49				3.0000	103.0000
50				3.0000	103.0000
51				3.0000	103.0000
52				3.0000	103.0000
53				3.0000	103.0000
54				3.0000	103.0000
55				3.0000	103.0000
56				3.0000	103.0000
57				3.0000	103.0000
58				3.0000	103.0000
59				3.0000	103.0000
60				3.0000	103.0000
61				3.0000	103.0000
62				3.0000	103.0000
63				3.0000	103.0000
64				3.0000	103.0000
Risk-Neutral Value =					104.7071

EXHIBIT 21.4 Probability-Weighted Discount Factors

Row Number	Current Time				
	0	1	2	3	4
1	1.0000	0.1246	0.0152	0.0018	0.0017
2		0.1246	0.0152	0.0018	0.0017
3		0.2493	0.0305	0.0036	0.0035
4		0.4985	0.0610	0.0073	0.0067
5			0.0155	0.0019	0.0018

EXHIBIT 21.4 (Continued)

Row Number	Current Time				
	0	1	2	3	4
6			0.0155	0.0019	0.0018
7			0.0309	0.0038	0.0037
8			0.0619	0.0075	0.0072
9			0.0308	0.0037	0.0035
10			0.0308	0.0037	0.0036
11			0.0616	0.0074	0.0072
12			0.1233	0.0149	0.0141
13			0.0615	0.0073	0.0069
14			0.0615	0.0073	0.0069
15			0.1230	0.0145	0.0139
16			0.2459	0.0291	0.0271
17				0.0019	0.0018
18				0.0019	0.0019
19				0.0038	0.0037
20				0.0076	0.0072
21				0.0019	0.0019
22				0.0019	0.0019
23				0.0038	0.0038
24				0.0077	0.0076
25				0.0038	0.0037
26				0.0038	0.0039
27				0.0077	0.0077
28				0.0153	0.0150
29				0.0076	0.0075
30				0.0076	0.0075
31				0.0153	0.0152
32				0.0306	0.0298
33				0.0038	0.0037
34				0.0038	0.0037
35				0.0076	0.0075
36				0.0151	0.0146
37				0.0038	0.0037
38				0.0038	0.0037
39				0.0076	0.0075
40				0.0152	0.0147
41				0.0076	0.0074
42				0.0076	0.0076
43				0.0153	0.0151
44				0.0306	0.0296
45				0.0150	0.0143
46				0.0150	0.0148
47				0.0300	0.0292
48				0.0599	0.0572

(Continued)

EXHIBIT 21.4 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
49				0.0075	0.0073
50				0.0075	0.0075
51				0.0150	0.0148
52				0.0301	0.0290
53				0.0076	0.0074
54				0.0076	0.0074
55				0.0151	0.0149
56				0.0302	0.0292
57				0.0152	0.0148
58				0.0152	0.0149
59				0.0304	0.0299
60				0.0608	0.0587
61				0.0298	0.0285
62				0.0298	0.0291
63				0.0596	0.0580
64				0.1192	0.1135

EXHIBIT 21.5 Zero-Coupon Bond Prices

State Name	State Number	Periods to Maturity			
		1	2	3	4
Time 0	0	0.9970	0.9841	0.9619	0.9309
Time 1 S-1	1	0.9788	0.9425	0.8945	
Time 1 S-2	2	0.9931	0.9816	0.9636	
Time 1 S-3	3	0.9891	0.9692	0.9404	
Time 1 S-4	4	0.9866	0.9640	0.9326	
Time 2 S-1, S-1	5	0.9557	0.8952		
Time 2 S-1, S-2	6	0.9861	0.9551		
Time 2 S-1, S-3	7	0.9746	0.9332		
Time 2 S-1, S-4	8	0.9531	0.8986		
Time 2 S-2, S-1	9	0.9780	0.9461		
Time 2 S-2, S-2	10	0.9941	0.9829		
Time 2 S-2, S-3	11	0.9919	0.9781		
Time 2 S-2, S-4	12	0.9877	0.9691		
Time 2 S-3, S-1	13	0.9811	0.9580		
Time 2 S-3, S-2	14	0.9860	0.9628		
Time 2 S-3, S-3	15	0.9915	0.9681		
Time 2 S-3, S-4	16	0.9721	0.9372		
Time 2 S-4, S-1	17	0.9784	0.9525		
Time 2 S-4, S-2	18	0.9833	0.9572		
Time 2 S-4, S-3	19	0.9887	0.9625		
Time 2 S-4, S-4	20	0.9694	0.9319		

To get the value of the bond as of time $T = 2$, we multiply the \$3 received at time $T = 3$ and the \$103 received at time $T = 4$ by the appropriate zero-coupon bond price in each state. (We assume that we are doing the analysis one second after the \$3 coupon due at time $T = 2$ has been paid.) We get a different value for the bond in each of the 16 states (see Exhibit 21.6).

Now that we know all of the relevant time $T = 2$ values for the reference instrument, we can value European options on this instrument.

EXAMPLE: EUROPEAN CALL OPTION ON COUPON-BEARING BOND

We first value a European call option that gives us the right, but not the obligation, to buy the 3 percent coupon bond at a price of 101 at time $T = 2$, 1 second after the time $T = 2$ coupon of \$3 has been paid. We assume in this chapter that we exercise this call option rationally. In later chapters, we discuss retail options where irrational exercise is a possibility. We exercise our option and buy the bond whenever the time $T = 2$ price is more than 101. The cash flow in this case will be the price of the bond minus the 101 exercise price. If exercise of the option is not rational, the cash flow is zero. The cash flows attributable to the option are shown in Exhibit 21.7.

We then drop these option values into the cash flow table and multiply them by the respective probability-weighted discount factor. The cash flow table is shown in Exhibit 21.8, with the time $T = 3$ and $T = 4$ states partially hidden because they are all zero. The value of the European option to buy the bond at time $T = 2$ at a price of 101 is 0.5785.

What if the European option were a *put* option, not a *call* option? That gives us the right to sell the bond at a price of 101 at time $T = 2$. The cash flows from rational exercise of this put option are shown in Exhibit 21.9.

EXHIBIT 21.6 Value of 3 Percent Bond Maturing at $T = 4$

State Name	State Number	Bond Value
Time 2 S-1, S-1	5	95.07198353
Time 2 S-1, S-2	6	101.3313841
Time 2 S-1, S-3	7	99.04814186
Time 2 S-1, S-4	8	95.41694972
Time 2 S-2, S-1	9	100.3865212
Time 2 S-2, S-2	10	104.2211847
Time 2 S-2, S-3	11	103.7175263
Time 2 S-2, S-4	12	102.7851907
Time 2 S-3, S-1	13	101.6150074
Time 2 S-3, S-2	14	102.1228783
Time 2 S-3, S-3	15	102.6855935
Time 2 S-3, S-4	16	99.45226715
Time 2 S-4, S-1	17	101.0392465
Time 2 S-4, S-2	18	101.5442397
Time 2 S-4, S-3	19	102.1037666
Time 2 S-4, S-4	20	98.88886506

EXHIBIT 21.7 Value of 3 Percent Bond Maturing at $T = 4$

State Name	State Number	Bond Value	Exercise Price	Option Cash Flow
Time 2 S-1, S-1	5	95.07198353	101	0.0000
Time 2 S-1, S-2	6	101.3313841	101	0.3314
Time 2 S-1, S-3	7	99.04814186	101	0.0000
Time 2 S-1, S-4	8	95.41694972	101	0.0000
Time 2 S-2, S-1	9	100.3865212	101	0.0000
Time 2 S-2, S-2	10	104.2211847	101	3.2212
Time 2 S-2, S-3	11	103.7175263	101	2.7175
Time 2 S-2, S-4	12	102.7851907	101	1.7852
Time 2 S-3, S-1	13	101.6150074	101	0.6150
Time 2 S-3, S-2	14	102.1228783	101	1.1229
Time 2 S-3, S-3	15	102.6855935	101	1.6856
Time 2 S-3, S-4	16	99.45226715	101	0.0000
Time 2 S-4, S-1	17	101.0392465	101	0.0392
Time 2 S-4, S-2	18	101.5442397	101	0.5442
Time 2 S-4, S-3	19	102.1037666	101	1.1038
Time 2 S-4, S-4	20	98.88886506	101	0.0000

EXHIBIT 21.8 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	0.0000	0.0000
2		0.0000	0.3314	0.0000	0.0000
3		0.0000	0.0000	0.0000	0.0000
4		0.0000	0.0000	0.0000	0.0000
5			0.0000	0.0000	0.0000
6			3.2212	0.0000	0.0000
7			2.7175	0.0000	0.0000
8			1.7852	0.0000	0.0000
9			0.6150	0.0000	0.0000
10			1.1229	0.0000	0.0000
11			1.6856	0.0000	0.0000
12			0.0000	0.0000	0.0000
13			0.0392	0.0000	0.0000
14			0.5442	0.0000	0.0000
15			1.1038	0.0000	0.0000
16			0.0000	0.0000	0.0000
17				0.0000	0.0000
18				0.0000	0.0000
Row-Neutral Value =					0.5785

EXHIBIT 21.9 Value of 3 Percent Bond Maturing at $T = 4$

State Name	State Number	Bond Value	Exercise Price	Option Cash Flow
Time 2 S-1, S-1	5	95.07198353	101	5.9280
Time 2 S-1, S-2	6	101.3313841	101	0.0000
Time 2 S-1, S-3	7	99.04814186	101	1.9519
Time 2 S-1, S-4	8	95.41694972	101	5.5831
Time 2 S-2, S-1	9	100.3865212	101	0.6135
Time 2 S-2, S-2	10	104.2211847	101	0.0000
Time 2 S-2, S-3	11	103.7175263	101	0.0000
Time 2 S-2, S-4	12	102.7851907	101	0.0000
Time 2 S-3, S-1	13	101.6150074	101	0.0000
Time 2 S-3, S-2	14	102.1228783	101	0.0000
Time 2 S-3, S-3	15	102.6855935	101	0.0000
Time 2 S-3, S-4	16	99.45226715	101	1.5477
Time 2 S-4, S-1	17	101.0392465	101	0.0000
Time 2 S-4, S-2	18	101.5442397	101	0.0000
Time 2 S-4, S-3	19	102.1037666	101	0.0000
Time 2 S-4, S-4	20	98.88886506	101	2.1111

We again drop these put option cash flows into the cash flow table and multiply them by the respective probability-weighted discount factor. We show the cash flow table in Exhibit 21.10, along with the put option value of 1.2099.

We can check that these options values are consistent with no arbitrage in a very simple way. Consider this trade:

1. Buy the 3 percent bond maturing in $T = 4$
2. Sell the call option, allowing the option holder to buy the bond from us at 101
3. Buy the put option, allowing us to sell the bond to the put option holder at 101

The cost of this strategy is $104.7071 - 0.5785 + 1.2099 = 105.3385$. Another way to receive these cash flows would be to buy a bond that pays \$3 coupons at time $T = 1$ and $T = 2$ and principal of 101 at $T = 2$. We drop these cash flows into the cash flow table in Exhibit 21.11 and verify that the value/cost of this bond is the same number: 105.3385.

EXAMPLE: COUPON-BEARING BOND WITH EMBEDDED EUROPEAN CALL OPTION

We now turn to another common structure for retail borrowers, corporate borrowers, and sovereign borrowers. That is, a structure that allows the borrower to prepay the borrowing. Let's continue with the case of the 3 percent bond due at time $T = 4$ with principal of 100. What if we grant the issuer of the bond the onetime only

EXHIBIT 21.10 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	5.9280	0.0000	0.0000
2		0.0000	0.0000	0.0000	0.0000
3		0.0000	1.9519	0.0000	0.0000
4		0.0000	5.5831	0.0000	0.0000
5			0.6135	0.0000	0.0000
6			0.0000	0.0000	0.0000
7			0.0000	0.0000	0.0000
8			0.0000	0.0000	0.0000
9			0.0000	0.0000	0.0000
10			0.0000	0.0000	0.0000
11			0.0000	0.0000	0.0000
12			1.5477	0.0000	0.0000
13			0.0000	0.0000	0.0000
14			0.0000	0.0000	0.0000
15			0.0000	0.0000	0.0000
16			2.1111	0.0000	0.0000
17				0.0000	0.0000
18				0.0000	0.0000
Risk-Neutral Value =					1.2099

EXHIBIT 21.11 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	104.0000	0.0000	0.0000
2		3.0000	104.0000	0.0000	0.0000
3		3.0000	104.0000	0.0000	0.0000
4		3.0000	104.0000	0.0000	0.0000
5			104.0000	0.0000	0.0000
6			104.0000	0.0000	0.0000
7			104.0000	0.0000	0.0000
8			104.0000	0.0000	0.0000
9			104.0000	0.0000	0.0000
10			104.0000	0.0000	0.0000
11			104.0000	0.0000	0.0000
12			104.0000	0.0000	0.0000
13			104.0000	0.0000	0.0000
14			104.0000	0.0000	0.0000
15			104.0000	0.0000	0.0000
16			104.0000	0.0000	0.0000
17				0.0000	0.0000
Risk-Neutral Value =					105.3385

right to prepay at 100 immediately after the time $T = 2$ coupon of \$3 is paid. If the issuer does not exercise the option, the bond is repaid as scheduled at time $T = 4$. What is the true nature of this arrangement? This structure is labeled with the misleading name of “callable bond”; in fact, it involves a put option held by the bond issuer. The issuer sells a bond maturing at time $T = 2$ with a 3 percent coupon and principal of 100 and buys the put option that allows the issuer to compel the holder to buy another two-year, 3 percent bond maturing at time $T = 4$ at the issuer’s option at time $T = 2$ at a price of 100. We can model this structure as the sum of the two-year bond less the value of the put option or we can model the entire structure together. The value of the two-year bond is easily calculated either using our input zero-coupon bonds or via an input to the cash flow table. We show the latter and the value of 104.3544 in Exhibit 21.12.

The value of the put option on bond that pays a \$3 coupon at time $T = 3$ and $T = 4$ and principal of 100 at $T = 4$ is just the value of the cash flow gains that accrue to the owner of the put option (the issuer). The put option is exercised in 5 of 16 scenarios, so the bond continues to time $T = 4$. In the other 10 scenarios, the bond is called or repaid at time $T = 2$. Those cash flows are shown in Exhibit 21.13.

The put cash flows are put into the cash flow table as usual, giving us a value of 0.7245 in Exhibit 21.14.

Alternatively, we can model the full table of cash flows over all four periods. In the four nodes at time $T = 1$, cash flow is \$3. The put option, extending the bonds to year 4, is exercised in five scenarios at time $T = 2$. The time 2 cash flows for these

EXHIBIT 21.12 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	103.0000	0.0000	0.0000
2		3.0000	103.0000	0.0000	0.0000
3		3.0000	103.0000	0.0000	0.0000
4		3.0000	103.0000	0.0000	0.0000
5			103.0000	0.0000	0.0000
6			103.0000	0.0000	0.0000
7			103.0000	0.0000	0.0000
8			103.0000	0.0000	0.0000
9			103.0000	0.0000	0.0000
10			103.0000	0.0000	0.0000
11			103.0000	0.0000	0.0000
12			103.0000	0.0000	0.0000
13			103.0000	0.0000	0.0000
14			103.0000	0.0000	0.0000
15			103.0000	0.0000	0.0000
16			103.0000	0.0000	0.0000
17				0.0000	0.0000
Risk-Neutral Value =					104.3544

EXHIBIT 21.13 Value of 3 Percent Bond Maturing at $T = 4$

State Name	State Number	Bond Value	Exercise Price	Option Cash Flow
Time 2 S-1, S-1	5	95.07198353	100	4.9280
Time 2 S-1, S-2	6	101.3313841	100	0.0000
Time 2 S-1, S-3	7	99.04814186	100	0.9519
Time 2 S-1, S-4	8	95.41694972	100	4.5831
Time 2 S-2, S-1	9	100.3865212	100	0.0000
Time 2 S-2, S-2	10	104.2211847	100	0.0000
Time 2 S-2, S-3	11	103.7175263	100	0.0000
Time 2 S-2, S-4	12	102.7851907	100	0.0000
Time 2 S-3, S-1	13	101.6150074	100	0.0000
Time 2 S-3, S-2	14	102.1228783	100	0.0000
Time 2 S-3, S-3	15	102.6855935	100	0.0000
Time 2 S-3, S-4	16	99.45226715	100	0.5477
Time 2 S-4, S-1	17	101.0392465	100	0.0000
Time 2 S-4, S-2	18	101.5442397	100	0.0000
Time 2 S-4, S-3	19	102.1037666	100	0.0000
Time 2 S-4, S-4	20	98.88886506	100	1.1111

EXHIBIT 21.14 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	4.9280	0.0000	0.0000
2		0.0000	0.0000	0.0000	0.0000
3		0.0000	0.9519	0.0000	0.0000
4		0.0000	4.5831	0.0000	0.0000
5			0.0000	0.0000	0.0000
6			0.0000	0.0000	0.0000
7			0.0000	0.0000	0.0000
8			0.0000	0.0000	0.0000
9			0.0000	0.0000	0.0000
10			0.0000	0.0000	0.0000
11			0.0000	0.0000	0.0000
12			0.5477	0.0000	0.0000
13			0.0000	0.0000	0.0000
14			0.0000	0.0000	0.0000
15			0.0000	0.0000	0.0000
16			1.1111	0.0000	0.0000
17				0.0000	0.0000
Risk-Neutral Value =					0.7245

five scenarios are \$3. For the $4 \times 5 = 20$ nodes at time $T = 3$ associated with the decision to extend the bonds, cash flow is also \$3. For the associated 20 nodes at time $T = 4$, cash flow is 103. Going back to time $T = 2$, for the 10 scenarios in which the put is not exercised, the bonds are repaid at a cash flow of 103. All other nodes at times 3 and 4 are set to zero because the bonds have been retired. The value of this callable bond is 103.6299 as shown in Exhibit 21.15.

EXHIBIT 21.15 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		3.0000	103.0000	3.0000	103.0000
3		3.0000	3.0000	3.0000	103.0000
4		3.0000	3.0000	3.0000	103.0000
5			103.0000	0.0000	0.0000
6			103.0000	0.0000	0.0000
7			103.0000	0.0000	0.0000
8			103.0000	0.0000	0.0000
9			103.0000	3.0000	103.0000
10			103.0000	3.0000	103.0000
11			103.0000	3.0000	103.0000
12			3.0000	3.0000	103.0000
13			103.0000	3.0000	103.0000
14			103.0000	3.0000	103.0000
15			103.0000	3.0000	103.0000
16			3.0000	3.0000	103.0000
17				0.0000	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000

(Continued)

EXHIBIT 21.15 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
34				0.0000	0.0000
35				0.0000	0.0000
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				3.0000	103.0000
46				3.0000	103.0000
47				3.0000	103.0000
48				3.0000	103.0000
49				0.0000	0.0000
50				0.0000	0.0000
51				0.0000	0.0000
52				0.0000	0.0000
53				0.0000	0.0000
54				0.0000	0.0000
55				0.0000	0.0000
56				0.0000	0.0000
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				3.0000	103.0000
62				3.0000	103.0000
63				3.0000	103.0000
64				3.0000	103.0000
Risk-Neutral Value =					103.6299

We can verify there is no arbitrage by showing that the value of the callable bond (103.6299) equals the value of the noncall two-year bond (104.3544) less the cost of the put option (0.7245). One can easily use the solver function in common spreadsheet software to answer this question: What coupon level causes a new issue bond that matures in four years with a European call option (i.e., the option to prepay or the put option to extend) to be valued at 100 percent of par value?

EUROPEAN OPTIONS ON DEFAULTABLE BONDS

This review of European options on bonds emphasizes the power and ease of use of the HJM framework. Once the bushy tree is set up or the Monte Carlo simulations have been generated, valuation is quick and easy. We have discussed ABC Corporation as a risky issuer and taken advantage of the Jarrow insight that the risky zero-coupon bonds of ABC can be used to value coupon-bearing bonds in the same manner as if ABC had no credit risk. In analyzing the callable bonds in the prior example, the issuer holds the option so there is no counterparty risk to be analyzed by the issuer of the bonds.

In the case of discrete call or put options (i.e., point-in-time credit default swaps) on the bonds, there is indeed the risk that one's counterparty on the options may default. We have ignored that risk here. In Chapter 20, we reviewed the Jarrow formula for the impact of counterparty risk. In general, we strongly recommend a simulation approach to counterparty risk as the only realistic way to accurately capture the joint impact of macroeconomic factors on the default risk and credit spreads of both the bond issuer and those who hold call or put options on those securities.

Note that this chapter provides worked examples of how to value the Jarrow-Merton put option on the entire organization or portfolio as a measure of comprehensive risk. The only difference in methodology is that cash flows are generated from a collection of instruments valued on a collection of yield curves. Modern computers are expert at adding numbers together. The process is no more complex than the examples above. Repetition does not make the task more complex. It merely takes longer.

We close with a brief review of how simpler and less realistic term structure models can be used to obtain closed form solutions for fixed income options values. We present them in this chapter for exposition purposes.

Model Risk Alert

We do not recommend using such one-factor term structure models due to their inaccuracy. This applies, too, to the one-factor Vasicek model described in the following section.

HJM SPECIAL CASE: EUROPEAN OPTIONS IN THE ONE-FACTOR VASICEK MODEL

The lack of accuracy of one-factor models was reviewed in great detail in Chapters 3 through 10. (Indeed, we recommend that anyone tempted to use what follows to review Chapters 3 through 10 again.) We present this section as an aid to understanding for those who are mathematically inclined—recall, once more, that those who prefer mathematics based on unrealistic assumptions over accuracy have been labeled “F9 Model Monkeys” by the ever-entertaining Felix Salmon (2012) of Reuters.

In this section, we focus on the Vasicek model rather than the HJM three-factor model used in the rest of the chapter and the rest of the book. We do so to illustrate

that closed form solutions are possible for options valuation if one makes enough simplifying (and realism destroying) assumptions.

In Chapter 13 on legacy interest rate models, we introduced the Vasicek model in which short-term interest rates move randomly according to the stochastic process

$$dr = \alpha(\mu - r)dt + \sigma dz$$

Using the hedging argument from Chapter 13, riskless arbitrage requires that zero-coupon bonds for all maturities are related to the market price of risk λ by the following equation:

$$\frac{P_r \alpha(\mu - r) + \frac{1}{2} \sigma^2 P_{rr} + P_t - rP}{\sigma P_r} = -\lambda$$

Jamshidian (1989) derived the value of a European call option on a zero-coupon bond. We assume that the call option is to be valued at current time t . It is an option to purchase a risk-free, zero-coupon bond with maturity at time T_2 that is exercisable at a purchase price K at time T_1 .

The value of the call option as of time t given the observable short rate r is

$$V(r, t, T_1, T_2, K) = P(r, t, T_2)N(h) - P(r, t, T_1)KN(h - \sigma_P)$$

where N is the standard cumulative normal distribution. We use the following definitions:

$$h = \frac{1}{\sigma_P} \ln \left[\frac{P(r, t, T_2)}{KP(r, t, T_1)} \right] + \frac{\sigma_P}{2}$$

$$\sigma_P = \nu F_1$$

and

$$\nu = \sqrt{\frac{\sigma^2}{2\alpha} [1 - e^{-2\alpha(T_1-t)}]}$$

$$F_1(T_1 - t) = \frac{1}{\alpha} \left(1 - e^{-\alpha(T_1-t)} \right)$$

The latter expression is the standard deviation of the price of the T_2 maturity zero-coupon bond's price.

Readers familiar with the Black-Scholes options model will notice a very strong similarity to the pricing formula in that model. The similarities occur because the zero-coupon bond's price is lognormally distributed since its yield is linear in the short rate of interest and the yield is therefore normally distributed. There are very important differences in this formulation compared to the Black-Scholes model that are worth noting:

- The volatility of the bond's price declines over time and reaches zero at maturity, which is fully captured by the Jamshidian formulation. It is not captured by the Black-Scholes model if applied to bond options, since the bond price volatility is

assumed to be constant in the Black-Scholes model. As we noted earlier, many traders in caps, floors, and swaptions erroneously live with the Black model's assumption of constant bond price volatility, which is a partial explanation for the volatility "smile" observed when using the Black model for caps, floors, and swaptions.

- The Jamshidian formulation looks through bond price fluctuations to the economic source of the fluctuations, random interest rates, and provides an internally consistent methodology for bond option valuation. Using the Black-Scholes model for bond options pricing relies on the dangerous inconsistency that bond options are being valued using a model that assumes interest rates are constant.
- The Jamshidian formula, in combination with the Vasicek or Extended Vasicek term structure model, provides a theoretically consistent approach that allows cross-hedging of a one-year option on a three-year, zero-coupon bond with the appropriate hedging amount of two-year bonds.

Model Risk Alert

That being said, the one-factor model is too simple, but relative to Black-Scholes, it is an improvement of considerable magnitude. The Black-Scholes model does not provide an explanation of the relationship between price changes on bonds with different maturities.

OPTIONS ON COUPON-BEARING BONDS

Jamshidian notes that the price of any security with a positive amount of contractual cash flows a_i at n points of time in the future is a monotonically decreasing function of the short rate of interest r . Therefore, the zero-coupon bond formulas for European options can be applied directly to the problem of European options on coupon-bearing bonds. For each cash flow date, the strike price K_i is set to equal the zero-coupon bond price as of the exercise date T such that the present value of the remaining payments on the security as of time T exactly equal the true strike price on the entire security K . The interest rate r^* is the short-term interest rate at which this is true:

$$\sum_{i=1}^n a_i P(r^*, T, T_i) = K$$

and

$$K_i = P(r^*, T, T_i)$$

Jamshidian shows that the value of a European option on the entire security is the weighted sum of options on each cash flow:

$$\text{Option}(r, t, T, K) = \sum_{i=1}^n a_i \text{Option}(r, t, T, T_i, K_i)$$

THE JARROW-MERTON PUT OPTION

The worked examples using the three-factor HJM framework and the single-factor Jamshidian formula for bond options in the legacy Vasicek model both serve an important purpose: to show that the Jarrow-Merton put option calculation for measurement of total risk is not complicated or mysterious. It is a calculation that modern enterprise risk management software should be able to accomplish with ease. As we shall see in later chapters, the Jarrow-Merton put option as a total risk measure is far superior to the varied uses of “value at risk” that have been called into question in the wake of financial institutions’ failures around the world in the 2006–2011 credit crisis and in more recent times. Before demonstrating that concretely, we continue to work through other security types to illustrate the power of the HJM approach.

Forward and Futures Contracts

In this chapter, we continue the process of valuing the most common instruments on the balance sheet of major financial institutions around the world. Valuation enables us to efficiently perform a multiperiod simulation of cash flows and future values for these instruments, both of which are critical to the ultimate objective—measuring interest rate risk, market risk, liquidity risk, and credit risk in a consistent framework via the Jarrow-Merton put option introduced in Chapter 1.

Our focus in this chapter is on forward contracts and futures contracts for interest rate-related instruments. Careful attention to this issue is important both from a correct hedging analysis and from an accounting perspective. This chapter lays the groundwork for the correlation requirements of those standards, both in terms of valuation and cash flow. In Chapter 20, we warned of the dangers of collusion and market manipulation in the market for credit default swaps. (In Chapter 17, on credit spread modeling, we mentioned the legal actions alleging similar manipulation in the LIBOR market. On June 27, 2012, agencies of the U.S. government and the UK's Financial Services Authority levied more than \$450 million in penalties for LIBOR manipulation against Barclays PLC.) We assume perfectly competitive markets in the following analysis, but we warn readers, as we have throughout this volume, to make the appropriate adjustments given the nature of the pricing index and the nature of the counterparty. It is a mistake to assume that (1) everyone you trade with is Mother Teresa and that (2) you are the smartest one in the room. Lehman veteran Larry McDonald is often quoted as saying “If you are playing poker and you can't find the sucker at the table, it's probably you.” Beware.

Financial forward and futures contracts are critical everyday tools of risk management and investment management. Their popularity is probably best summarized by the interest rate futures contracts traded on the Chicago Mercantile Exchange. Open interest in these contracts has long exceeded trillions of dollars, even for the most long-dated contracts. Eurodollar futures on June 28, 2012, traded at maturities as long as nine years forward. Exhibit 22.1 from the CME Group summarizes the open interest in interest rate futures as of May 2012.

This chapter illustrates the important distinctions between forward and futures contracts of different types, using the tools of Chapter 21 that we developed in the context of valuing options on zero-coupon and coupon-bearing bonds. Note that we don't delve into the detailed contract specifications of specific futures contracts because the information is overwhelmingly voluminous. Instead, we concentrate on analytical methods that can be applied to any forward or futures contract with appropriate adjustments for the details of the instrument. Daily cash flows and

EXHIBIT 22.1 CME Group Exchange Open Interest Report—Monthly

	Open Interest MAY 2012	Open Interest MAY 2011	%CHG	Open Interest APR2012	%CHG
Interest Rate Futures					
30-YR BOND	741,599	740,953	0.1%	576,675	28.6%
10-YR NOTE	2,012,970	1,973,136	2.0%	1,847,462	9.0%
5-YR NOTE	1,300,391	1,555,969	−16.4%	1,332,831	−2.4%
2-YR NOTE	991,488	1,046,815	−5.3%	950,995	4.3%
FED FUND	553,694	751,787	−26.3%	532,994	3.9%
10-YR SWAP	8,633	13,823	−37.5%	7,973	8.3%
5-YR SWAP	10,210	28,377	−64.0%	8,880	15.0%
MINI					
EURODOLLAR	0	1	−100.0%	0	—
EUROYEN	263	825	−68.1%	516	−49.0%
EURODOLLARS	8,658,588	10,202,720	−15.1%	8,811,377	−17.7%
EMINI ED	7	11	−36.4%	7	0.0%
7-YR SWAP	36	631	−94.3%	39	−1.7%
3-MONTH OIS	0	0	—	0	—
3-YR NOTE	0	0	—	0	—
30-YR SWAP	122	640	−80.9%	140	−12.9%
1-MONTH					
EURODOLLAR	18,632	10,318	80.6%	18,050	3.2%
ULTRA T-BOND	386,551	378,832	2.0%	357,091	8.2%
OTR 10-YR NOTE	0	410	−100.0%	2	−100.0%
OTR 5-YR NOTE	0	47	−100.0%	0	—
OTR 2-YR NOTE	0	44	−100.0%	0	—
SOVEREIGN					
YIELD FUTURE	0	0	0.0%	0	—
EURO BONDS	294	0	—	532	−44.7%

Source: www.cmegroup.com/wrappedpages/web_monthly_report/Web_OI_Report_CMEG.pdf.

margin requirements are critical considerations with interest rate futures. We use the three-factor Heath, Jarrow, and Morton (HJM) bushy tree from Chapter 9 again in this chapter for ease of exposition. In practice, we would use daily time steps and a Monte Carlo implementation of the HJM bushy tree. We start first with various forward contracts and then cover futures contracts.

FORWARD CONTRACTS ON ZERO-COUPON BONDS

The most basic forward contract is a forward contract on the price of a zero-coupon bond. In this section and the rest of the chapter, we need to make an important distinction between the *value* of an existing contract on the books of an investment manager, insurance firm, or bank and the *price* of the forward contract quoted in the market. Before making this distinction clear, a real-world example of a forward contract on zero-coupon bonds is useful. In this case, the best example is provided by

the 90-day bank bill futures contract traded on the Sydney Futures Exchange. For purposes of this section, we ignore daily mark-to-market requirements and analyze this contract as a forward, not a futures, contract. (We perform the true futures analysis later in the chapter.) Our example has the following criteria:

Contract: 90-day Bank Bills, Sydney Futures Exchange

Quotation basis: 100 – yield in percent, quoted to two decimal places

Assumed maturity of underlying bill: 90 days

Valuation formula:

$$\text{Bank bill contract price} = \frac{365(1000000)}{365 + \frac{90 * \text{Yield}}{100}}$$

The bank bill contract assumes the underlying instrument pays simple interest on an actual/365-day basis and has a maturity of 90 days. Essentially, the price of the contract is the price of a 90-day zero-coupon bond with principal amount in Australian dollars of A\$1 million at maturity. The calculation is the same as those we used in Chapter 4. One of the key features of the Sydney Bank Bill contract, as opposed to the Eurodollar contracts we will discuss, is that the value of a 0.01 percent change in yield depends on the level of interest rates rather than being a constant dollar amount like the Eurodollar contract:

Yield (%)	A\$ Value of 0.01% Yield Increase
9.00%	A\$23.60
10.00%	A\$23.49
11.00%	A\$23.37
12.00%	A\$23.26

Note that the quotation method (100 – Yield) is irrelevant to valuation—only the underlying cash flow described by the valuation formula matters for valuation purposes.

For notational convenience, we do the valuation analysis assuming that the principal amount on the underlying instrument is A\$1 instead of A\$1 million. We know from Chapter 4 that the forward price of this bond in a no-arbitrage world with perfect competition should be the ratio of zero-coupon bond prices:

$$\text{Forward bond price}(r, t, T_1, T_2) = \frac{P(r, t, T_2)}{P(r, t, T_1)}$$

where T_2 is the maturity date of the underlying instrument (in years), T_1 is the maturity date of the forward contract, and t is the current time. This result doesn't depend on any particular term structure model since it can be derived using arbitrage arguments alone. If a financial institution has bought the forward contract at an exercise price of K , then at the maturity of the forward (time $T = T_1$), the gain will be

$P(r, T_1, T_2) - K$. Let's assume that the forward contract we are analyzing is a three-year forward contract on the one-year zero-coupon bond maturing at time $T = 4$. We can use the bushy tree of Chapter 9 and the related one-year zero-coupon bond prices as of time $T = 3$ to analyze cash flows to the purchaser of the forward contract for any given exercise price K . For an exercise price $K = 0.97$, the cash flows at time $T = 3$ on the 64 nodes are shown in Exhibit 22.2.

EXHIBIT 22.2 Gain or Loss on Three-Year Forward Contract on Zero-Coupon Bond Maturing in Year 2

State Name	State Number	Zero Coupon	Forward Price	Forward Cash Flow
Time 3 S-1, S-1, S-1	21	0.942412737	0.97	-0.0276
Time 3 S-1, S-1, S-2	22	0.947268253	0.97	-0.0227
Time 3 S-1, S-1, S-3	23	0.950832139	0.97	-0.0192
Time 3 S-1, S-1, S-4	24	0.925604805	0.97	-0.0444
Time 3 S-1, S-2, S-1	25	0.969893771	0.97	-0.0001
Time 3 S-1, S-2, S-2	26	0.974741292	0.97	0.0047
Time 3 S-1, S-2, S-3	27	0.98011229	0.97	0.0101
Time 3 S-1, S-2, S-4	28	0.960957928	0.97	-0.0090
Time 3 S-1, S-3, S-1	29	0.950380623	0.97	-0.0196
Time 3 S-1, S-3, S-2	30	0.980599699	0.97	0.0106
Time 3 S-1, S-3, S-3	31	0.969182195	0.97	-0.0008
Time 3 S-1, S-3, S-4	32	0.947852422	0.97	-0.0221
Time 3 S-1, S-4, S-1	33	0.948539838	0.97	-0.0215
Time 3 S-1, S-4, S-2	34	0.953426923	0.97	-0.0166
Time 3 S-1, S-4, S-3	35	0.957013979	0.97	-0.0130
Time 3 S-1, S-4, S-4	36	0.93162263	0.97	-0.0384
Time 3 S-2, S-1, S-1	37	0.960137528	0.97	-0.0099
Time 3 S-2, S-1, S-2	38	0.990666843	0.97	0.0207
Time 3 S-2, S-1, S-3	39	0.979132124	0.97	0.0091
Time 3 S-2, S-1, S-4	40	0.957583373	0.97	-0.0124
Time 3 S-2, S-2, S-1	41	0.978439239	0.97	0.0084
Time 3 S-2, S-2, S-2	42	0.994499441	0.97	0.0245
Time 3 S-2, S-2, S-3	43	0.992391065	0.97	0.0224
Time 3 S-2, S-2, S-4	44	0.988138659	0.97	0.0181
Time 3 S-2, S-3, S-1	45	0.974439351	0.97	0.0044
Time 3 S-2, S-3, S-2	46	1.003747972	0.97	0.0337
Time 3 S-2, S-3, S-3	47	0.997413498	0.97	0.0274
Time 3 S-2, S-3, S-4	48	0.978781276	0.97	0.0088
Time 3 S-2, S-4, S-1	49	0.982545874	0.97	0.0125
Time 3 S-2, S-4, S-2	50	0.987456629	0.97	0.0175
Time 3 S-2, S-4, S-3	51	0.992897691	0.97	0.0229
Time 3 S-2, S-4, S-4	52	0.973493464	0.97	0.0035
Time 3 S-3, S-1, S-1	53	0.977704061	0.97	0.0077
Time 3 S-3, S-1, S-2	54	0.982590617	0.97	0.0126
Time 3 S-3, S-1, S-3	55	0.988004867	0.97	0.0180
Time 3 S-3, S-1, S-4	56	0.96869626	0.97	-0.0013

EXHIBIT 22.2 (Continued)

State Name	State Number	Zero Coupon	Forward Price	Forward Cash Flow
Time 3 S-3, S-2, S-1	57	0.9777704061	0.97	0.0077
Time 3 S-3, S-2, S-2	58	0.982590617	0.97	0.0126
Time 3 S-3, S-2, S-3	59	0.988004867	0.97	0.0180
Time 3 S-3, S-2, S-4	60	0.96869626	0.97	−0.0013
Time 3 S-3, S-3, S-1	61	0.964921197	0.97	−0.0051
Time 3 S-3, S-3, S-2	62	0.993943537	0.97	0.0239
Time 3 S-3, S-3, S-3	63	0.987670937	0.97	0.0177
Time 3 S-3, S-3, S-4	64	0.969220711	0.97	−0.0008
Time 3 S-3, S-4, S-1	65	0.956873373	0.97	−0.0131
Time 3 S-3, S-4, S-2	66	0.987298898	0.97	0.0173
Time 3 S-3, S-4, S-3	67	0.975803393	0.97	0.0058
Time 3 S-3, S-4, S-4	68	0.954327901	0.97	−0.0157
Time 3 S-4, S-1, S-1	69	0.966171633	0.97	−0.0038
Time 3 S-4, S-1, S-2	70	0.996892813	0.97	0.0269
Time 3 S-4, S-1, S-3	71	0.985285603	0.97	0.0153
Time 3 S-4, S-1, S-4	72	0.963601426	0.97	−0.0064
Time 3 S-4, S-2, S-1	73	0.974819161	0.97	0.0048
Time 3 S-4, S-2, S-2	74	0.979691298	0.97	0.0097
Time 3 S-4, S-2, S-3	75	0.985089571	0.97	0.0151
Time 3 S-4, S-2, S-4	76	0.965837939	0.97	−0.0042
Time 3 S-4, S-3, S-1	77	0.974819161	0.97	0.0048
Time 3 S-4, S-3, S-2	78	0.979691298	0.97	0.0097
Time 3 S-4, S-3, S-3	79	0.985089571	0.97	0.0151
Time 3 S-4, S-3, S-4	80	0.965837939	0.97	−0.0042
Time 3 S-4, S-4, S-1	81	0.956882114	0.97	−0.0131
Time 3 S-4, S-4, S-2	82	0.977817245	0.97	0.0078
Time 3 S-4, S-4, S-3	83	0.973471088	0.97	0.0035
Time 3 S-4, S-4, S-4	84	0.952172548	0.97	−0.0178

If we drop these cash flows into the cash flow table and multiply by the proper probability-weighted discount factors, we find in Exhibit 22.3 that the value of this forward position is -0.0022 .

We know, however, that we can replicate the cash flows of this forward contract by making these transactions:

- Step 1: Buy zero-coupon bond maturing at time $T = 4$.
- Step 2: Agree to issue a zero-coupon bond with a face value of 0.97 maturing at time $T = 3$. The proceeds of this loan for 0.97, at the current yield curve, are 0.9330.
- Step 3: At time $T = 3$, sell zero-coupon bond maturing at time $T = 4$ and pay off the loan of 0.97 as it matures. There will be a gain or a loss from doing so that is identical to the payoff on the forward contract, so the net cash flow of this alternative strategy to the forward contract should also be -0.0022 .

EXHIBIT 22.3 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	−0.0276	0.0000
2		0.0000	0.0000	−0.0227	0.0000
3		0.0000	0.0000	−0.0192	0.0000
4		0.0000	0.0000	−0.0444	0.0000
5			0.0000	−0.0001	0.0000
6			0.0000	0.0047	0.0000
7			0.0000	0.0101	0.0000
8			0.0000	−0.0090	0.0000
9			0.0000	−0.0196	0.0000
10			0.0000	0.0106	0.0000
11			0.0000	−0.0008	0.0000
12			0.0000	−0.0221	0.0000
13			0.0000	−0.0215	0.0000
14			0.0000	−0.0166	0.0000
15			0.0000	−0.0130	0.0000
16			0.0000	−0.0384	0.0000
17				−0.0099	0.0000
18				0.0207	0.0000
19				0.0091	0.0000
20				−0.0124	0.0000
21				0.0084	0.0000
22				0.0245	0.0000
23				0.0224	0.0000
24				0.0181	0.0000
25				0.0044	0.0000
26				0.0337	0.0000
27				0.0274	0.0000
28				0.0088	0.0000
29				0.0125	0.0000
30				0.0175	0.0000
31				0.0229	0.0000
32				0.0035	0.0000
33				0.0077	0.0000
34				0.0126	0.0000
35				0.0180	0.0000
36				−0.0013	0.0000
37				0.0077	0.0000
38				0.0126	0.0000
39				0.0180	0.0000
40				−0.0013	0.0000
41				−0.0051	0.0000
42				0.0239	0.0000
43				0.0177	0.0000

EXHIBIT 22.3 (Continued)

Row Number	Current Time				
	0	1	2	3	4
44				−0.0008	0.0000
45				−0.0131	0.0000
46				0.0173	0.0000
47				0.0058	0.0000
48				−0.0157	0.0000
49				−0.0038	0.0000
50				0.0269	0.0000
51				0.0153	0.0000
52				−0.0064	0.0000
53				0.0048	0.0000
54				0.0097	0.0000
55				0.0151	0.0000
56				−0.0042	0.0000
57				0.0048	0.0000
58				0.0097	0.0000
59				0.0151	0.0000
60				−0.0042	0.0000
61				−0.0131	0.0000
62				0.0078	0.0000
63				0.0035	0.0000
64				−0.0178	0.0000
Risk-Neutral Value =					−0.0022

We can confirm that, indeed,

$$\begin{aligned}
 \text{Forward}(0, 3, 4, .97) &= P(0, 4) - P(0, 3)K \\
 &= 0.9308550992 - 0.9618922376(0.97) \\
 &= -0.0022.
 \end{aligned}$$

This result stems from no-arbitrage arguments and holds for any term structure model. Leaving aside the existing forward contract purchased for $K = 0.97$, what should the exercise price K of a new forward contract be today? Since no initial cash is necessary to buy a forward at market, the value of the forward should be zero so that

$$0 = P(0, 4) - P(0, 3)K$$

or

$$K = P(0, 4)/P(0, 3)$$

We can confirm this result by using the HJM cash flow table and solving for the level of K that makes the value of the existing futures contract zero.

Turning back to the Sydney forward contract, the observable market price (in zero-coupon bond terms) equals the forward bond price. In yield terms, as quoted on the Sydney Futures Exchange, the quoted yield will be

$$Y^* = \frac{36500}{90} \left[\frac{1000000}{\text{Forward zero-coupon bond price}} - 1 \right]$$

Later in this chapter, we deal with the implications of daily margin requirements and the possibility that there could be a default (1) either by the Exchange itself or (2) on the instrument underlying the forward contract in keeping with our interest in an integrated treatment of interest rate risk, market risk, liquidity risk, and credit risk. The analysis in this section implicitly assumes that both the Exchange and the underlying instrument are default free.

Forward Rate Agreements

The design of the Sydney Futures Exchange Bank Bill Contract is one of a relatively small number of contracts where the underlying instrument is a zero-coupon bond. It is most common for the underlying instrument to be expressed in terms of an interest rate. Perhaps the most common forward contract with this structure is the forward rate agreement (FRA), where the counterparties agree on a strike rate K at maturity T_1 based on an underlying rate for an instrument maturing at T_2 . The cash flow at time T_2 from the FRA is proportional to the FRA rate at maturity and the strike rate K :

$$\text{Cash flow} = B(\text{FRA rate}[T_1] - K)$$

where the coefficient B reflects the notional principal amount, the maturity of the underlying instrument, and the interest quotation method used for the FRA contract (i.e., actual/360 days, actual/365 days, etc.). In this section, we analyze the case where this interest differential is paid at the maturity of the underlying instrument, T_2 . Conventional FRAs normally pay in an economically equivalent way, paying the present value of this interest rate differential at the maturity of the FRA contract, T_1 .

We value the contract as if cash flowed at time T_1 by taking the simple present value of T_2 cash flow from the perspective of time T_1 (the standard FRA convention):

$$\begin{aligned} \text{Present value} &= P(\tau = T_2 - T_1)B[\text{FRA rate} - K] \\ &= P(\tau)B \left[\frac{1}{P(\tau)} - 1 - K \right] \\ &= B[1 - P(\tau)(1 + K)] \\ &= B - BP(\tau)(1 + K) \end{aligned}$$

Let's assume that B is A\$25 per 1 percent rate change (A\$0.25 per basis point) and that the FRA exercise price is $K = 4$ percent. You receive any excess of the actual spot rate over 4 percent and you pay any shortfall of the actual spot rate under 4 percent. The cash flows that are paid at time $T = 4$ on a one-year underlying FRA are set out in Exhibit 22.4 for each of the relevant 64 nodes.

EXHIBIT 22.4 FRA at 4 Percent on One-Year Instrument

State Name	State Number	Spot Rate at Time 3	Exercise Spot Rate	Cash Flow Time $T = 4$
Time 3 S-1, S-1, S-1	21	6.11%	4.00%	52.77
Time 3 S-1, S-1, S-2	22	5.57%	4.00%	39.17
Time 3 S-1, S-1, S-3	23	5.17%	4.00%	29.28
Time 3 S-1, S-1, S-4	24	8.04%	4.00%	100.94
Time 3 S-1, S-2, S-1	25	3.10%	4.00%	-22.40
Time 3 S-1, S-2, S-2	26	2.59%	4.00%	-35.22
Time 3 S-1, S-2, S-3	27	2.03%	4.00%	-49.27
Time 3 S-1, S-2, S-4	28	4.06%	4.00%	1.57
Time 3 S-1, S-3, S-1	29	5.22%	4.00%	30.53
Time 3 S-1, S-3, S-2	30	1.98%	4.00%	-50.54
Time 3 S-1, S-3, S-3	31	3.18%	4.00%	-20.51
Time 3 S-1, S-3, S-4	32	5.50%	4.00%	37.54
Time 3 S-1, S-4, S-1	33	5.43%	4.00%	35.63
Time 3 S-1, S-4, S-2	34	4.88%	4.00%	22.12
Time 3 S-1, S-4, S-3	35	4.49%	4.00%	12.29
Time 3 S-1, S-4, S-4	36	7.34%	4.00%	83.49
Time 3 S-2, S-1, S-1	37	4.15%	4.00%	3.79
Time 3 S-2, S-1, S-2	38	0.94%	4.00%	-76.45
Time 3 S-2, S-1, S-3	39	2.13%	4.00%	-46.72
Time 3 S-2, S-1, S-4	40	4.43%	4.00%	10.74
Time 3 S-2, S-2, S-1	41	2.20%	4.00%	-44.91
Time 3 S-2, S-2, S-2	42	0.55%	4.00%	-86.17
Time 3 S-2, S-2, S-3	43	0.77%	4.00%	-80.83
Time 3 S-2, S-2, S-4	44	1.20%	4.00%	-69.99
Time 3 S-2, S-3, S-1	45	2.62%	4.00%	-34.42
Time 3 S-2, S-3, S-2	46	-0.37%	4.00%	-109.33
Time 3 S-2, S-3, S-3	47	0.26%	4.00%	-93.52
Time 3 S-2, S-3, S-4	48	2.17%	4.00%	-45.80
Time 3 S-2, S-4, S-1	49	1.78%	4.00%	-55.59
Time 3 S-2, S-4, S-2	50	1.27%	4.00%	-68.24
Time 3 S-2, S-4, S-3	51	0.72%	4.00%	-82.12
Time 3 S-2, S-4, S-4	52	2.72%	4.00%	-31.93
Time 3 S-3, S-1, S-1	53	2.28%	4.00%	-42.99
Time 3 S-3, S-1, S-2	54	1.77%	4.00%	-55.71
Time 3 S-3, S-1, S-3	55	1.21%	4.00%	-69.65
Time 3 S-3, S-1, S-4	56	3.23%	4.00%	-19.21
Time 3 S-3, S-2, S-1	57	2.28%	4.00%	-42.99
Time 3 S-3, S-2, S-2	58	1.77%	4.00%	-55.71
Time 3 S-3, S-2, S-3	59	1.21%	4.00%	-69.65
Time 3 S-3, S-2, S-4	60	3.23%	4.00%	-19.21
Time 3 S-3, S-3, S-1	61	3.64%	4.00%	-9.11
Time 3 S-3, S-3, S-2	62	0.61%	4.00%	-84.77

(Continued)

EXHIBIT 22.4 (Continued)

State Name	State Number	Spot Rate at Time 3	Exercise Spot Rate	Cash Flow Time $T = 4$
Time 3 S-3, S-3, S-3	63	1.25%	4.00%	−68.79
Time 3 S-3, S-3, S-4	64	3.18%	4.00%	−20.61
Time 3 S-3, S-4, S-1	65	4.51%	4.00%	12.68
Time 3 S-3, S-4, S-2	66	1.29%	4.00%	−67.84
Time 3 S-3, S-4, S-3	67	2.48%	4.00%	−38.01
Time 3 S-3, S-4, S-4	68	4.79%	4.00%	19.64
Time 3 S-4, S-1, S-1	69	3.50%	4.00%	−12.47
Time 3 S-4, S-1, S-2	70	0.31%	4.00%	−92.21
Time 3 S-4, S-1, S-3	71	1.49%	4.00%	−62.66
Time 3 S-4, S-1, S-4	72	3.78%	4.00%	−5.57
Time 3 S-4, S-2, S-1	73	2.58%	4.00%	−35.42
Time 3 S-4, S-2, S-2	74	2.07%	4.00%	−48.18
Time 3 S-4, S-2, S-3	75	1.51%	4.00%	−62.16
Time 3 S-4, S-2, S-4	76	3.54%	4.00%	−11.57
Time 3 S-4, S-3, S-1	77	2.58%	4.00%	−35.42
Time 3 S-4, S-3, S-2	78	2.07%	4.00%	−48.18
Time 3 S-4, S-3, S-3	79	1.51%	4.00%	−62.16
Time 3 S-4, S-3, S-4	80	3.54%	4.00%	−11.57
Time 3 S-4, S-4, S-1	81	4.51%	4.00%	12.65
Time 3 S-4, S-4, S-2	82	2.27%	4.00%	−43.29
Time 3 S-4, S-4, S-3	83	2.73%	4.00%	−31.87
Time 3 S-4, S-4, S-4	84	5.02%	4.00%	25.57

If we drop these into the cash flow table, we get the value −15.4927 in Exhibit 22.5.

Note that we can replicate the net present value that will be received under this transaction at time $T = 3$. We know $B = A\$2,500$ (since rates are expressed as a decimal) and the FRA exercise price is 0.04. We replicate this time $T = 3$ net present value payoff with the following steps:

Step 1: Buy A\$2,500 in principal amount of the zero-coupon bond maturing at time $T = 3$. The cost of this investment is $P(0,3)2,500$.

Step 2: Borrow in the form of a zero-coupon bond with principal A\$2,500 $(1 + 0.04)$ to mature at time $T = 4$. The proceeds of this borrowing are $P(0,4)[2,500](1.04)$.

Step 3: At time $T = 3$ either liquidate the position or invest the proceeds of the bond purchased in step 1 (A\$2,500) in the one-year spot rate maturing at time $T = 4$.

No arbitrage requires that the net cost of this strategy be the same as −15.4927 and we can confirm that is the case:

$$P(0, 3)2, 500 - P(0, 4)[2, 500](1.04) = -15.4927$$

EXHIBIT 22.5 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	0.0000	52.7655
2		0.0000	0.0000	0.0000	39.1679
3		0.0000	0.0000	0.0000	29.2759
4		0.0000	0.0000	0.0000	100.9367
5			0.0000	0.0000	−22.3981
6			0.0000	0.0000	−35.2169
7			0.0000	0.0000	−49.2719
8			0.0000	0.0000	1.5707
9			0.0000	0.0000	30.5250
10			0.0000	0.0000	−50.5397
11			0.0000	0.0000	−20.5056
12			0.0000	0.0000	37.5414
13			0.0000	0.0000	35.6299
14			0.0000	0.0000	22.1202
15			0.0000	0.0000	12.2920
16			0.0000	0.0000	83.4900
17				0.0000	3.7937
18				0.0000	−76.4473
19				0.0000	−46.7184
20				0.0000	10.7387
21				0.0000	−44.9103
22				0.0000	−86.1725
23				0.0000	−80.8318
24				0.0000	−69.9907
25				0.0000	−34.4222
26				0.0000	−109.3349
27				0.0000	−93.5170
28				0.0000	−45.8032
29				0.0000	−55.5895
30				0.0000	−68.2432
31				0.0000	−82.1172
32				0.0000	−31.9293
33				0.0000	−42.9890
34				0.0000	−55.7054
35				0.0000	−69.6481
36				0.0000	−19.2117
37				0.0000	−42.9890
38				0.0000	−55.7054
39				0.0000	−69.6481
40				0.0000	−19.2117
41				0.0000	−9.1149

(Continued)

EXHIBIT 22.5 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
42				0.0000	−84.7666
43				0.0000	−68.7926
44				0.0000	−20.6082
45				0.0000	12.6759
46				0.0000	−67.8388
47				0.0000	−38.0085
48				0.0000	19.6447
49				0.0000	−12.4680
50				0.0000	−92.2078
51				0.0000	−62.6646
52				0.0000	−5.5663
53				0.0000	−35.4218
54				0.0000	−48.1758
55				0.0000	−62.1597
56				0.0000	−11.5740
57				0.0000	−35.4218
58				0.0000	−48.1758
59				0.0000	−62.1597
60				0.0000	−11.5740
61				0.0000	12.6520
62				0.0000	−43.2850
63				0.0000	−31.8703
64				0.0000	25.5745
Risk-Neutral Value =					−15.4927

We can also solve for the implied “price” of new FRA transactions K , recognizing that an FRA transaction at time zero is costless and therefore should have a value of zero to avoid arbitrage possibilities:

$$K = P(0, 3)/P(0, 4) - 1$$

Again, this formula holds for any term structure model because the replication of the FRA is independent of the term structure model selected. What if the rate differential between the FRA rate and K is paid on an undiscounted basis at time T_1 ? We answer that question in the next section.

EURODOLLAR FUTURES-TYPE FORWARD CONTRACTS

What if the market convention regarding FRA payment was that the forward rate agreement pays its cash flow:

$$\text{Cash flow} = B(\text{FRA rate}[T_1] - K)$$

at the T_1 maturity of the FRA contract instead its present value at time T_1 or payment in cash at the T_2 maturity of the underlying instrument? An FRA that pays in this manner is essentially equivalent to a Eurodollar Futures contract with no mark-to-market margin (i.e., a Eurodollar futures-type forward). Such a Eurodollar futures-type forward, if modeled after the Chicago Mercantile Exchange Eurodollar futures contract, would pay \$25 for every one basis point differential between the settlement yield on the futures-type forward and the strike price. Unlike the Sydney Futures Exchange Bank Bill future discussed in the previous section, the dollar value of a basis point change in quoted yield is the same for all levels of interest rates under this type of contract. The cash flow at time T_1 under this kind of contract is

$$\text{Cash flow} = B \left(\frac{1}{P(\tau)} - 1 - K \right) = \frac{B}{P(\tau)} - (1 + K)B$$

The cash flows are the same as before but they arrive at time $T = 3$ instead of time $T = 4$. We can calculate the value of this kind of structure by dropping the cash flows into a different column, the column for time $T = 3$, in the cash flow table. The value is -15.4146 , as shown in Exhibit 22.6.

EXHIBIT 22.6 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	52.7655	0.0000
2		0.0000	0.0000	39.1679	0.0000
3		0.0000	0.0000	29.2759	0.0000
4		0.0000	0.0000	100.9367	0.0000
5			0.0000	-22.3981	0.0000
6			0.0000	-35.2169	0.0000
7			0.0000	-49.2719	0.0000
8			0.0000	1.5707	0.0000
9			0.0000	30.5250	0.0000
10			0.0000	-50.5397	0.0000
11			0.0000	-20.5056	0.0000
12			0.0000	37.5414	0.0000
13			0.0000	35.6299	0.0000
14			0.0000	22.1202	0.0000
15			0.0000	12.2920	0.0000
16			0.0000	83.4900	0.0000
17				3.7937	0.0000
18				-76.4473	0.0000
19				-46.7184	0.0000
20				10.7387	0.0000
21				-44.9103	0.0000
22				-86.1725	0.0000

(Continued)

EXHIBIT 22.6 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
23				−80.8318	0.0000
24				−69.9907	0.0000
25				−34.4222	0.0000
26				−109.3349	0.0000
27				−93.5170	0.0000
28				−45.8032	0.0000
29				−55.5895	0.0000
30				−68.2432	0.0000
31				−82.1172	0.0000
32				−31.9293	0.0000
33				−42.9890	0.0000
34				−55.7054	0.0000
35				−69.6481	0.0000
36				−19.2117	0.0000
37				−42.9890	0.0000
38				−55.7054	0.0000
39				−69.6481	0.0000
40				−19.2117	0.0000
41				−9.1149	0.0000
42				−84.7666	0.0000
43				−68.7926	0.0000
44				−20.6082	0.0000
45				12.6759	0.0000
46				−67.8388	0.0000
47				−38.0085	0.0000
48				19.6447	0.0000
49				−12.4680	0.0000
50				−92.2078	0.0000
51				−62.6646	0.0000
52				−5.5663	0.0000
53				−35.4218	0.0000
54				−48.1758	0.0000
55				−62.1597	0.0000
56				−11.5740	0.0000
57				−35.4218	0.0000
58				−48.1758	0.0000
59				−62.1597	0.0000
60				−11.5740	0.0000
61				12.6520	0.0000
62				−43.2850	0.0000
63				−31.8703	0.0000
64				25.5745	0.0000
Risk-Neutral Value =					−15.4146

The valuation formulas for this instrument are more complex than the simple no-arbitrage results given previously.

FUTURES ON ZERO-COUPON BONDS: THE SYDNEY FUTURES EXCHANGE BANK BILL CONTRACT

From this point on in the chapter, we deal with true futures contracts. A true futures contract has daily mark-to-market requirements, with cash flow on an existing contract occurring continuously through its life. Ending cash flow is essentially zero, since the full amount of the change in observable futures prices over the life of the contract will have already been paid in cash through the continuous mark-to-market requirements. We will ignore initial margin payments and the return of these initial payments since they represent a simple zero-coupon bond-type adjustment to the formulas that follow.

Like the previous three sections, there is an important distinction between the market value of an existing futures position V and the observable market price of a futures contract Q . The first three sections showed the distinct differences between V and Q for three different types of forward contracts. In this section and the remainder of the chapter, we note the following:

- The market value V of an existing futures position will always be zero since (1) the market price of a new futures contract must be zero to avoid arbitrage and (2) a previously taken futures position is continuously made equivalent to a new position by continuous margin payments. The observable market price Q of the future contract and the initial strike rate K on the futures transaction, put in place at time t_0 , are related by the fact that cash payments of at least $M = K - Q$, the cumulative net loss (gain) on the futures contract, must have been made to (withdrawn from) the margin account since initiation of the transaction at time t_0 . The balance of the margin account is unknown since parties with a positive margin balance may use the cash proceeds for any purpose.
- This important comment notwithstanding, a trader is always thinking about their gains or losses in terms of the original rate at which the futures contract was purchased K and the current market price Q . It is critical to note that this differential will have already impacted the financial institution's balance sheet (i.e., the cash account) and the trader's focus on the differential $K - Q$ reflects only a partial view of the impact of the futures contract on the institution. Most sophisticated institutions are taking both views of the impact of a futures contract, the trader's view and the true daily cash flow pattern via the margin accounts.

These comments require a careful vocabulary to bridge the gap between the jargon of market participants and financial reality. To a market participant, the value of a futures contract is $Q - K$, even though the margin account balance would be $\text{Max}(0, K - Q)$ rather than always being $Q - K$. These same market participants are rarely aware of the margin account balance, which is normally managed by an operational group. To a financial economist, the value is zero. Fortunately, the interest rate risk of a futures position to both groups is the change in Q that results

from a change in interest rates, as captured by movements in the single stochastic factor r (if the Vasicek term structure model applies) or by movements in the three factors driving our HJM model from Chapter 9. Taking care to remember this distinction between viewpoints, we ignore the market value of an existing futures position V and concentrate on determining the observable market price Q .¹

In modeling interest rate futures contracts, the most important difference versus the forward rate contracts discussed earlier in the chapter is easy to summarize. Let's assume at time $T = 0$ one takes a position in a futures contract maturing at time $T = 3$ on the zero-coupon bond maturing at time $T = 4$. Obviously, at time $T = 3$, the price of the maturing futures contract will be forced by arbitrage to exactly equal the time $T = 3$ price of the zero-coupon bond maturing at time $T = 4$. This boundary condition is no different than it is for forward contracts. The difference comes in what happens at times $T = 1$ and $T = 2$. In the case of a forward contract maturing at time $T = 3$, nothing happens at times 1 and 2, which is why the cash flows above were zero for those time periods. We write the price of a futures contract as of time t with maturity at $T = 3$ on the zero-coupon bond maturing at $T = 4$ as $F(t, 3, 4)$. The time zero price of the futures contract is $F(0, 3, 4)$. At time 1, there are four nodes on the bushy tree and the futures price will be taken on one of four values, $F(1, 3, 4, s_i)$ for $i = 1, 2, 3$, and 4. At each of these four nodes, there will be cash flow equal to $F(1, 3, 4, s_i) - F(0, 3, 4)$. This cash flow either represents the withdrawal of a gain from a margin account or the contribution of a loss back into the margin account. After this cash flow has occurred at time $T = 1$, of course, an existing futures position and a new futures position are identical. Both have zero dollars in the margin account.

At time $T = 2$, there are 16 nodes and there will be 16 gains or losses on the futures position relative to the four futures values on the nodes at $T = 1$.

The analytical derivation of the futures price evolution is nicely done by Chen (1992) in the Vasicek framework. As the number of risk factors driving interest rates in an HJM framework gets larger and larger for greater realism, futures prices and their evolution in general can be most usefully derived by numerical methods. See Jarrow and Turnbull (1996) for an excellent summary of the general principles that govern the evolution of futures prices. The need for numerical methods is particularly true in light of the "cheapest to deliver option" embedded in many bond futures contracts. We turn to that issue now.

FUTURES ON COUPON-BEARING BONDS: DEALING WITH THE CHEAPEST TO DELIVER OPTION

In this section, we consider the contract specifications of the Singapore Exchange's Japanese government bond (JGB) contract. Let's assume the coupon rate on the futures contract is 2 percent:

Notional principal: ¥50,000,000

Maturity: 10 years

Coupon rate: 2 percent

In actual operation, there will be a number of JGB issues that are deliverable under the Singapore contract. The seller of the futures contract has a delivery option

which allows the seller of the futures contract to deliver the cheapest bond if the seller of the futures contract chooses to hold the contract to maturity. If there are N deliverable bonds, the exchange will calculate a delivery factor F_j , which specifies the ratio of the principal amount on the i th deliverable bond to the notional principal of ¥50,000,000 that is necessary to satisfy delivery requirements.

One of the bonds at any given time is the cheapest to deliver. After multiplying the yen coupon amounts and yen principal amount on this bond times the delivery factor, we will have the j maturities and cash flow amounts underlying the bond future. The market price of the bond future (as long as only one bond is deliverable) is the sum of zero-coupon bond futures prices (assuming one yen of principal on each zero-coupon bond future) times these cash flow amounts C_i :

$$Q_{\text{Bond}}(t, T_0) = \sum_{i=1}^j Q(r, t, T_i) C_i$$

T_0 is the maturity date of the bond futures contract. The sensitivity analysis of bond futures, assuming only one bond is deliverable, is the sum of the sensitivities for the zero-coupon bond futures portfolio, which replicates the cash flows on the bond.

What about the value of the delivery option? Let us assume that there are two deliverable bonds and the one-factor Vasicek model is a realistic description of rate movements. (Model Risk Alert: It is not realistic.) Using the insights of Jamshidian's (1989) approach to coupon-bearing bond options valuation, we know that there will be a level of the short rate r , which we label s^* , such that the two bonds are equally attractive to deliver as of time T_f . Above s^* , bond one is cheaper to deliver. Below s^* , bond two is cheaper to deliver. We can solve the partial differential equation for futures valuation subject to the boundary conditions that the futures price converges to bond one's deliverable value above s^* and to bond two's deliverable value below s^* . We also impose the condition that the first and second derivatives of the futures price are smooth at a short rate level of s^* . The result is a special weighted average of the futures prices that would have prevailed if each bond was the only deliverable bond under the futures contract. There are other options embedded in common U.S. and Japanese bond contracts that can be analyzed in a related way. The process is much more complex in a realistic HJM framework with the 5 to 10 risk factors that we showed were necessary for realism in Chapter 3.

EURODOLLAR AND EUROYEN FUTURES CONTRACTS

Money market futures have evolved to a fairly standard structure around the world. The Singapore Exchange futures contracts in U.S. dollars, yen, and Singapore dollars are all on three-month instruments with a constant currency amount paid at maturity per basis point change in the price of the contract. The per-basis-point value of rate changes under each contract is as follows:

U.S. dollars: US\$25.00

Yen (Tokyo and London): ¥2,500

Singapore dollars: SG\$25.00

In each case, the futures price must be equal to the underlying three-month interest rates, multiplied by the notional principal amount and the appropriate day count factor to convert the rate to an annual basis. At maturity, the futures price (actually, the futures rate times notional principal) Q must meet this boundary condition:

$$Q(r, T_1, T_1, T_2) = B \left[\frac{1}{P(r, T_1, T_2)} - 1 \right]$$

B represents the notional principal and annualization factor, T_1 is the maturity of the futures contract in years, and T_2 is the maturity of the underlying three-month instrument so T_2 is roughly 0.25 greater than T_1 .

The market price of the futures contract in the Vasicek model context (ignoring the annualization factor and notional principal embedded in B) is one over the forward bond price, multiplied by an adjustment factor, minus one. This formula reduces to the forward rate if interest rate volatility σ is zero. In general, the formulas for these money market futures contracts (when they can be obtained) are highly dependent on the term structure model selected.

DEFAULTABLE FORWARD AND FUTURES CONTRACTS

Jarrow and Turnbull (1995) show that, in special circumstances, the adjustment made in the price of a defaultable instrument is a simple ratio times the value of the same instrument if it were not defaultable. Many market participants simulate future values and cash flows of financial forwards and futures on the implicit assumption that default will not occur. The realization that default of the futures exchange itself is a possibility, however, should not come as a surprise to anyone who is familiar with the high correlation in the default probabilities between major financial institutions that we discussed in the introduction to this book in the 2006–2011 credit crisis. How should this risk be analyzed? We need to use the simulation technique outlined in Chapters 19 and 20 to handle this accurately. Accuracy requires us to use simulation methods outlined in those chapters, although explicit formulaic solutions can be obtained for highly simplified assumptions.

A careful Monte Carlo simulation like that specified in Chapters 19 and 20 is essential to recognizing the slim, but real possibility of the default of the futures exchanges in times of crisis, like the collapse of futures specialist MF Global on October 30, 2011. We return to this theme in Chapters 36 through 41 in detail.

NOTE

1. For an excellent introduction to this topic, see Chen (1992) and Jarrow and Turnbull (1996), especially Chapters 6 and 13 of the latter.

European Options on Forward and Futures Contracts

This chapter continues our analysis of standard instruments found on the balance sheet of large financial institutions as we continue with our objective of establishing a unified measure of interest rate risk, market risk, liquidity risk, and credit risk. In order to value the Jarrow-Merton put option proposed in Chapter 1 as the best measure of total risk, we need a methodology for valuation and for simulation of cash flows and values at many future dates. As in Chapter 22, we need this capability for European options on forward and futures contracts for three reasons:

- To understand correct hedging amounts from a shareholder value-added perspective, as discussed in the early chapters of this book
- To meet the requirements of constantly evolving accounting standards for hedging
- To meet the requirements of regulatory capital calculations like Basel II, Basel III, and Solvency II

This chapter combines the valuation formulas for futures and forwards from Chapter 22 and the options approach from Chapter 21. As in Chapters 21 and 22, we use the three-factor Heath, Jarrow, and Morton (HJM) framework to value options on forwards and futures. We remind the reader that Chapter 3 illustrated the need for an interest rate model with 5 to 10 driving risk factors. We use the three-factor HJM model from Chapter 9 for expositional purposes. In the first edition of this book, we showed in detail how the one-factor Vasicek model can be used to derive closed form valuation formulas for this important class of instruments. While this approach provides many useful insights, we don't believe that a one-factor model is sufficiently accurate to emphasize in this volume. Similarly, the \$450 million settlement announced by Barclays PLC on June 27, 2012, for manipulation of the LIBOR market calls into question the pervasive use of the LIBOR market model for money market instruments. A more general approach like the HJM approach provides the most reliable and accurate framework for analysis.

VALUING OPTIONS ON FORWARDS AND FUTURES: NOTATIONS AND USEFUL FORMULAS

Valuing options on forward and futures contracts can be done in using one of two methods. First, we can employ the HJM framework using a bushy tree or Monte

Interest Rates Futures & Options Products - CME Group		
Block Trades		
STIR (CME)		
Eurodollar	FUT OPT	
Euribor	FUT	
Mid-Curve Options	FUT OPT	
1-month Eurodollar	FUT OPT	
Euroyen TBOR	FUT OPT	
3-Month OIS Futures	FUT OPT	
Eurodollar Calendar Spread Options	OPT	
STIR (CBOT)		
30-Day Federal Funds	FUT OPT	
Cleared OTC		
Interest Rate Swaps		
Sovereign Yield Spreads (CME)		
US-UK 10-Year Sovys Futures	FUT	
US-DE 10-Year Sovys Futures	FUT	
US-FR 10-Year Sovys Futures	FUT	
US-IT 10-Year Sovys Futures	FUT	
US-Hd 10-Year Sovys Futures	FUT	
UK-DE 10-Year Sovys Futures	FUT	
UK-FR 10-Year Sovys Futures	FUT	
UK-IT 10-Year Sovys Futures	FUT	
UK-Hd 10-Year Sovys Futures	FUT	
DE-FR 10-Year Sovys Futures	FUT	
DE-IT 10-Year Sovys Futures	FUT	
DE-Hd 10-Year Sovys Futures	FUT	
U.S. Treasury Futures and Options (CBOT)		
Ultra T-Bond	FUT OPT	
U.S. Treasury Bond	FUT OPT	
10-Year U.S. Treasury Note	FUT OPT	
5-Year U.S. Treasury Note	FUT OPT	
3-Year U.S. Treasury Note	FUT OPT	
2-Year U.S. Treasury Note	FUT OPT	
OTR 2-Year U.S. Treasury Note	FUT	
OTR 5-Year U.S. Treasury Note	FUT	
OTR 10-Year U.S. Treasury Note	FUT	
U.S. Treasury Futures and Options (CME)		
13-week T-bill	FUT	
Intercommodity Spreads (CBOT)		
Treasury and Swap Spreads	FUT	
Swap Futures and Options (CBOT)		
5-Year Interest Rate Swap	FUT OPT	
7-Year Interest Rate Swap	FUT OPT	
10-Year Interest Rate Swap	FUT OPT	
30-Year Interest Rate Swap	FUT OPT	
Interest Rate Indexes (CME)		
Barclays U.S. Aggregate Bond Index Futures	FUT	
Eurozone HCP Futures	FUT	

EXHIBIT 23.1 Interest Rate Pictures and Options Products—CME Group

Carlo simulation, properly parameterized, to obtain accurate values and hedges using a term structure model with a realistic number of risk factors (5 to 10). We use the Chapter 9 HJM three-factor example for exposition purposes in this chapter. A second approach, which the authors admire but now judge insufficiently accurate, is to take a simple term structure model with one or two factors and derive analytical valuation and hedging formulas. For example, we can take advantage of the Jamshidian general solution for security pricing under the Vasicek model from Chapter 21 to obtain analytical solutions.

Market participants have historically preferred analytical solutions over numerical solutions because of a desire for total control over the calculation on the part of the trader/analyst. In the modern era, with 64-bit operating systems and powerful enterprise risk systems, this relentless urge by Felix Salmon's "F9 Model Monkeys" is not in the best interest of most financial services firms. For that reason we emphasize the HJM numerical solutions in this chapter.

Exhibit 23.1 summarizes the current list of interest rate futures and options contracts on interest rates available at www.cmegroup.com.

Because of the wide array of interest rate options currently available, this chapter will emphasize options on stylized types of instruments. Space does not permit a detailed analysis of each contract, and, because contract specifications are rapidly evolving, such a product-specific focus would soon be rendered obsolete. We start with options on forward contracts on zero-coupon bonds, as in Chapter 22.

EUROPEAN OPTIONS ON FORWARD CONTRACTS ON ZERO-COUPON BONDS

In this section, we return to our example from Chapter 22 of a three-year forward contract on a zero-coupon bond maturing in year 4. The forward contract had an exercise price of 0.97. The cash flows on such a contract, using the HJM 3 factor model of Chapter 9, were shown in Exhibits 22.2 and 22.3 of Chapter 22. We reproduce the cash flow table for the forward contract as Exhibit 23.2, along with the forward contract's value of -0.0022 .

EXHIBIT 23.2 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	−0.0276	0.0000
2		0.0000	0.0000	−0.0227	0.0000
3		0.0000	0.0000	−0.0192	0.0000
4		0.0000	0.0000	−0.0444	0.0000
5			0.0000	−0.0001	0.0000
6			0.0000	0.0047	0.0000
7			0.0000	0.0101	0.0000
8			0.0000	−0.0090	0.0000
9			0.0000	−0.0196	0.0000
10			0.0000	0.0106	0.0000
11			0.0000	−0.0008	0.0000
12			0.0000	−0.0221	0.0000
13			0.0000	−0.0215	0.0000
14			0.0000	−0.0166	0.0000
15			0.0000	−0.0130	0.0000
16			0.0000	−0.0384	0.0000
17				−0.0099	0.0000
18				0.0207	0.0000
19				0.0091	0.0000
20				−0.0124	0.0000
21				0.0084	0.0000
22				0.0245	0.0000
23				0.0224	0.0000
24				0.0181	0.0000
25				0.0044	0.0000
26				0.0337	0.0000
27				0.0274	0.0000
28				0.0088	0.0000
29				0.0125	0.0000
30				0.0175	0.0000
31				0.0229	0.0000
32				0.0035	0.0000
33				0.0077	0.0000
34				0.0126	0.0000
35				0.0180	0.0000
36				−0.0013	0.0000
37				0.0077	0.0000
38				0.0126	0.0000
39				0.0180	0.0000
40				−0.0013	0.0000
41				−0.0051	0.0000

(Continued)

EXHIBIT 23.2 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
42				0.0239	0.0000
43				0.0177	0.0000
44				−0.0008	0.0000
45				−0.0131	0.0000
46				0.0173	0.0000
47				0.0058	0.0000
48				−0.0157	0.0000
49				−0.0038	0.0000
50				0.0269	0.0000
51				0.0153	0.0000
52				−0.0064	0.0000
53				0.0048	0.0000
54				0.0097	0.0000
55				0.0151	0.0000
56				−0.0042	0.0000
57				0.0048	0.0000
58				0.0097	0.0000
59				0.0151	0.0000
60				−0.0042	0.0000
61				−0.0131	0.0000
62				0.0078	0.0000
63				0.0035	0.0000
64				−0.0178	0.0000
Risk-Neutral Value =					−0.0022

Now consider the simplest possible call option on this forward contract. Let the option be a European call option with an exercise period of three years, exactly the maturity of the forward contract. Let the exercise price of the option on the forward contract be zero. In this case, since the forward contract will converge to the zero-coupon bond price minus 0.97 at maturity, the call option will be exercised only when there are gains at the maturity of the futures contract. To value the call option, we modify the cash flow table above so that it registers only values equal to Maximum (0, cash flow). When we make this modification, we see the cash flows for the three-year European call option on this forward contract in Exhibit 23.3.

We drop these cash flows into the cash flow table and multiply times the corresponding probability-weighted discount factors from Chapter 9. When we do this, we see that the value of the three-year call option on this forward contract is 0.0048 in Exhibit 23.4.

What if the option contract were a put option with the same three-year term? Then the cash flows that would be relevant would be Minimum (0, cash flow). If we

EXHIBIT 23.3 Gain or Loss on Three-Year Forward Contract on Zero-Coupon Bond Maturing in Year 4

State Name	State Number	Zero-Coupon Bond Value	Forward Price	Forward Cash Flow	Call Option Cash Flow
Time 3 S-1, S-1, S-1	21	0.942412737	0.97	-0.0276	0.0000
Time 3 S-1, S-1, S-2	22	0.947268253	0.97	-0.0227	0.0000
Time 3 S-1, S-1, S-3	23	0.950832139	0.97	-0.0192	0.0000
Time 3 S-1, S-1, S-4	24	0.925604805	0.97	-0.0444	0.0000
Time 3 S-1, S-2, S-1	25	0.969893771	0.97	-0.0001	0.0000
Time 3 S-1, S-2, S-2	26	0.974741292	0.97	0.0047	0.0047
Time 3 S-1, S-2, S-3	27	0.98011229	0.97	0.0101	0.0101
Time 3 S-1, S-2, S-4	28	0.960957928	0.97	-0.0090	0.0000
Time 3 S-1, S-3, S-1	29	0.950380623	0.97	-0.0196	0.0000
Time 3 S-1, S-3, S-2	30	0.980599699	0.97	0.0106	0.0106
Time 3 S-1, S-3, S-3	31	0.969182195	0.97	-0.0008	0.0000
Time 3 S-1, S-3, S-4	32	0.947852422	0.97	-0.0221	0.0000
Time 3 S-1, S-4, S-1	33	0.948539838	0.97	-0.0215	0.0000
Time 3 S-1, S-4, S-2	34	0.953426923	0.97	-0.0166	0.0000
Time 3 S-1, S-4, S-3	35	0.957013979	0.97	-0.0130	0.0000
Time 3 S-1, S-4, S-4	36	0.93162263	0.97	-0.0384	0.0000
Time 3 S-2, S-1, S-1	37	0.960137528	0.97	-0.0099	0.0000
Time 3 S-2, S-1, S-2	38	0.990666843	0.97	0.0207	0.0207
Time 3 S-2, S-1, S-3	39	0.979132124	0.97	0.0091	0.0091
Time 3 S-2, S-1, S-4	40	0.957583373	0.97	-0.0124	0.0000
Time 3 S-2, S-2, S-1	41	0.978439239	0.97	0.0084	0.0084
Time 3 S-2, S-2, S-2	42	0.994499441	0.97	0.0245	0.0245
Time 3 S-2, S-2, S-3	43	0.992391065	0.97	0.0224	0.0224
Time 3 S-2, S-2, S-4	44	0.988138659	0.97	0.0181	0.0181
Time 3 S-2, S-3, S-1	45	0.974439351	0.97	0.0044	0.0044
Time 3 S-2, S-3, S-2	46	1.003747972	0.97	0.0337	0.0337
Time 3 S-2, S-3, S-3	47	0.997413498	0.97	0.0274	0.0274
Time 3 S-2, S-3, S-4	48	0.978781276	0.97	0.0088	0.0088
Time 3 S-2, S-4, S-1	49	0.982545874	0.97	0.0125	0.0125
Time 3 S-2, S-4, S-2	50	0.987456629	0.97	0.0175	0.0175
Time 3 S-2, S-4, S-3	51	0.992897691	0.97	0.0229	0.0229
Time 3 S-2, S-4, S-4	52	0.973493464	0.97	0.0035	0.0035
Time 3 S-3, S-1, S-1	53	0.977704061	0.97	0.0077	0.0077
Time 3 S-3, S-1, S-2	54	0.982590617	0.97	0.0126	0.0126
Time 3 S-3, S-1, S-3	55	0.988004867	0.97	0.0180	0.0180
Time 3 S-3, S-1, S-4	56	0.96869626	0.97	-0.0013	0.0000
Time 3 S-3, S-2, S-1	57	0.977704061	0.97	0.0077	0.0077
Time 3 S-3, S-2, S-2	58	0.982590617	0.97	0.0126	0.0126
Time 3 S-3, S-2, S-3	59	0.988004867	0.97	0.0180	0.0180
Time 3 S-3, S-2, S-4	60	0.96869626	0.97	-0.0013	0.0000
Time 3 S-3, S-3, S-1	61	0.964921197	0.97	-0.0051	0.0000

(Continued)

EXHIBIT 23.3 (Continued)

State Name	State Number	Zero-Coupon Bond Value	Forward Price	Forward Cash Flow	Call Option Cash Flow
Time 3 S-3, S-3, S-2	62	0.993943537	0.97	0.0239	0.0239
Time 3 S-3, S-3, S-3	63	0.987670937	0.97	0.0177	0.0177
Time 3 S-3, S-3, S-4	64	0.969220711	0.97	-0.0008	0.0000
Time 3 S-3, S-4, S-1	65	0.956873373	0.97	-0.0131	0.0000
Time 3 S-3, S-4, S-2	66	0.987298898	0.97	0.0173	0.0173
Time 3 S-3, S-4, S-3	67	0.975803393	0.97	0.0058	0.0058
Time 3 S-3, S-4, S-4	68	0.954327901	0.97	-0.0157	0.0000
Time 3 S-4, S-1, S-1	69	0.966171633	0.97	-0.0038	0.0000
Time 3 S-4, S-1, S-2	70	0.996892813	0.97	0.0269	0.0269
Time 3 S-4, S-1, S-3	71	0.985285603	0.97	0.0153	0.0153
Time 3 S-4, S-1, S-4	72	0.963601426	0.97	-0.0064	0.0000
Time 3 S-4, S-2, S-1	73	0.974819161	0.97	0.0048	0.0048
Time 3 S-4, S-2, S-2	74	0.979691298	0.97	0.0097	0.0097
Time 3 S-4, S-2, S-3	75	0.985089571	0.97	0.0151	0.0151
Time 3 S-4, S-2, S-4	76	0.965837939	0.97	-0.0042	0.0000
Time 3 S-4, S-3, S-1	77	0.974819161	0.97	0.0048	0.0048
Time 3 S-4, S-3, S-2	78	0.979691298	0.97	0.0097	0.0097
Time 3 S-4, S-3, S-3	79	0.985089571	0.97	0.0151	0.0151
Time 3 S-4, S-3, S-4	80	0.965837939	0.97	-0.0042	0.0000
Time 3 S-4, S-4, S-1	81	0.956882114	0.97	-0.0131	0.0000
Time 3 S-4, S-4, S-2	82	0.977817245	0.97	0.0078	0.0078
Time 3 S-4, S-4, S-3	83	0.973471088	0.97	0.0035	0.0035
Time 3 S-4, S-4, S-4	84	0.952172548	0.97	-0.0178	0.0000

EXHIBIT 23.4 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	0.0000	0.0000
2		0.0000	0.0000	0.0000	0.0000
3		0.0000	0.0000	0.0000	0.0000
4		0.0000	0.0000	0.0000	0.0000
5			0.0000	0.0000	0.0000
6			0.0000	0.0047	0.0000
7			0.0000	0.0101	0.0000
8			0.0000	0.0000	0.0000
9			0.0000	0.0000	0.0000
10			0.0000	0.0106	0.0000
11			0.0000	0.0000	0.0000
12			0.0000	0.0000	0.0000
13			0.0000	0.0000	0.0000

EXHIBIT 23.4 (Continued)

Row Number	Current Time				
	0	1	2	3	4
14			0.0000	0.0000	0.0000
15			0.0000	0.0000	0.0000
16			0.0000	0.0000	0.0000
17				0.0000	0.0000
18				0.0207	0.0000
19				0.0091	0.0000
20				0.0000	0.0000
21				0.0084	0.0000
22				0.0245	0.0000
23				0.0224	0.0000
24				0.0181	0.0000
25				0.0044	0.0000
26				0.0337	0.0000
27				0.0274	0.0000
28				0.0088	0.0000
29				0.0125	0.0000
30				0.0175	0.0000
31				0.0229	0.0000
32				0.0035	0.0000
33				0.0077	0.0000
34				0.0126	0.0000
35				0.0180	0.0000
36				0.0000	0.0000
37				0.0077	0.0000
38				0.0126	0.0000
39				0.0180	0.0000
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0239	0.0000
43				0.0177	0.0000
44				0.0000	0.0000
45				0.0000	0.0000
46				0.0173	0.0000
47				0.0058	0.0000
48				0.0000	0.0000
49				0.0000	0.0000
50				0.0269	0.0000
51				0.0153	0.0000
52				0.0000	0.0000
53				0.0048	0.0000
54				0.0097	0.0000
55				0.0151	0.0000

(Continued)

EXHIBIT 23.4 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
56				0.0000	0.0000
57				0.0048	0.0000
58				0.0097	0.0000
59				0.0151	0.0000
60				0.0000	0.0000
61				0.0000	0.0000
62				0.0078	0.0000
63				0.0035	0.0000
64				0.0000	0.0000
Risk-Neutral Value =					0.0048

calculate these cash flows, drop them in the cash flow table, and multiply times the relevant probability-weighted discount factors, we see in Exhibit 23.5 that the value of this European put option is -0.0070 .

We now check to make sure that these valuations are consistent with no arbitrage. A forward contract with a three-year maturity and a strike price of 0.97 has a value of -0.0022 . We can replicate this forward contract by (1) buying a three-year call option on the forward contract at a strike price of zero and (2) selling a three-year put option on the forward contract at a strike price of zero. The cash flow from this replication strategy would be $-0.0048 + 0.0070 = 0.0022$, exactly the amount of cash a rational investor would need to receive before the investor would take on a forward contract with a current value of -0.0022 .

EUROPEAN OPTIONS ON FORWARD RATE AGREEMENTS

We now examine the valuation of a European call option on a forward rate agreement with the same terms and cash flows as in Exhibits 22.4 and 22.5 in Chapter 22. Recall that the FRA from Chapter 22 paid \$25 per basis point for every basis point over the exercise price of 4.00%. Time zero value of this FRA was -15.4927 , as shown in Exhibit 22.5 of Chapter 22. What would be the value of a European call option to buy this FRA contract with a strike price of zero and an exercise period of three years? The call option would earn the holder the FRA cash flows such that the cash flows were Maximum (0, FRA cash flow); that is, only the positive cash flows from the ownership of the FRA. These cash flows are shown in Exhibit 23.6.

The value of the call option on the FRA is the value of these cash flows. We drop them into the cash flow table and multiply times the probability-weighted discount factors to get a call option value of 9.0716 shown in Exhibit 23.7.

If instead we valued a put option on the FRA, the cash flows would be Minimum (0, FRA cash flow). The value of the put option is shown in Exhibit 23.8 as -24.5643 .

EXHIBIT 23.5 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	−0.0276	0.0000
2		0.0000	0.0000	−0.0227	0.0000
3		0.0000	0.0000	−0.0192	0.0000
4		0.0000	0.0000	−0.0444	0.0000
5			0.0000	−0.0001	0.0000
6			0.0000	0.0000	0.0000
7			0.0000	0.0000	0.0000
8			0.0000	−0.0090	0.0000
9			0.0000	−0.0196	0.0000
10			0.0000	0.0000	0.0000
11			0.0000	−0.0008	0.0000
12			0.0000	−0.0221	0.0000
13			0.0000	−0.0215	0.0000
14			0.0000	−0.0166	0.0000
15			0.0000	−0.0130	0.0000
16			0.0000	−0.0384	0.0000
17				−0.0099	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				−0.0124	0.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000
36				−0.0013	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				−0.0013	0.0000

(Continued)

EXHIBIT 23.5 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
41				−0.0051	0.0000
42				0.0000	0.0000
43				0.0000	0.0000
44				−0.0008	0.0000
45				−0.0131	0.0000
46				0.0000	0.0000
47				0.0000	0.0000
48				−0.0157	0.0000
49				−0.0038	0.0000
50				0.0000	0.0000
51				0.0000	0.0000
52				−0.0064	0.0000
53				0.0000	0.0000
54				0.0000	0.0000
55				0.0000	0.0000
56				−0.0042	0.0000
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				−0.0042	0.0000
61				−0.0131	0.0000
62				0.0000	0.0000
63				0.0000	0.0000
64				−0.0178	0.0000
Risk-Neutral Value =					−0.0070

We can again confirm the absence of arbitrage by noting that the total of the cash flows from a replication strategy (buy the call option and sell the put option, $-9.0716 + 24.5643$) equals the cash flow that an investor would receive for taking on an existing FRA position (15.4927) currently valued at a loss of 15.4927.

EUROPEAN OPTIONS ON A EURODOLLAR FUTURES-TYPE FORWARD CONTRACT

In Exhibit 22.6 in Chapter 22, we dealt with a forward contract that pays like the standard money market futures contracts; that is, at the expiration of the futures contract or forward contract instead of on the interest payment date of the underlying instrument. Options on this type of forward contract are exactly the same as the analysis in Exhibits 23.7 and 23.8 except for the timing of the cash flows. These valuations are left as an exercise for the interested reader.

EXHIBIT 23.6 FRA at 4 Percent on One-Year Instrument

State Name	State Number	Spot Rate at Time 3	Exercise Spot Rate	FRA Cash Flow Time $T = 4$	Call Option Cash Flow
Time 3 S-1, S-1, S-1	21	6.11%	4.00%	52.77	52.77
Time 3 S-1, S-1, S-2	22	5.57%	4.00%	39.17	39.17
Time 3 S-1, S-1, S-3	23	5.17%	4.00%	29.28	29.28
Time 3 S-1, S-1, S-4	24	8.04%	4.00%	100.94	100.94
Time 3 S-1, S-2, S-1	25	3.10%	4.00%	-22.40	0.00
Time 3 S-1, S-2, S-2	26	2.59%	4.00%	-35.22	0.00
Time 3 S-1, S-2, S-3	27	2.03%	4.00%	-49.27	0.00
Time 3 S-1, S-2, S-4	28	4.06%	4.00%	1.57	1.57
Time 3 S-1, S-3, S-1	29	5.22%	4.00%	30.53	30.53
Time 3 S-1, S-3, S-2	30	1.98%	4.00%	-50.54	0.00
Time 3 S-1, S-3, S-3	31	3.18%	4.00%	-20.51	0.00
Time 3 S-1, S-3, S-4	32	5.50%	4.00%	37.54	37.54
Time 3 S-1, S-4, S-1	33	5.43%	4.00%	35.63	35.63
Time 3 S-1, S-4, S-2	34	4.88%	4.00%	22.12	22.12
Time 3 S-1, S-4, S-3	35	4.49%	4.00%	12.29	12.29
Time 3 S-1, S-4, S-4	36	7.34%	4.00%	83.49	83.49
Time 3 S-2, S-1, S-1	37	4.15%	4.00%	3.79	3.79
Time 3 S-2, S-1, S-2	38	0.94%	4.00%	-76.45	0.00
Time 3 S-2, S-1, S-3	39	2.13%	4.00%	-46.72	0.00
Time 3 S-2, S-1, S-4	40	4.43%	4.00%	10.74	10.74
Time 3 S-2, S-2, S-1	41	2.20%	4.00%	-44.91	0.00
Time 3 S-2, S-2, S-2	42	0.55%	4.00%	-86.17	0.00
Time 3 S-2, S-2, S-3	43	0.77%	4.00%	-80.83	0.00
Time 3 S-2, S-2, S-4	44	1.20%	4.00%	-69.99	0.00
Time 3 S-2, S-3, S-1	45	2.62%	4.00%	-34.42	0.00
Time 3 S-2, S-3, S-2	46	-0.37%	4.00%	-109.33	0.00
Time 3 S-2, S-3, S-3	47	0.26%	4.00%	-93.52	0.00
Time 3 S-2, S-3, S-4	48	2.17%	4.00%	-45.80	0.00
Time 3 S-2, S-4, S-1	49	1.78%	4.00%	-55.59	0.00
Time 3 S-2, S-4, S-2	50	1.27%	4.00%	-68.24	0.00
Time 3 S-2, S-4, S-3	51	0.72%	4.00%	-82.12	0.00
Time 3 S-2, S-4, S-4	52	2.72%	4.00%	-31.93	0.00
Time 3 S-3, S-1, S-1	53	2.28%	4.00%	-42.99	0.00
Time 3 S-3, S-1, S-2	54	1.77%	4.00%	-55.71	0.00
Time 3 S-3, S-1, S-3	55	1.21%	4.00%	-69.65	0.00
Time 3 S-3, S-1, S-4	56	3.23%	4.00%	-19.21	0.00
Time 3 S-3, S-2, S-1	57	2.28%	4.00%	-42.99	0.00
Time 3 S-3, S-2, S-2	58	1.77%	4.00%	-55.71	0.00
Time 3 S-3, S-2, S-3	59	1.21%	4.00%	-69.65	0.00
Time 3 S-3, S-2, S-4	60	3.23%	4.00%	-19.21	0.00
Time 3 S-3, S-3, S-1	61	3.64%	4.00%	-9.11	0.00

(Continued)

EXHIBIT 23.6 (Continued)

State Name	State Number	Spot Rate at Time 3	Exercise Spot Rate	FRA Cash Flow Time $T = 4$	Call Option Cash Flow
Time 3 S-3, S-3, S-2	62	0.61%	4.00%	-84.77	0.00
Time 3 S-3, S-3, S-3	63	1.25%	4.00%	-68.79	0.00
Time 3 S-3, S-3, S-4	64	3.18%	4.00%	-20.61	0.00
Time 3 S-3, S-4, S-1	65	4.51%	4.00%	12.68	12.68
Time 3 S-3, S-4, S-2	66	1.29%	4.00%	-67.84	0.00
Time 3 S-3, S-4, S-3	67	2.48%	4.00%	-38.01	0.00
Time 3 S-3, S-4, S-4	68	4.79%	4.00%	19.64	19.64
Time 3 S-4, S-1, S-1	69	3.50%	4.00%	-12.47	0.00
Time 3 S-4, S-1, S-2	70	0.31%	4.00%	-92.21	0.00
Time 3 S-4, S-1, S-3	71	1.49%	4.00%	-62.66	0.00
Time 3 S-4, S-1, S-4	72	3.78%	4.00%	-5.57	0.00
Time 3 S-4, S-2, S-1	73	2.58%	4.00%	-35.42	0.00
Time 3 S-4, S-2, S-2	74	2.07%	4.00%	-48.18	0.00
Time 3 S-4, S-2, S-3	75	1.51%	4.00%	-62.16	0.00
Time 3 S-4, S-2, S-4	76	3.54%	4.00%	-11.57	0.00
Time 3 S-4, S-3, S-1	77	2.58%	4.00%	-35.42	0.00
Time 3 S-4, S-3, S-2	78	2.07%	4.00%	-48.18	0.00
Time 3 S-4, S-3, S-3	79	1.51%	4.00%	-62.16	0.00
Time 3 S-4, S-3, S-4	80	3.54%	4.00%	-11.57	0.00
Time 3 S-4, S-4, S-1	81	4.51%	4.00%	12.65	12.65
Time 3 S-4, S-4, S-2	82	2.27%	4.00%	-43.29	0.00
Time 3 S-4, S-4, S-3	83	2.73%	4.00%	-31.87	0.00
Time 3 S-4, S-4, S-4	84	5.02%	4.00%	25.57	25.57

EXHIBIT 23.7 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	0.0000	52.7655
2		0.0000	0.0000	0.0000	39.1679
3		0.0000	0.0000	0.0000	29.2759
4		0.0000	0.0000	0.0000	100.9367
5			0.0000	0.0000	0.0000
6			0.0000	0.0000	0.0000
7			0.0000	0.0000	0.0000
8			0.0000	0.0000	1.5707
9			0.0000	0.0000	30.5250
10			0.0000	0.0000	0.0000
11			0.0000	0.0000	0.0000
12			0.0000	0.0000	37.5414

EXHIBIT 23.7 (Continued)

Row Number	Current Time				
	0	1	2	3	4
13			0.0000	0.0000	35.6299
14			0.0000	0.0000	22.1202
15			0.0000	0.0000	12.2920
16			0.0000	0.0000	83.4900
17				0.0000	3.7937
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	10.7387
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				0.0000	12.6759
46				0.0000	0.0000
47				0.0000	0.0000
48				0.0000	19.6447
49				0.0000	0.0000
50				0.0000	0.0000
51				0.0000	0.0000
52				0.0000	0.0000
53				0.0000	0.0000
54				0.0000	0.0000

(Continued)

EXHIBIT 23.7 (Continued)

Row Number	Current Time				
	0	1	2	3	4
55				0.0000	0.0000
56				0.0000	0.0000
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				0.0000	12.6520
62				0.0000	0.0000
63				0.0000	0.0000
64				0.0000	25.5745
Risk-Neutral Value =					9.0716

EXHIBIT 23.8 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	0.0000	0.0000
2		0.0000	0.0000	0.0000	0.0000
3		0.0000	0.0000	0.0000	0.0000
4		0.0000	0.0000	0.0000	0.0000
5			0.0000	0.0000	−22.3981
6			0.0000	0.0000	−35.2169
7			0.0000	0.0000	−49.2719
8			0.0000	0.0000	0.0000
9			0.0000	0.0000	0.0000
10			0.0000	0.0000	−50.5397
11			0.0000	0.0000	−20.5056
12			0.0000	0.0000	0.0000
13			0.0000	0.0000	0.0000
14			0.0000	0.0000	0.0000
15			0.0000	0.0000	0.0000
16			0.0000	0.0000	0.0000
17				0.0000	0.0000
18				0.0000	−76.4473
19				0.0000	−46.7184
20				0.0000	0.0000
21				0.0000	−44.9103
22				0.0000	−86.1725
23				0.0000	−80.8318
24				0.0000	−69.9907

EXHIBIT 23.8 (Continued)

Row Number	Current Time				
	0	1	2	3	4
25				0.0000	−34.4222
26				0.0000	−109.3349
27				0.0000	−93.5170
28				0.0000	−45.8032
29				0.0000	−55.5895
30				0.0000	−68.2432
31				0.0000	−82.1172
32				0.0000	−31.9293
33				0.0000	−42.9890
34				0.0000	−55.7054
35				0.0000	−69.6481
36				0.0000	−19.2117
37				0.0000	−42.9890
38				0.0000	−55.7054
39				0.0000	−69.6481
40				0.0000	−19.2117
41				0.0000	−9.1149
42				0.0000	−84.7666
43				0.0000	−68.7926
44				0.0000	−20.6082
45				0.0000	0.0000
46				0.0000	−67.8388
47				0.0000	−38.0085
48				0.0000	0.0000
49				0.0000	−12.4680
50				0.0000	−92.2078
51				0.0000	−62.6646
52				0.0000	−5.5663
53				0.0000	−35.4218
54				0.0000	−48.1758
55				0.0000	−62.1597
56				0.0000	−11.5740
57				0.0000	−35.4218
58				0.0000	−48.1758
59				0.0000	−62.1597
60				0.0000	−11.5740
61				0.0000	0.0000
62				0.0000	−43.2850
63				0.0000	−31.8703
64				0.0000	0.0000
Risk-Neutral Value =					−24.5643

European Options on Futures on Coupon-Bearing Bonds

In Chapter 22, we argued that the valuation of futures contracts on coupon-bearing bonds is best accomplished using numerical methods such as the HJM three-factor bushy tree or related Monte Carlo simulation. There were three reasons for this conclusion: First, as we saw in Chapters 3, 5 to 10 risk factors drive interest rates and there is, in general, no closed form solution for futures in this case; second, daily mark-to-market cash flows compound the complexity of the futures contract beyond the complexity of the cash flows on the underlying instrument; third, the cheapest-to-deliver option embedded in most bond futures contracts requires numerical methods for an accurate valuation. We saw in Chapter 21 that the Jamshidian bond option formula makes use of the critical short rate level that triggers the “moneyness” of the bond option. A similar technique can be used in a one-factor term structure model to identify the transition points at which the bond that is cheapest to deliver changes from bond J to bond $J + 1$. In a term structure model where many factors drive interest rates, an option on a bond future that has an embedded cheapest-to-deliver option is a compound option. In general, this kind of compound option is best analyzed using the HJM framework rather than via closed form approximations.

European Options on Money Market Futures Contracts

For similar reasons, the HJM framework is the best approach to value and hedge positions that involve money market futures contracts as well. This challenging task has been further complicated by the LIBOR manipulation scandal in two dimensions. First, the manipulation of LIBOR that Barclays admitted on June 27, 2012, pollutes any historical data set that one would use for LIBOR-based interest rate modeling. Second, until there is a dramatic reform of setting short-term interest rates for pricing caps, floors, and swaps, the potential for future manipulation remains. To ignore this in modeling options on futures would be naïve at best and suicidal at worst. The best modeling strategy will be dependent on further revelations that emerge from ongoing investigations into LIBOR manipulation.

DEFAULTABLE OPTIONS ON FORWARD AND FUTURES CONTRACTS

As we noted in Chapter 22, Jarrow and Turnbull (1995) show that, in special circumstances, the adjustment made to the price of a defaultable instrument is a simple ratio times the value of the same instrument if it was not defaultable. Can we analyze the possibility of a default on the options in this chapter in the same way?

As we noted in Chapter 22, the answer is no. Since the macro factors impact options prices, forward prices, futures prices, and the default probabilities of the counterparties dealing in the forwards and futures, the neat separation of the credit-adjusted value of futures and forwards that Jarrow and Turnbull found for simpler instruments will not generally apply. This doesn't mean that this approach can't be used as a first approximation, but it will significantly understate risk. As JPMorgan Chase found out in 1998 in a well-known incident with S.K. Securities of Korea, the default probability of the counterparty can be dangerously correlated with the amount of money they owe you.¹ This is because the same macroeconomic factor

drives both the legally required payoff on the security at hand and the default probability of the counterparty. A true credit-adjusted valuation requires the kind of scenario-specific default probabilities that are embedded in the rich flexibility of the reduced form models discussed in detail in Chapter 16.

How do we simulate future values and cash flows on options on futures and forwards in a true default-adjusted fashion? Again, a careful Monte Carlo simulation like that specified in Chapters 19 and 20 is essential to recognizing the slim, but real possibility of the default of the futures exchanges in times of crisis. Similarly, the default of our counterparty on an over-the-counter option on futures and forwards could be analyzed analytically or numerically. As suggested in the previous chapter, rather than doing this mathematical derivation, a carefully structured Monte Carlo simulation can capture the same effects. This is the principal topic in Chapters 36 to 41.

NOTE

1. See van Deventer and Imai (2003) for a detailed discussion of this incident.

Caps and Floors

This chapter was drafted on July 2, 2012, the day that the Barclays PLC chief executive officer resigned in the wake of the \$451 million in fines levied on Barclays for the manipulation of LIBOR. Barclays Chairman Marcus Agius resigned on July 1. Given that LIBOR is the most common reference rate that has historically been used as the reference rate in caps and floors, big changes lie (no pun intended) ahead for the LIBOR market and the related market in caps and floors. LIBOR and the market for interest rate derivatives are so tightly tied together that a Heath, Jarrow, and Morton (HJM) variation called the “LIBOR market model” is described by some as the industry standard model for pricing interest rate derivatives. The LIBOR market model is a very useful point of reference and we recommend that serious readers review Riccardo Rebonato’s excellent book *Modern Pricing of Interest-Rate Derivatives: The LIBOR Market Model and Beyond* (John Wiley & Sons, 2012) for more on that specific implementation. In this chapter, the nature of the reference rate is irrelevant to the analysis, provided that the reference rate is determined in a competitive market with transparent price discovery and the usual assumptions of perfect competition. This description, obviously, does not apply to LIBOR. Therefore any results from the following methods have to be adjusted for the result of criminal activity and market manipulation. The authors recommend that readers initiate no new contracts tied to LIBOR or any other reference rate that does not represent actual transactions. Even in that case, sham transactions have to be eliminated from the reference rate calculation.

We now turn from the ethics of financial markets to more pleasant topics. We are making step-by-step progress in valuing the financial instruments that might be found on the balance sheet of a major financial institution. To reiterate, our ultimate objective is the full valuation of the Jarrow-Merton put option on the balance sheet of the financial institution as the best measure of integrated credit risk, market risk, liquidity risk, and interest rate risk. The Jarrow-Merton put option is also an excellent basis for capital allocation, which is discussed in Chapters 36 through 41. In this chapter, we turn to interest rate caps and floors using the tools from Chapter 23 extensively. We continue to employ the three-factor HJM model implementation from Chapter 9. Normally, in practice, one would use the parameters from Chapter 9 to value caps and floors using a Monte Carlo simulation. For ease of exposition, we use the related bushy tree from Chapter 9.

Putting aside the 2012 LIBOR manipulation scandal, interest rate caps and floors have become, perhaps, the most popular interest rate option-based instrument. They provide a valuable tool for financial managers interested in controlling their exposure

to extremely high or extremely low interest rates. They also provide an excellent basis for the estimation of term structure model parameters as we discussed in Chapter 14. Common uses of caps and floors include the following examples:

- Limiting the maximum rate paid on a floating-rate loan by purchasing a cap
- Limiting the exposure to an effective floor¹ on nonmaturity consumer deposit rates paid by the bank by purchasing a floor in the open market
- Hedging the prepayment risk on a mortgage portfolio
- Hedging the cancellation risk on a portfolio of life insurance policies
- Insuring that the performance of a fixed income portfolio can only improve as rates rise

The purchaser of a cap effectively purchases an insurance policy against interest rate increases. Let's consider a hedger who has a floating-rate loan tied to a three-month U.S. dollar reference rate. By purchasing a 3 percent cap, the hedger insures that they will never pay more than 3 percent on the loan after subtracting payments from the seller of the cap. If rates rise to 5 percent, the hedger pays 5 percent on the loan, but receives 2 percent from the person who sold the cap, for a net expense of 3 percent.

Caps and floors are both traded outright and embedded in common banking and insurance industry assets and liabilities. A number of market conventions deserve mention here. First, payment for the cap or floor typically takes place at time zero when the contract is traded outright. Embedded in a loan or deposit, the cap or floor price would be disguised as an interest rate differential relative to the identical loan or deposit with no cap or floor. Second, the cap or floor strike rate is stated on the same basis as the underlying interest rate. For example, the buyer of the cap above is quoted a strike rate of 3 percent versus a money market reference rate. Most money market interest rates are quoted on an actual/360 basis, so the cap strike rate interest amounts are calculated on an actual/360 day basis also. In other markets, actual/365 day calculations are common. Both the money market reference rate and the cap or floor rate will normally be set from the same reference point (the same "screen" provided by a vendor of real-time information, such as Reuters, Telerate, or Bloomberg) and use the same method for counting business days and the same holiday convention.

We abstract from these market conventions in this chapter to ease exposition. For purposes of this chapter, a strike rate quoted as 3 percent for a money market rate cap for a 91-day period would be converted to a decimal strike rate K using the following:

$$K = 0.03 \frac{91}{360}$$

In applying the formulas, the strike rate should be converted to this K format to allow us to abstract from day count and accrual methods in what follows.

Another important market convention is the use of Black's (1976) futures model for the quotation of caps and floors prices that we have mentioned often in previous chapters. The Black-Scholes (1973) option model was perhaps the biggest innovation in financial theory in the twentieth century, but its well-justified popularity should not obscure the reasons cited in Chapters 3 through 13 for using a term structure model approach instead to model interest rate derivatives. Hull (1993) presents a very lucid explanation of the use of the Black (1976) approximation to cap and floor valuation and the reasons why a term structure model approach works better. We

ignore the Black model in what follows, consistent with our de-emphasis on historical approaches that have been passed by.

At this point, we take advantage of our work in previous chapters to illustrate cap and floor pricing under a term structure model approach using the three-factor HJM model of Chapter 9.

CAPS AS EUROPEAN OPTIONS ON FORWARD RATE AGREEMENTS

In Chapter 23, we analyzed the European option to buy a forward rate agreement at the FRA strike rate of K . Looking at the cash flows for an option on an FRA shows that a cap is really just a special case of an option on an FRA, where the exercise date of the option on the FRA is the same as the expiration date of the FRA. In a single-factor Vasicek model, the bond option approach of Jamshidian [1989] from Chapter 21 can be used to derive a closed form solution for caps and floors. We know from our statistical analysis in Chapter 3 that a one-factor model is not an accurate description of yield curve movements, so this chapter focuses on numerical methods using the three-factor HJM model for expository purposes.

FORMING OTHER CAP-RELATED SECURITIES

A wide variety of cap and floor derivatives can be constructed with this basic building block as a tool:

- Floor prices are calculated in an equivalent way using the formulas for options on a zero-coupon bond in Chapter 23.
- Longer-term caps are constructed as the sum of a number of caplets.
- Longer-term floors are constructed as the sum of a number of floorlets.
- Collars represent a combination of caps and floors (typically purchasing a cap and selling a floor).
- Costless collars are created by finding the floor rate K_f such that the value of the floor exactly equals the value of the cap with cap rate K_c (or vice versa).

We now supplement our examples from Chapter 23 with additional examples of caps, floors, and related instruments.

Valuing a Cap

Recall that from Chapter 9, we know the one-year spot rate of interest at times 0, 1, 2, and 3. There is 1 node at time 0, 4 nodes at time 1, 16 nodes at time 2, and 64 nodes at time 3. We will use the one-year spot rate of interest at each of these nodes as the reference interest rate in valuing a cap. The reference rate values are given in Exhibit 24.1.

Assume we want to buy a cap on this reference rate with a three-year maturity that pays us the difference between the ultimate value of the reference rate and our strike price, which we assume is 3 percent. We assume the reference amount is \$1,000. If we are owed a payment, this payment will be received in arrears at time 4, the point in time at which interest is paid on the reference security. The cash flows are calculated as $\text{Maximum}(0, 1000 * (\text{Spot rate} - 3\%))$. The cash flows associated with this cap are shown in Exhibit 24.2.

EXHIBIT 24.1 Spot Rate Process

State	Row	0	1	2	3
S-1, S-1, S-1	1	0.3003%	2.1653%	4.6388%	6.1106%
S-1, S-1, S-2	2		0.6911%	1.4142%	5.5667%
S-1, S-1, S-3	3		1.0972%	2.6089%	5.1710%
S-1, S-1, S-4	4		1.3611%	4.9179%	8.0375%
S-1, S-2, S-1	5			2.2496%	3.1041%
S-1, S-2, S-2	6			0.5984%	2.5913%
S-1, S-2, S-3	7			0.8121%	2.0291%
S-1, S-2, S-4	8			1.2459%	4.0628%
S-1, S-3, S-1	9			1.9217%	5.2210%
S-1, S-3, S-2	10			1.4149%	1.9784%
S-1, S-3, S-3	11			0.8591%	3.1798%
S-1, S-3, S-4	12			2.8695%	5.5017%
S-1, S-4, S-1	13			2.2088%	5.4252%
S-1, S-4, S-2	14			1.7005%	4.8848%
S-1, S-4, S-3	15			1.1432%	4.4917%
S-1, S-4, S-4	16			3.1593%	7.3396%
S-2, S-1, S-1	17				4.1517%
S-2, S-1, S-2	18				0.9421%
S-2, S-1, S-3	19				2.1313%
S-2, S-1, S-4	20				4.4295%
S-2, S-2, S-1	21				2.2036%
S-2, S-2, S-2	22				0.5531%
S-2, S-2, S-3	23				0.7667%
S-2, S-2, S-4	24				1.2004%
S-2, S-3, S-1	25				2.6231%
S-2, S-3, S-2	26				−0.3734%
S-2, S-3, S-3	27				0.2593%
S-2, S-3, S-4	28				2.1679%
S-2, S-4, S-1	29				1.7764%
S-2, S-4, S-2	30				1.2703%
S-2, S-4, S-3	31				0.7153%
S-2, S-4, S-4	32				2.7228%
S-3, S-1, S-1	33				2.2804%
S-3, S-1, S-2	34				1.7718%
S-3, S-1, S-3	35				1.2141%
S-3, S-1, S-4	36				3.2315%
S-3, S-2, S-1	37				2.2804%
S-3, S-2, S-2	38				1.7718%
S-3, S-2, S-3	39				1.2141%
S-3, S-2, S-4	40				3.2315%
S-3, S-3, S-1	41				3.6354%
S-3, S-3, S-2	42				0.6093%
S-3, S-3, S-3	43				1.2483%
S-3, S-3, S-4	44				3.1757%

(Continued)

EXHIBIT 24.1 (Continued)

State	Row	0	1	2	3
S-3, S-4, S-1	45				4.5070%
S-3, S-4, S-2	46				1.2864%
S-3, S-4, S-3	47				2.4797%
S-3, S-4, S-4	48				4.7858%
S-4, S-1, S-1	49				3.5013%
S-4, S-1, S-2	50				0.3117%
S-4, S-1, S-3	51				1.4934%
S-4, S-1, S-4	52				3.7773%
S-4, S-2, S-1	53				2.5831%
S-4, S-2, S-2	54				2.0730%
S-4, S-2, S-3	55				1.5136%
S-4, S-2, S-4	56				3.5370%
S-4, S-3, S-1	57				2.5831%
S-4, S-3, S-2	58				2.0730%
S-4, S-3, S-3	59				1.5136%
S-4, S-3, S-4	60				3.5370%
S-4, S-4, S-1	61				4.5061%
S-4, S-4, S-2	62				2.2686%
S-4, S-4, S-3	63				2.7252%
S-4, S-4, S-4	64				5.0230%

EXHIBIT 24.2 Time Period

State	Row	0	1	2	3
S-1, S-1, S-1	1		0	0	31.1062
S-1, S-1, S-2	2		0	0	25.6672
S-1, S-1, S-3	3		0	0	21.7103
S-1, S-1, S-4	4		0	0	50.3747
S-1, S-2, S-1	5			0	1.0407
S-1, S-2, S-2	6			0	0.0000
S-1, S-2, S-3	7			0	0.0000
S-1, S-2, S-4	8			0	10.6283
S-1, S-3, S-1	9			0	22.2100
S-1, S-3, S-2	10			0	0.0000
S-1, S-3, S-3	11			0	1.7977
S-1, S-3, S-4	12			0	25.0166
S-1, S-4, S-1	13			0	24.2520
S-1, S-4, S-2	14			0	18.8481
S-1, S-4, S-3	15			0	14.9168
S-1, S-4, S-4	16			0	43.3960
S-2, S-1, S-1	17				11.5175
S-2, S-1, S-2	18				0.0000
S-2, S-1, S-3	19				0.0000

EXHIBIT 24.2 (Continued)

State	Row	0	1	2	3
S-2, S-1, S-4	20				14.2955
S-2, S-2, S-1	21				0.0000
S-2, S-2, S-2	22				0.0000
S-2, S-2, S-3	23				0.0000
S-2, S-2, S-4	24				0.0000
S-2, S-3, S-1	25				0.0000
S-2, S-3, S-2	26				0.0000
S-2, S-3, S-3	27				0.0000
S-2, S-3, S-4	28				0.0000
S-2, S-4, S-1	29				0.0000
S-2, S-4, S-2	30				0.0000
S-2, S-4, S-3	31				0.0000
S-2, S-4, S-4	32				0.0000
S-3, S-1, S-1	33				0.0000
S-3, S-1, S-2	34				0.0000
S-3, S-1, S-3	35				0.0000
S-3, S-1, S-4	36				2.3153
S-3, S-2, S-1	37				0.0000
S-3, S-2, S-2	38				0.0000
S-3, S-2, S-3	39				0.0000
S-3, S-2, S-4	40				2.3153
S-3, S-3, S-1	41				6.3541
S-3, S-3, S-2	42				0.0000
S-3, S-3, S-3	43				0.0000
S-3, S-3, S-4	44				1.7567
S-3, S-4, S-1	45				15.0704
S-3, S-4, S-2	46				0.0000
S-3, S-4, S-3	47				0.0000
S-3, S-4, S-4	48				17.8579
S-4, S-1, S-1	49				5.0128
S-4, S-1, S-2	50				0.0000
S-4, S-1, S-3	51				0.0000
S-4, S-1, S-4	52				7.7735
S-4, S-2, S-1	53				0.0000
S-4, S-2, S-2	54				0.0000
S-4, S-2, S-3	55				0.0000
S-4, S-2, S-4	56				5.3704
S-4, S-3, S-1	57				0.0000
S-4, S-3, S-2	58				0.0000
S-4, S-3, S-3	59				0.0000
S-4, S-3, S-4	60				5.3704
S-4, S-4, S-1	61				15.0608
S-4, S-4, S-2	62				0.0000
S-4, S-4, S-3	63				0.0000
S-4, S-4, S-4	64				20.2298

We drop these cash flows into the cash flow table and multiply times the probability-weighted discount factors from Chapter 9. Exhibit 24.3 shows these cash flows and produces the value of the cap: \$7.7028. Note that the cash flows occur at time 4 even though the reference interest rate is known at time 3 because interest is paid in arrears, the normal money market practice.

Valuing a Floor

What if we had agreed to pay our counterparty whenever the reference interest rate was below 1.5 percent on the same terms as before? Those terms were \$1,000 in the reference money market instrument to be issued at time $T = 3$ for maturity at time $T = 4$. The cash flows on our promise to pay can be written $\text{Minimum}(0, 1000 \times (\text{Spot rate} - 1.5\%))$. Those cash flows are given in Exhibit 24.4.

EXHIBIT 24.3 Maturity of Cash Flow Received

Row	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	0.0000	31.1062
2		0.0000	0.0000	0.0000	25.6672
3		0.0000	0.0000	0.0000	21.7103
4		0.0000	0.0000	0.0000	50.3747
5			0.0000	0.0000	1.0407
6			0.0000	0.0000	0.0000
7			0.0000	0.0000	0.0000
8			0.0000	0.0000	10.6283
9			0.0000	0.0000	22.2100
10			0.0000	0.0000	0.0000
11			0.0000	0.0000	1.7977
12			0.0000	0.0000	25.0166
13			0.0000	0.0000	24.2520
14			0.0000	0.0000	18.8481
15			0.0000	0.0000	14.9168
16			0.0000	0.0000	43.3960
17				0.0000	11.5175
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	14.2955
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000

EXHIBIT 24.3 (Continued)

Row	Current Time				
	0	1	2	3	4
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000
36				0.0000	2.3153
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	2.3153
41				0.0000	6.3541
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	1.7567
45				0.0000	15.0704
46				0.0000	0.0000
47				0.0000	0.0000
48				0.0000	17.8579
49				0.0000	5.0128
50				0.0000	0.0000
51				0.0000	0.0000
52				0.0000	7.7735
53				0.0000	0.0000
54				0.0000	0.0000
55				0.0000	0.0000
56				0.0000	5.3704
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	5.3704
61				0.0000	15.0608
62				0.0000	0.0000
63				0.0000	0.0000
64				0.0000	20.2298
Risk-Neutral Value =					7.7028

EXHIBIT 24.4 Time Period

State	Row Number	0	1	2	3
S-1, S-1, S-1	1		0	0	0.0000
S-1, S-1, S-2	2		0	0	0.0000
S-1, S-1, S-3	3		0	0	0.0000
S-1, S-1, S-4	4		0	0	0.0000
S-1, S-2, S-1	5			0	0.0000
S-1, S-2, S-2	6			0	0.0000
S-1, S-2, S-3	7			0	0.0000
S-1, S-2, S-4	8			0	0.0000
S-1, S-3, S-1	9			0	0.0000
S-1, S-3, S-2	10			0	0.0000
S-1, S-3, S-3	11			0	0.0000
S-1, S-3, S-4	12			0	0.0000
S-1, S-4, S-1	13			0	0.0000
S-1, S-4, S-2	14			0	0.0000
S-1, S-4, S-3	15			0	0.0000
S-1, S-4, S-4	16			0	0.0000
S-2, S-1, S-1	17				0.0000
S-2, S-1, S-2	18				-5.5789
S-2, S-1, S-3	19				0.0000
S-2, S-1, S-4	20				0.0000
S-2, S-2, S-1	21				0.0000
S-2, S-2, S-2	22				-9.4690
S-2, S-2, S-3	23				-7.3327
S-2, S-2, S-4	24				-2.9963
S-2, S-3, S-1	25				0.0000
S-2, S-3, S-2	26				-18.7340
S-2, S-3, S-3	27				-12.4068
S-2, S-3, S-4	28				0.0000
S-2, S-4, S-1	29				0.0000
S-2, S-4, S-2	30				-2.2973
S-2, S-4, S-3	31				-7.8469
S-2, S-4, S-4	32				0.0000
S-3, S-1, S-1	33				0.0000
S-3, S-1, S-2	34				0.0000
S-3, S-1, S-3	35				-2.8592
S-3, S-1, S-4	36				0.0000
S-3, S-2, S-1	37				0.0000
S-3, S-2, S-2	38				0.0000
S-3, S-2, S-3	39				-2.8592
S-3, S-2, S-4	40				0.0000
S-3, S-3, S-1	41				0.0000
S-3, S-3, S-2	42				-8.9066
S-3, S-3, S-3	43				-2.5170

EXHIBIT 24.4 (Continued)

State	Row Number	0	1	2	3
S-3, S-3, S-4	44				0.0000
S-3, S-4, S-1	45				0.0000
S-3, S-4, S-2	46				−2.1355
S-3, S-4, S-3	47				0.0000
S-3, S-4, S-4	48				0.0000
S-4, S-1, S-1	49				0.0000
S-4, S-1, S-2	50				−11.8831
S-4, S-1, S-3	51				−0.0659
S-4, S-1, S-4	52				0.0000
S-4, S-2, S-1	53				0.0000
S-4, S-2, S-2	54				0.0000
S-4, S-2, S-3	55				0.0000
S-4, S-2, S-4	56				0.0000
S-4, S-3, S-1	57				0.0000
S-4, S-3, S-2	58				0.0000
S-4, S-3, S-3	59				0.0000
S-4, S-3, S-4	60				0.0000
S-4, S-4, S-1	61				0.0000
S-4, S-4, S-2	62				0.0000
S-4, S-4, S-3	63				0.0000
S-4, S-4, S-4	64				0.0000

If we drop these cash flows into the cash flow table shown in Exhibit 24.5, we can calculate the value of our promise to pay on this floor contract to be -0.6528 .

Valuing a Floating Rate Loan with a Cap

We now embed the caps and floors in a loan contract. First, we revert to Chapter 19 and value the cash flows of a loan that pays the spot rate of interest in every period (see Exhibit 24.1) in arrears and pays the principal of \$1,000 at time $T = 4$. As we showed in Chapter 19, the value of that loan is the par value of \$1,000 because we are using a valuation yield curve for that borrower, ABC Company, derived from observable zero-coupon bonds that correctly reflect ABC's default risk and related market risk aversion. Exhibit 24.6 shows those cash flows and the resulting valuation, exactly par value of \$1,000, which results from multiplying each cash flow times the relevant probability-weighted discount factor from Chapter 9. Note that, as in Chapter 9, there is one node with a negative interest rate in row 26 so the period 4 payment is less than par value.

Now let's embed a cap on the interest rate in the loan of 3.00 percent, applicable only to the interest payment made at time $T = 4$ based on the reference rate at time $T = 3$. This changes the cash flows to those shown in Exhibit 24.7 and produces a valuation of \$992.2972.

EXHIBIT 24.5 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.0000	0.0000	0.0000	0.0000
2		0.0000	0.0000	0.0000	0.0000
3		0.0000	0.0000	0.0000	0.0000
4		0.0000	0.0000	0.0000	0.0000
5			0.0000	0.0000	0.0000
6			0.0000	0.0000	0.0000
7			0.0000	0.0000	0.0000
8			0.0000	0.0000	0.0000
9			0.0000	0.0000	0.0000
10			0.0000	0.0000	0.0000
11			0.0000	0.0000	0.0000
12			0.0000	0.0000	0.0000
13			0.0000	0.0000	0.0000
14			0.0000	0.0000	0.0000
15			0.0000	0.0000	0.0000
16			0.0000	0.0000	0.0000
17				0.0000	0.0000
18				0.0000	−5.5789
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000
22				0.0000	−9.4690
23				0.0000	−7.3327
24				0.0000	−2.9963
25				0.0000	0.0000
26				0.0000	−18.7340
27				0.0000	−12.4068
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	−2.2973
31				0.0000	−7.8469
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	−2.8592
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	−2.8592
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0000	−8.9066

EXHIBIT 24.5 (Continued)

Row Number	Current Time				
	0	1	2	3	4
43				0.0000	−2.5170
44				0.0000	0.0000
45				0.0000	0.0000
46				0.0000	−2.1355
47				0.0000	0.0000
48				0.0000	0.0000
49				0.0000	0.0000
50				0.0000	−11.8831
51				0.0000	−0.0659
52				0.0000	0.0000
53				0.0000	0.0000
54				0.0000	0.0000
55				0.0000	0.0000
56				0.0000	0.0000
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				0.0000	0.0000
62				0.0000	0.0000
63				0.0000	0.0000
64				0.0000	0.0000
Risk-Neutral Value =					−0.6528

EXHIBIT 24.6 Maturity of Cash Flow Received

Row	Current Time				
	0	1	2	3	4
1	0.0000	3.0032	21.6527	46.3884	1,061.1062
2		3.0032	21.6527	46.3884	1,055.6672
3		3.0032	21.6527	46.3884	1,051.7103
4		3.0032	21.6527	46.3884	1,080.3747
5			6.9105	14.1419	1,031.0407
6			6.9105	14.1419	1,025.9132
7			6.9105	14.1419	1,020.2913
8			6.9105	14.1419	1,040.6283
9			10.9723	26.0890	1,052.2100

(Continued)

EXHIBIT 24.6 *(Continued)*

Row	Current Time				
	0	1	2	3	4
10			10.9723	26.0890	1,019.7841
11			10.9723	26.0890	1,031.7977
12			10.9723	26.0890	1,055.0166
13			13.6107	49.1794	1,054.2520
14			13.6107	49.1794	1,048.8481
15			13.6107	49.1794	1,044.9168
16			13.6107	49.1794	1,073.3960
17				22.4959	1,041.5175
18				22.4959	1,009.4211
19				22.4959	1,021.3126
20				22.4959	1,044.2955
21				5.9835	1,022.0359
22				5.9835	1,005.5310
23				5.9835	1,007.6673
24				5.9835	1,012.0037
25				8.1208	1,026.2311
26				8.1208	996.2660
27				8.1208	1,002.5932
28				8.1208	1,021.6787
29				12.4592	1,017.7642
30				12.4592	1,012.7027
31				12.4592	1,007.1531
32				12.4592	1,027.2283
33				19.2173	1,022.8044
34				19.2173	1,017.7178
35				19.2173	1,012.1408
36				19.2173	1,032.3153
37				14.1486	1,022.8044
38				14.1486	1,017.7178
39				14.1486	1,012.1408
40				14.1486	1,032.3153
41				8.5911	1,036.3541
42				8.5911	1,006.0934
43				8.5911	1,012.4830
44				8.5911	1,031.7567
45				28.6949	1,045.0704
46				28.6949	1,012.8645
47				28.6949	1,024.7966
48				28.6949	1,047.8579
49				22.0883	1,035.0128
50				22.0883	1,003.1169
51				22.0883	1,014.9341
52				22.0883	1,037.7735
53				17.0053	1,025.8313

EXHIBIT 24.6 (Continued)

Row	Current Time				
	0	1	2	3	4
54				17.0053	1,020.7297
55				17.0053	1,015.1361
56				17.0053	1,035.3704
57				11.4322	1,025.8313
58				11.4322	1,020.7297
59				11.4322	1,015.1361
60				11.4322	1,035.3704
61				31.5926	1,045.0608
62				31.5926	1,022.6860
63				31.5926	1,027.2519
64				31.5926	1,050.2298
Risk-Neutral Value =					1,000.0000

EXHIBIT 24.7 Maturity of Cash Flow Received

Row	Current Time				
	0	1	2	3	4
1	0.0000	3.0032	21.6527	46.3884	1,030.0000
2		3.0032	21.6527	46.3884	1,030.0000
3		3.0032	21.6527	46.3884	1,030.0000
4		3.0032	21.6527	46.3884	1,030.0000
5			6.9105	14.1419	1,030.0000
6			6.9105	14.1419	1,025.9132
7			6.9105	14.1419	1,020.2913
8			6.9105	14.1419	1,030.0000
9			10.9723	26.0890	1,030.0000
10			10.9723	26.0890	1,019.7841
11			10.9723	26.0890	1,030.0000
12			10.9723	26.0890	1,030.0000
13			13.6107	49.1794	1,030.0000
14			13.6107	49.1794	1,030.0000
15			13.6107	49.1794	1,030.0000
16			13.6107	49.1794	1,030.0000
17				22.4959	1,030.0000
18				22.4959	1,009.4211
19				22.4959	1,021.3126
20				22.4959	1,030.0000
21				5.9835	1,022.0359
22				5.9835	1,005.5310

(Continued)

EXHIBIT 24.7 (Continued)

Row	Current Time				
	0	1	2	3	4
23				5.9835	1,007.6673
24				5.9835	1,012.0037
25				8.1208	1,026.2311
26				8.1208	996.2660
27				8.1208	1,002.5932
28				8.1208	1,021.6787
29				12.4592	1,017.7642
30				12.4592	1,012.7027
31				12.4592	1,007.1531
32				12.4592	1,027.2283
33				19.2173	1,022.8044
34				19.2173	1,017.7178
35				19.2173	1,012.1408
36				19.2173	1,030.0000
37				14.1486	1,022.8044
38				14.1486	1,017.7178
39				14.1486	1,012.1408
40				14.1486	1,030.0000
41				8.5911	1,030.0000
42				8.5911	1,006.0934
43				8.5911	1,012.4830
44				8.5911	1,030.0000
45				28.6949	1,030.0000
46				28.6949	1,012.8645
47				28.6949	1,024.7966
48				28.6949	1,030.0000
49				22.0883	1,030.0000
50				22.0883	1,003.1169
51				22.0883	1,014.9341
52				22.0883	1,030.0000
53				17.0053	1,025.8313
54				17.0053	1,020.7297
55				17.0053	1,015.1361
56				17.0053	1,030.0000
57				11.4322	1,025.8313
58				11.4322	1,020.7297
59				11.4322	1,015.1361
60				11.4322	1,030.0000
61				31.5926	1,030.0000
62				31.5926	1,022.6860
63				31.5926	1,027.2519
64				31.5926	1,030.0000
Risk-Neutral Value =					992.2972

To check that there is no arbitrage between the loan market and cap market, the value of the loan with no cap (\$1000.000) less the value of the cap that the lender has given to the borrower (\$7.7028) should equal the value of the loan. We can confirm that the loan value of \$992.2972 matches exactly.

VALUE OF A LOAN WITH A CAP AND A FLOOR

Now we modify the terms of the loan one more time. We assume that the borrower agrees to a floor of 1.5 percent on the interest payment to be made at time $T = 4$ based on the reference rate levels prevailing at time $T = 3$. In that case, the cash flows are as shown in Exhibit 24.8 and the resulting value of the loan is \$992.9501.

In the absence of arbitrage, the loan value should increase by 0.6528 (which is the value of the floor agreed by the borrower to the lender's favor) and we confirm that $\$992.2972 + 0.6528 = 992.9501$ (after rounding).

EXHIBIT 24.8 Maturity of Cash Flow Received

Row	Current Time				
	0	1	2	3	4
1	0.0000	3.0032	21.6527	46.3884	1,030.0000
2		3.0032	21.6527	46.3884	1,030.0000
3		3.0032	21.6527	46.3884	1,030.0000
4		3.0032	21.6527	46.3884	1,030.0000
5			6.9105	14.1419	1,030.0000
6			6.9105	14.1419	1,025.9132
7			6.9105	14.1419	1,020.2913
8			6.9105	14.1419	1,030.0000
9			10.9723	26.0890	1,030.0000
10			10.9723	26.0890	1,019.7841
11			10.9723	26.0890	1,030.0000
12			10.9723	26.0890	1,030.0000
13			13.6107	49.1794	1,030.0000
14			13.6107	49.1794	1,030.0000
15			13.6107	49.1794	1,030.0000
16			13.6107	49.1794	1,030.0000
17				22.4959	1,030.0000
18				22.4959	1,015.0000
19				22.4959	1,021.3126
20				22.4959	1,030.0000
21				5.9835	1,022.0359
22				5.9835	1,015.0000
23				5.9835	1,015.0000
24				5.9835	1,015.0000

(Continued)

EXHIBIT 24.8 *(Continued)*

Row	Current Time				
	0	1	2	3	4
25				8.1208	1,026.2311
26				8.1208	1,015.0000
27				8.1208	1,015.0000
28				8.1208	1,021.6787
29				12.4592	1,017.7642
30				12.4592	1,015.0000
31				12.4592	1,015.0000
32				12.4592	1,027.2283
33				19.2173	1,022.8044
34				19.2173	1,017.7178
35				19.2173	1,015.0000
36				19.2173	1,030.0000
37				14.1486	1,022.8044
38				14.1486	1,017.7178
39				14.1486	1,015.0000
40				14.1486	1,030.0000
41				8.5911	1,030.0000
42				8.5911	1,015.0000
43				8.5911	1,015.0000
44				8.5911	1,030.0000
45				28.6949	1,030.0000
46				28.6949	1,015.0000
47				28.6949	1,024.7966
48				28.6949	1,030.0000
49				22.0883	1,030.0000
50				22.0883	1,015.0000
51				22.0883	1,015.0000
52				22.0883	1,030.0000
53				17.0053	1,025.8313
54				17.0053	1,020.7297
55				17.0053	1,015.1361
56				17.0053	1,030.0000
57				11.4322	1,025.8313
58				11.4322	1,020.7297
59				11.4322	1,015.1361
60				11.4322	1,030.0000
61				31.5926	1,030.0000
62				31.5926	1,022.6860
63				31.5926	1,027.2519
64				31.5926	1,030.0000
Risk-Neutral Value =					992.9501

Variations on Caps and Floors

We saw in Chapters 6 through 9 that it is easy to value a digital option on interest rates in the HJM framework. Indeed, there are an infinite variety of “derivatives” that can be priced in the HJM framework. Any function of interest rates can be dropped into the cash flow table and valued on a risk-neutral basis. Extending the coverage of the caps and floors to the first three years of the loan is very simple. This means that the interest rate on the loan will vary only between 1.5 percent and 3.00 percent for the life of the loan. The variety of structures that can be priced (and hedged) is truly mind-boggling. That being said, only a few of the structures one can imagine have real value to both borrowers and lenders so there is a practical limit to the complexity that one sees in the marketplace. As we saw in Chapter 20, in the case of the CDO market, the level of complexity exceeded the level of complexity that could be quickly and accurately priced, and the CDO structure has gone into hibernation, perhaps permanently.

In valuing and hedging complex interest rate derivatives, the analyst must be sure that the HJM framework is properly adjusted. For example, the simple four-step bushy tree from the three-factor HJM model in Chapter 9 has only four nodes at time $T = 1$. If one is pricing instruments with a short-term maturity, the granularity of the analysis has to be much greater than four outcomes in a bushy tree or Monte Carlo simulation. The degree of granularity necessary can be calculated explicitly once an analyst can state the accuracy of valuation that is necessary. One such definition of accuracy is the sentence “I need the model valuation to be within the bid-offer spread for this instrument 95 percent of the time.” This defines the maximum sampling error that can be tolerated, and that in turn defines the number of nodes on a bushy tree or the number of Monte Carlo scenarios that are needed. See Fishman (1996) for an example of this calculation.

In the remainder of this book, we cover the most practical variations on interest rate structures.

MEASURING THE CREDIT RISK OF COUNTERPARTIES ON CAPS AND FLOORS

All of our comments in Chapters 24 to 26 regarding the valuation of “vulnerable” derivatives apply here as well. Since the default probability of our counterparty on the cap or floor is very likely to be driven by interest rate levels, the credit risk-adjusted valuation of caps and floors will involve either a complex closed form solution in a reduced form model or (more accurately) the equivalent Monte Carlo simulation with explicit interest rate drivers of our counterparties’ default probabilities. This is a standard capability of a best practice enterprise risk management system. Clearly, a 10-year cap on interest rates purchased in the early 1980s from a soon-to-be bankrupt savings and loan association (driven into bankruptcy by high interest rates in the United States) is much less valuable than the same cap on U.S. interest rates provided by a non-U.S. financial institution since their exposure to U.S. interest rate risk is so much less than a savings and loan association in the early 1980s, just prior to the worst spike in interest rates in U.S. history.

Best practice requires that we take this into account in order to avoid the \$1 trillion in losses that U.S. taxpayers suffered from failed financial institutions from the savings and loan crisis. Similarly, a systematic risk factor like home prices drove much of the second \$1 trillion bailout in the United States in the 2006 to 2011 period. The FDIC Loss Distribution Model announced on December 10, 2003 (and discussed frequently in this book), does exactly this. Given that the 2006 to 2011 credit crisis firmly validated the FDIC modeling approach (which unfortunately came too late to stave off the crisis), one would hope that the probability of the next crisis can be calculated precisely and early. Perhaps this will lead to preventative action by bankers and regulators. Then again, perhaps it will not. The readers of this book will have much to say about that.

NOTE

1. For many years, U.S. bankers felt that there was an interest rate floor on consumer savings deposits, a rate below which savings deposit rates could not go. Now that we have extensive experience with extremely low interest rates in Japan, the United States, and other countries, it is apparent that this floor is effectively zero. What many initially thought was a floor was more likely the lagged response of consumer savings deposit rates to changing market rates. See Chapter 30 for more on this issue.

Interest Rate Swaps and Swaptions

Interest rate swaps have become the most successful over-the-counter derivative security in the world, and for this reason it is particularly important that we include their credit-adjusted valuation in this book. We emphasize, in the wake of the scandal that enveloped Barclays for LIBOR manipulation, that nothing in this chapter is specific to LIBOR. As in Chapter 24, we assume that the short-term reference rate and other point on the yield curve from which swap pricing is derived are determined in a perfectly competitive, transparent market. LIBOR does not fit that description and it should be rejected as a pricing index by sophisticated (and honest) market participants.

Even when swap counterparties offer transactions through an AAA-rated special purpose vehicle, credit-adjusted valuation is critical because the default probabilities of major financial institutions are highly correlated, as the credit crisis chronology in the Introduction documents. This means that a guarantee by Bank A of another investment bank's special purpose vehicle has much less value than a guarantee by another entity with a much lower correlation in default probabilities and events of default.

Valuation and simulation of interest rate swaps and swaptions is also critical to meet the accounting requirements for hedges under the constantly evolving hedge accounting requirements of Generally Accepted Accounting Principles (GAAP) in the United States and International Financial Reporting Standards (IFRS) around the world. In this chapter, we will continue to employ the three-factor Heath, Jarrow, and Morton (HJM) framework of Chapter 9 for valuation analysis.

INTEREST RATE SWAP BASICS

The most common plain vanilla interest rate swap between two counterparties calls for counterparty A to make fixed rate payments at even intervals (normally semi-annual) to counterparty B and for counterparty B to make floating-rate payments, typically at three-month or six-month intervals, to counterparty A. Seen from the perspective of counterparty A, the swap has the same net cash flow as if counterparty A issued a fixed rate bond and purchased a floating-rate bond (except in the case of default, where the principal amount is relevant in the bond case but not in the swap case). From the perspective of counterparty B, the swap cash flows are identical to the case where counterparty B purchases a fixed rate bond with the proceeds of a floating-rate bond issue (again, except in the case of default).

Why are interest rate swaps so popular? The primary reasons include the vast liquidity of the swap markets in the major currencies, the low issuance costs compared to traditional bonds even under a medium-term note program or shelf registration, the ease of reversing a position with an offsetting transaction, and the (historical) lack of daily mark-to-market margin requirements for swap market participants of good quality. In the wake of the credit crisis of 2006 to 2011, however, nearly all counterparties now face daily margin requirements, but at least margins are settled on a net basis rather than on a transaction by transaction basis. While the cash flows would be the same in either case, the operational efficiency of margining on net exposure is overwhelming. Operationally, a swap has historically been much less work than a similar position in interest rate futures with the daily mark-to-market requirements. In the wake of the pervasive troubles of the largest swap counterparties, however, the move toward exchange-settled transactions may well reverse this trend in the years ahead.

Many market participants value swaps as if both the fixed and the floating side involved principal payments at maturity. Since the amounts net to zero, this is a good approximation when both parties have zero credit risk. As the credit risk of each party diverges from zero, this becomes a more dangerous assumption, so we won't use it in this chapter.

VALUING THE INTEREST RATE SWAPS

To illustrate the valuation of interest rate swaps in the general HJM valuation framework, we use the assumptions of the three-factor HJM example in Chapter 9. We assume that the short-term, floating-rate side of the interest rate swap pays at one-year intervals at the one-year spot rate. We know from Exhibit 9.22 in Chapter 9 that the value of a bond paying 3 percent interest at annual intervals has a present value of 104.7071. This is the value of the fixed side of a swap when the annual interest rate is 3 percent and we include an exchange of principal. We know from Exhibit 19.9 in Chapter 19 that the present value of a floating rate transaction with \$100 principal amount that matures at time $T = 4$ is par: 100.0000. The value of this swap to the party who receives the fixed rate side and pays the floating-rate side is $104.7071 - 100.0000 = 4.7071$. If this were not true, our no-arbitrage assumptions about the marketplace would be proven false.

This example assumed that the principal amounts were exchanged at maturity. If we do away with that assumption, we know that the value of \$3 received annually will be equal to the value of the bond (104.7071) less the present value of \$100 received in four years. Using our standard zero-coupon bonds as input, we know the present value of a dollar received in four years is 0.9308550992. This means the value of \$3 received annually is $104.7071 - 93.0855 = 11.6216$. We can confirm this by simply valuing the fixed coupons using the zero-coupon bond prices we are given for maturity at times 1, 2, 3, and 4 or, as shown in Exhibit 25.1, we can insert \$3 in the cash flow table at every node on the bushy tree to confirm that the value of the fixed payments is indeed 11.6216.

Similarly, the floating rate payments, in the absence of arbitrage, should total $100.0000 - 93.0855 = 6.9145$. We can confirm this by inserting the floating-rate

EXHIBIT 25.1 Maturity of Cash Flow Received

Row	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	3.0000
2		3.0000	3.0000	3.0000	3.0000
3		3.0000	3.0000	3.0000	3.0000
4		3.0000	3.0000	3.0000	3.0000
5			3.0000	3.0000	3.0000
6			3.0000	3.0000	3.0000
7			3.0000	3.0000	3.0000
8			3.0000	3.0000	3.0000
9			3.0000	3.0000	3.0000
10			3.0000	3.0000	3.0000
11			3.0000	3.0000	3.0000
12			3.0000	3.0000	3.0000
13			3.0000	3.0000	3.0000
14			3.0000	3.0000	3.0000
15			3.0000	3.0000	3.0000
16			3.0000	3.0000	3.0000
17				3.0000	3.0000
18				3.0000	3.0000
19				3.0000	3.0000
20				3.0000	3.0000
21				3.0000	3.0000
22				3.0000	3.0000
23				3.0000	3.0000
24				3.0000	3.0000
25				3.0000	3.0000
26				3.0000	3.0000
27				3.0000	3.0000
28				3.0000	3.0000
29				3.0000	3.0000
30				3.0000	3.0000
31				3.0000	3.0000
32				3.0000	3.0000
33				3.0000	3.0000
34				3.0000	3.0000
35				3.0000	3.0000
36				3.0000	3.0000
37				3.0000	3.0000
38				3.0000	3.0000
39				3.0000	3.0000
40				3.0000	3.0000
41				3.0000	3.0000

(Continued)

EXHIBIT 25.1 *(Continued)*

Row	Current Time				
	0	1	2	3	4
42				3.0000	3.0000
43				3.0000	3.0000
44				3.0000	3.0000
45				3.0000	3.0000
46				3.0000	3.0000
47				3.0000	3.0000
48				3.0000	3.0000
49				3.0000	3.0000
50				3.0000	3.0000
51				3.0000	3.0000
52				3.0000	3.0000
53				3.0000	3.0000
54				3.0000	3.0000
55				3.0000	3.0000
56				3.0000	3.0000
57				3.0000	3.0000
58				3.0000	3.0000
59				3.0000	3.0000
60				3.0000	3.0000
61				3.0000	3.0000
62				3.0000	3.0000
63				3.0000	3.0000
64				3.0000	3.0000
Risk-Neutral Value =					11.62159

payments (but no principal payments) in the cash flow table and confirming the 6.9145 value in Exhibit 25.2.

Alternatively, we can calculate the net cash flow received by the counterparty who receives the fixed rate payment.

Model Risk Alert

This netting is only appropriate if the two swap counterparties are riskless.

If we calculate the fixed payment less the floating payment at every node, we get the cash flows in Exhibit 25.3. As shown, the value of the swap should be the value of the fixed payments less the value of the floating payments, $11.6216 - 6.9145 = 4.7071$.

We now ask the question “What level of interest rate should prevail on the fixed side of the swap” for new swaps?

EXHIBIT 25.2 Time Period

State	Row	0	1	2	3	4
S-1, S-1, S-1	1		0.3003	2.1653	4.6388	6.1106
S-1, S-1, S-2	2		0.3003	2.1653	4.6388	5.5667
S-1, S-1, S-3	3		0.3003	2.1653	4.6388	5.1710
S-1, S-1, S-4	4		0.3003	2.1653	4.6388	8.0375
S-1, S-2, S-1	5			0.6911	1.4142	3.1041
S-1, S-2, S-2	6			0.6911	1.4142	2.5913
S-1, S-2, S-3	7			0.6911	1.4142	2.0291
S-1, S-2, S-4	8			0.6911	1.4142	4.0628
S-1, S-3, S-1	9			1.0972	2.6089	5.2210
S-1, S-3, S-2	10			1.0972	2.6089	1.9784
S-1, S-3, S-3	11			1.0972	2.6089	3.1798
S-1, S-3, S-4	12			1.0972	2.6089	5.5017
S-1, S-4, S-1	13			1.3611	4.9179	5.4252
S-1, S-4, S-2	14			1.3611	4.9179	4.8848
S-1, S-4, S-3	15			1.3611	4.9179	4.4917
S-1, S-4, S-4	16			1.3611	4.9179	7.3396
S-2, S-1, S-1	17				2.2496	4.1517
S-2, S-1, S-2	18				2.2496	0.9421
S-2, S-1, S-3	19				2.2496	2.1313
S-2, S-1, S-4	20				2.2496	4.4295
S-2, S-2, S-1	21				0.5984	2.2036
S-2, S-2, S-2	22				0.5984	0.5531
S-2, S-2, S-3	23				0.5984	0.7667
S-2, S-2, S-4	24				0.5984	1.2004
S-2, S-3, S-1	25				0.8121	2.6231
S-2, S-3, S-2	26				0.8121	-0.3734
S-2, S-3, S-3	27				0.8121	0.2593
S-2, S-3, S-4	28				0.8121	2.1679
S-2, S-4, S-1	29				1.2459	1.7764
S-2, S-4, S-2	30				1.2459	1.2703
S-2, S-4, S-3	31				1.2459	0.7153
S-2, S-4, S-4	32				1.2459	2.7228
S-3, S-1, S-1	33				1.9217	2.2804
S-3, S-1, S-2	34				1.9217	1.7718
S-3, S-1, S-3	35				1.9217	1.2141
S-3, S-1, S-4	36				1.9217	3.2315
S-3, S-2, S-1	37				1.4149	2.2804
S-3, S-2, S-2	38				1.4149	1.7718
S-3, S-2, S-3	39				1.4149	1.2141
S-3, S-2, S-4	40				1.4149	3.2315
S-3, S-3, S-1	41				0.8591	3.6354
S-3, S-3, S-2	42				0.8591	0.6093
S-3, S-3, S-3	43				0.8591	1.2483

(Continued)

EXHIBIT 25.2 (Continued)

State	Row	0	1	2	3	4
S-3, S-3, S-4	44				0.8591	3.1757
S-3, S-4, S-1	45				2.8695	4.5070
S-3, S-4, S-2	46				2.8695	1.2864
S-3, S-4, S-3	47				2.8695	2.4797
S-3, S-4, S-4	48				2.8695	4.7858
S-4, S-1, S-1	49				2.2088	3.5013
S-4, S-1, S-2	50				2.2088	0.3117
S-4, S-1, S-3	51				2.2088	1.4934
S-4, S-1, S-4	52				2.2088	3.7773
S-4, S-2, S-1	53				1.7005	2.5831
S-4, S-2, S-2	54				1.7005	2.0730
S-4, S-2, S-3	55				1.7005	1.5136
S-4, S-2, S-4	56				1.7005	3.5370
S-4, S-3, S-1	57				1.1432	2.5831
S-4, S-3, S-2	58				1.1432	2.0730
S-4, S-3, S-3	59				1.1432	1.5136
S-4, S-3, S-4	60				1.1432	3.5370
S-4, S-4, S-1	61				3.1593	4.5061
S-4, S-4, S-2	62				3.1593	2.2686
S-4, S-4, S-3	63				3.1593	2.7252
S-4, S-4, S-4	64				3.1593	5.0230
Risk-Neutral Value =						6.91449

EXHIBIT 25.3 Maturity of Cash Flow Received

Row	Current Time				
	0	1	2	3	4
1	0.0000	2.6997	0.8347	−1.6388	−3.1106
2		2.6997	0.8347	−1.6388	−2.5667
3		2.6997	0.8347	−1.6388	−2.1710
4		2.6997	0.8347	−1.6388	−5.0375
5			2.3089	1.5858	−0.1041
6			2.3089	1.5858	0.4087
7			2.3089	1.5858	0.9709
8			2.3089	1.5858	−1.0628
9			1.9028	0.3911	−2.2210
10			1.9028	0.3911	1.0216
11			1.9028	0.3911	−0.1798
12			1.9028	0.3911	−2.5017
13			1.6389	−1.9179	−2.4252
14			1.6389	−1.9179	−1.8848

EXHIBIT 25.3 (Continued)

Row	Current Time				
	0	1	2	3	4
15			1.6389	−1.9179	−1.4917
16			1.6389	−1.9179	−4.3396
17				0.7504	−1.1517
18				0.7504	2.0579
19				0.7504	0.8687
20				0.7504	−1.4295
21				2.4016	0.7964
22				2.4016	2.4469
23				2.4016	2.2333
24				2.4016	1.7996
25				2.1879	0.3769
26				2.1879	3.3734
27				2.1879	2.7407
28				2.1879	0.8321
29				1.7541	1.2236
30				1.7541	1.7297
31				1.7541	2.2847
32				1.7541	0.2772
33				1.0783	0.7196
34				1.0783	1.2282
35				1.0783	1.7859
36				1.0783	−0.2315
37				1.5851	0.7196
38				1.5851	1.2282
39				1.5851	1.7859
40				1.5851	−0.2315
41				2.1409	−0.6354
42				2.1409	2.3907
43				2.1409	1.7517
44				2.1409	−0.1757
45				0.1305	−1.5070
46				0.1305	1.7136
47				0.1305	0.5203
48				0.1305	−1.7858
49				0.7912	−0.5013
50				0.7912	2.6883
51				0.7912	1.5066
52				0.7912	−0.7773
53				1.2995	0.4169
54				1.2995	0.9270
55				1.2995	1.4864
56				1.2995	−0.5370
57				1.8568	0.4169

(Continued)

EXHIBIT 25.3 (Continued)

Row	Current Time				
	0	1	2	3	4
58				1.8568	0.9270
59				1.8568	1.4864
60				1.8568	-0.5370
61				-0.1593	-1.5061
62				-0.1593	0.7314
63				-0.1593	0.2748
64				-0.1593	-2.0230
Risk-Neutral Value =					4.7071

THE OBSERVABLE FIXED RATE IN THE SWAP MARKET

We know from the prior section and Chapter 19 that the floating side of a new swap should have a market value of par value, 100.0000, if the principal exchange at maturity is included. We also know, from no-arbitrage principles, that for a new swap the fixed side should also have a net present value of 100.0000 so that the net value to the receiver of the fixed rate is $100.0000 - 100.0000 = 0$. If this were not the case, there would be arbitrage opportunities. For what fixed rate will market value exactly equal par value? We know the answer from Chapter 4.

$$\text{Dollar coupon amount of fixed coupon bond} = \frac{\text{Value} - P(t_n) \text{ Principal}}{\sum_{i=1}^n P(t_i)}$$

When Value = 100, using the zero-coupon bond prices for time 0 given in Chapter 9, the par coupon level is 1.78490814. We can verify that a four-year bond with annual payments at this level and principal of 100 has a net present value of 100.0000 either using the original zero-coupon bond values and formulas from Chapter 4 or using the bushy tree from Chapter 9. If we do the latter, we get the cash flows in Exhibit 25.4 and confirm a value of 100.0000.

Similarly, we can do away with the exchange of principal and value the fixed payments minus the floating-rate payments. When we do so, we confirm in Exhibit 25.5 that the at-market fixed rate swap rate produces a mark-to-market value of the swap equal to par value of 100.0000.

AN INTRODUCTION TO SWAPTIONS

A *swaption* is an agreement between two counterparties that allows counterparty A to enter into an interest rate swap with a previously agreed fixed rate payment C at the option of counterparty A. Swaptions come in two forms. The constant maturity form of swaption prescribes that the interest rate swap has a set maturity, say five

EXHIBIT 25.4 Maturity of Cash Flow Received

Row	Current Time				
	0	1	2	3	4
1	0.0000	1.7849	1.7849	1.7849	101.78490814
2		1.7849	1.7849	1.7849	101.78490814
3		1.7849	1.7849	1.7849	101.78490814
4		1.7849	1.7849	1.7849	101.78490814
5			1.7849	1.7849	101.78490814
6			1.7849	1.7849	101.78490814
7			1.7849	1.7849	101.78490814
8			1.7849	1.7849	101.78490814
9			1.7849	1.7849	101.78490814
10			1.7849	1.7849	101.78490814
11			1.7849	1.7849	101.78490814
12			1.7849	1.7849	101.78490814
13			1.7849	1.7849	101.78490814
14			1.7849	1.7849	101.78490814
15			1.7849	1.7849	101.78490814
16			1.7849	1.7849	101.78490814
17				1.7849	101.78490814
18				1.7849	101.78490814
19				1.7849	101.78490814
20				1.7849	101.78490814
21				1.7849	101.78490814
22				1.7849	101.78490814
23				1.7849	101.78490814
24				1.7849	101.78490814
25				1.7849	101.78490814
26				1.7849	101.78490814
27				1.7849	101.78490814
28				1.7849	101.78490814
29				1.7849	101.78490814
30				1.7849	101.78490814
31				1.7849	101.78490814
32				1.7849	101.78490814
33				1.7849	101.78490814
34				1.7849	101.78490814
35				1.7849	101.78490814
36				1.7849	101.78490814
37				1.7849	101.78490814
38				1.7849	101.78490814
39				1.7849	101.78490814
40				1.7849	101.78490814
41				1.7849	101.78490814

(Continued)

EXHIBIT 25.4 *(Continued)*

Row	Current Time				
	0	1	2	3	4
42				1.7849	101.78490814
43				1.7849	101.78490814
44				1.7849	101.78490814
45				1.7849	101.78490814
46				1.7849	101.78490814
47				1.7849	101.78490814
48				1.7849	101.78490814
49				1.7849	101.78490814
50				1.7849	101.78490814
51				1.7849	101.78490814
52				1.7849	101.78490814
53				1.7849	101.78490814
54				1.7849	101.78490814
55				1.7849	101.78490814
56				1.7849	101.78490814
57				1.7849	101.78490814
58				1.7849	101.78490814
59				1.7849	101.78490814
60				1.7849	101.78490814
61				1.7849	101.78490814
62				1.7849	101.78490814
63				1.7849	101.78490814
64				1.7849	101.78490814
Risk-Neutral Value =					100.0000

EXHIBIT 25.5 Maturity of Cash Flow Received

Row	Current Time				
	0	1	2	3	4
1	0.0000	1.4846	−0.3804	−2.8539	−4.3257
2		1.4846	−0.3804	−2.8539	−3.7818
3		1.4846	−0.3804	−2.8539	−3.3861
4		1.4846	−0.3804	−2.8539	−6.2526
5			1.0939	0.3707	−1.3192
6			1.0939	0.3707	−0.8064
7			1.0939	0.3707	−0.2442
8			1.0939	0.3707	−2.2779
9			0.6877	−0.8240	−3.4361
10			0.6877	−0.8240	−0.1935
11			0.6877	−0.8240	−1.3949

EXHIBIT 25.5 (Continued)

Row	Current Time				
	0	1	2	3	4
12			0.6877	−0.8240	−3.7167
13			0.4238	−3.1330	−3.6403
14			0.4238	−3.1330	−3.0999
15			0.4238	−3.1330	−2.7068
16			0.4238	−3.1330	−5.5547
17				−0.4647	−2.3668
18				−0.4647	0.8428
19				−0.4647	−0.3464
20				−0.4647	−2.6446
21				1.1866	−0.4187
22				1.1866	1.2318
23				1.1866	1.0182
24				1.1866	0.5845
25				0.9728	−0.8382
26				0.9728	2.1583
27				0.9728	1.5256
28				0.9728	−0.3830
29				0.5390	0.0085
30				0.5390	0.5146
31				0.5390	1.0696
32				0.5390	−0.9379
33				−0.1368	−0.4955
34				−0.1368	0.0131
35				−0.1368	0.5708
36				−0.1368	−1.4466
37				0.3700	−0.4955
38				0.3700	0.0131
39				0.3700	0.5708
40				0.3700	−1.4466
41				0.9258	−1.8505
42				0.9258	1.1756
43				0.9258	0.5366
44				0.9258	−1.3908
45				−1.0846	−2.7221
46				−1.0846	0.4985
47				−1.0846	−0.6948
48				−1.0846	−3.0009
49				−0.4239	−1.7164
50				−0.4239	1.4732
51				−0.4239	0.2915
52				−0.4239	−1.9924
53				0.0844	−0.7982

(Continued)

EXHIBIT 25.5 (Continued)

Row	Current Time				
	0	1	2	3	4
54				0.0844	−0.2881
55				0.0844	0.2713
56				0.0844	−1.7521
57				0.6417	−0.7982
58				0.6417	−0.2881
59				0.6417	0.2713
60				0.6417	−1.7521
61				−1.3744	−2.7212
62				−1.3744	−0.4837
63				−1.3744	−0.9403
64				−1.3744	−3.2381
Risk-Neutral Value =					0.0000

years, regardless of when the swaption was exercised. For example, if the exercise period on the swaption is two years, the underlying swap would end up being a five-year swap regardless of whether the swaption was entered into after three months, 12 months, or two years. The fixed maturity date swaption prescribes a maturity date for the swap at the signing date of the contract, and therefore the effective maturity of the swap will shorten if exercise is delayed. Consider a three-year swaption of the fixed maturity type on a swap with an original maturity of 10 years. If the swaption is exercised after two years, the swap that is entered into will have an eight-year life.

Both forms are common. If the swaption is a European swaption, the two methods are equivalent.

Valuation of European Swaptions

The valuation of European swaptions is a direct application of European bond options that we covered in Chapter 21. In that chapter, we showed how to use the HJM three-factor model from Chapter 9 to value such options. The valuation in a swap context is nearly identical. If one has the lucky (and extremely unlikely) circumstance where a one-factor term structure model accurately describes interest rates, Jamshidian's formula for options on coupon-bearing bonds could be employed.

Model Risk Alert

A one-factor model will probably never be accurate enough for practical use.

As shown previously in this chapter, the market value of a new swap must be such that the fixed rate on the swap would cause a coupon-bearing bond with the same payment frequency and maturity date to trade at par.¹ This means we can ignore the floating rate side of the swap for purposes of valuation.² Therefore, a swaption that allows the holder to receive the fixed rate side of the swap is a call

option to buy the equivalent underlying bond. A swaption that allows the holder to pay the fixed rate side of the swap is a put option on the equivalent underlying bond. Having made this translation, the formula is identical to that of Chapter 21's options on coupon-bearing bonds.

Valuation of American Swaptions

Because of the intervening cash flows on swaps, there is the possibility of early exercise on an American swaption and, therefore, the European swaption valuation approach is only a rough approximation to the value of an American swaption. Correct valuation requires the numerical methods we analyze in Chapter 27.

Defaultable Interest Rate Swaps and Swaptions

As discussed in the introduction of this chapter, even the existence of AAA-rated special purpose vehicles does not insulate a financial market participant from the risk of default by a counterparty on an interest rate swap or swaption. Many firms found this out the hard way when Lehman Brothers filed for bankruptcy in September 2008. The 2006–2011 credit crises provides an almost uncountable number of cases where interest rates and other macroeconomic factors (most notably home prices) drove correlated default risk among counterparties. Therefore, we need to take the same approach as that discussed in Chapters 19 through 24, using Monte Carlo simulation of interest rate-driven (and other macro-factor driven) default probabilities to measure the valuation and future distribution of values and cash flows on a credit-adjusted basis. What this means for swap valuation is that different yield curves and different HJM bushy trees should be used to value the different legs of a swap transaction, even if the reference curve used is the same for both halves of the swap. It is the difference in the value of a promise to pay by Company A and Company B that causes this slight complication to be necessary.

We now turn to more exotic swap and option structures.

NOTES

1. And for purposes of this chapter, credit risk.
2. We could not ignore the floating side if the floating side involved a non-zero spread from the pricing index.

Exotic Swap and Options Structures

A difference of opinion makes a market, and that is the major rationale for the development of the market in exotic swaps and exotic options. A second and less happy rationale is the willingness of some market participants to buy and sell securities without much knowledge of their true value. This is particularly dangerous in markets such as the LIBOR market (in which manipulation has been admitted) and the credit default swap market (in which manipulation is suspected and investigations are under way).

In our quest for an integrated measure of interest rate risk and credit risk, we need to handle these exotic structures because they are found on the balance sheet of many institutions both on the buy side and the sell side. Furthermore, two of the authors are ex-investment bankers. For that reason, we recognize the endless quest of investment bankers to invent structures that they can value more accurately than their clients. The collateralized debt obligation market, which we discussed in Chapter 20, is a classic example. The real purpose of this chapter is not only to value specific structures but to show that this approach is general enough to apply to new structures as they emerge. That reflects the great power of the Heath, Jarrow, and Morton (HJM) framework that we demonstrated in increasingly realistic worked examples in Chapters 6 through 9. In this chapter, we again employ the three-factor HJM example from Chapter 9.

In the first sections of this chapter, we show how the HJM framework can be used to value securities that, at least at one point in their life, were considered exotic. “Exotic” means both strange and difficult to value, and it is the latter that often motivates the originators of exotic structures. The examples in this chapter are taken from actual deals found in the swap portfolio of some major U.S. financial institutions.

At the end of the chapter, we discuss how to incorporate the potential default of our counterparty so that we can incorporate these kinds of transactions in our valuation of the Jarrow-Merton put option as a comprehensive risk measure. We return to this discussion in great detail in Chapters 36 through 41.

ARREARS SWAPS

In most interest rate swaps, the floating-rate payment is based on LIBOR (we use LIBOR until an alternative rate based on actual transactions in a competitive and

transparent market replaces LIBOR as an index) in the relevant currency. If a normal swap calls for a six-month LIBOR, the LIBOR rate is set on a reference date six months before the cash interest payment is normally made. In an arrears swap, the LIBOR rate and payment amount will be made in arrears, that is, shortly before the cash payment must be made. Therefore, instead of being known six months in advance and paid at the effective maturity of the underlying index, the rate is set a few days before payment is due at the end of the six-month period. How do we value this kind of swap?

As in previous chapters, we can use either simple term structure models (like Vasicek or Hull and White) to get closed form or algebraic answers (Model Risk Alert), or we can use a more accurate approach with many risk factors driving interest rates. Such an approach, which requires numerical methods, is what we do here based on the framework of the three-factor HJM model. We assume that the fixed rate side of the swap is the same 3 percent interest rate on the \$100 notional amount that we analyzed in Chapter 25. For exposition purposes, we shorten the maturity from four years (in Chapter 25) to three years. We again drop those cash flows into our cash flow table and confirm, after multiplying by the correct probability-weighted discount factors, that the value of the fixed rate side of the swap is 8.8290 as shown in Exhibit 26.1.

EXHIBIT 26.1 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	0.0000
2		3.0000	3.0000	3.0000	0.0000
3		3.0000	3.0000	3.0000	0.0000
4		3.0000	3.0000	3.0000	0.0000
5			3.0000	3.0000	0.0000
6			3.0000	3.0000	0.0000
7			3.0000	3.0000	0.0000
8			3.0000	3.0000	0.0000
9			3.0000	3.0000	0.0000
10			3.0000	3.0000	0.0000
11			3.0000	3.0000	0.0000
12			3.0000	3.0000	0.0000
13			3.0000	3.0000	0.0000
14			3.0000	3.0000	0.0000
15			3.0000	3.0000	0.0000
16			3.0000	3.0000	0.0000
17				3.0000	0.0000
18				3.0000	0.0000
19				3.0000	0.0000
20				3.0000	0.0000
21				3.0000	0.0000

(Continued)

EXHIBIT 26.1 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
22				3.0000	0.0000
23				3.0000	0.0000
24				3.0000	0.0000
25				3.0000	0.0000
26				3.0000	0.0000
27				3.0000	0.0000
28				3.0000	0.0000
29				3.0000	0.0000
30				3.0000	0.0000
31				3.0000	0.0000
32				3.0000	0.0000
33				3.0000	0.0000
34				3.0000	0.0000
35				3.0000	0.0000
36				3.0000	0.0000
37				3.0000	0.0000
38				3.0000	0.0000
39				3.0000	0.0000
40				3.0000	0.0000
41				3.0000	0.0000
42				3.0000	0.0000
43				3.0000	0.0000
44				3.0000	0.0000
45				3.0000	0.0000
46				3.0000	0.0000
47				3.0000	0.0000
48				3.0000	0.0000
49				3.0000	0.0000
50				3.0000	0.0000
51				3.0000	0.0000
52				3.0000	0.0000
53				3.0000	0.0000
54				3.0000	0.0000
55				3.0000	0.0000
56				3.0000	0.0000
57				3.0000	0.0000
58				3.0000	0.0000
59				3.0000	0.0000
60				3.0000	0.0000
61				3.0000	0.0000
62				3.0000	0.0000
63				3.0000	0.0000
64				3.0000	0.0000
Risk-Neutral Value =					8.829025

As always, we could have calculated this value using the present value factors for times 1, 2, and 3 instead of the bushy tree. In the normal swap structure, the floating-rate cash flow at time $T = 1$ is determined by the relevant reference rate prevailing at time $T = 0$. In this example, however, the reference rate is effectively determined at time $T = 1$ for payment (ignoring a few days lag) at time $T = 1$. Because we have a three-factor bushy tree from Chapter 9, we have four different reference rate values at time $T = 1$, 16 at time $T = 2$, and 64 at time $T = 3$. The floating rate interest payments are simply \$100 in notional principal times the prevailing short rate at each node, paid immediately, not one period later. These cash flows are shown in Exhibit 26.2 and their risk-neutral value is 6.8243.

EXHIBIT 26.2 Time Period

State	Row Number	0	1	2	3	4
S-1, S-1, S-1	1		2.1653	4.6388	6.1106	0.0000
S-1, S-1, S-2	2		0.6911	1.4142	5.5667	0.0000
S-1, S-1, S-3	3		1.0972	2.6089	5.1710	0.0000
S-1, S-1, S-4	4		1.3611	4.9179	8.0375	0.0000
S-1, S-2, S-1	5			2.2496	3.1041	0.0000
S-1, S-2, S-2	6			0.5984	2.5913	0.0000
S-1, S-2, S-3	7			0.8121	2.0291	0.0000
S-1, S-2, S-4	8			1.2459	4.0628	0.0000
S-1, S-3, S-1	9			1.9217	5.2210	0.0000
S-1, S-3, S-2	10			1.4149	1.9784	0.0000
S-1, S-3, S-3	11			0.8591	3.1798	0.0000
S-1, S-3, S-4	12			2.8695	5.5017	0.0000
S-1, S-4, S-1	13			2.2088	5.4252	0.0000
S-1, S-4, S-2	14			1.7005	4.8848	0.0000
S-1, S-4, S-3	15			1.1432	4.4917	0.0000
S-1, S-4, S-4	16			3.1593	7.3396	0.0000
S-2, S-1, S-1	17				4.1517	0.0000
S-2, S-1, S-2	18				0.9421	0.0000
S-2, S-1, S-3	19				2.1313	0.0000
S-2, S-1, S-4	20				4.4295	0.0000
S-2, S-2, S-1	21				2.2036	0.0000
S-2, S-2, S-2	22				0.5531	0.0000
S-2, S-2, S-3	23				0.7667	0.0000
S-2, S-2, S-4	24				1.2004	0.0000
S-2, S-3, S-1	25				2.6231	0.0000
S-2, S-3, S-2	26				-0.3734	0.0000
S-2, S-3, S-3	27				0.2593	0.0000
S-2, S-3, S-4	28				2.1679	0.0000
S-2, S-4, S-1	29				1.7764	0.0000
S-2, S-4, S-2	30				1.2703	0.0000
S-2, S-4, S-3	31				0.7153	0.0000

(Continued)

EXHIBIT 26.2 (Continued)

State	Row Number	0	1	2	3	4
S-2, S-4, S-4	32				2.7228	0.0000
S-3, S-1, S-1	33				2.2804	0.0000
S-3, S-1, S-2	34				1.7718	0.0000
S-3, S-1, S-3	35				1.2141	0.0000
S-3, S-1, S-4	36				3.2315	0.0000
S-3, S-2, S-1	37				2.2804	0.0000
S-3, S-2, S-2	38				1.7718	0.0000
S-3, S-2, S-3	39				1.2141	0.0000
S-3, S-2, S-4	40				3.2315	0.0000
S-3, S-3, S-1	41				3.6354	0.0000
S-3, S-3, S-2	42				0.6093	0.0000
S-3, S-3, S-3	43				1.2483	0.0000
S-3, S-3, S-4	44				3.1757	0.0000
S-3, S-4, S-1	45				4.5070	0.0000
S-3, S-4, S-2	46				1.2864	0.0000
S-3, S-4, S-3	47				2.4797	0.0000
S-3, S-4, S-4	48				4.7858	0.0000
S-4, S-1, S-1	49				3.5013	0.0000
S-4, S-1, S-2	50				0.3117	0.0000
S-4, S-1, S-3	51				1.4934	0.0000
S-4, S-1, S-4	52				3.7773	0.0000
S-4, S-2, S-1	53				2.5831	0.0000
S-4, S-2, S-2	54				2.0730	0.0000
S-4, S-2, S-3	55				1.5136	0.0000
S-4, S-2, S-4	56				3.5370	0.0000
S-4, S-3, S-1	57				2.5831	0.0000
S-4, S-3, S-2	58				2.0730	0.0000
S-4, S-3, S-3	59				1.5136	0.0000
S-4, S-3, S-4	60				3.5370	0.0000
S-4, S-4, S-1	61				4.5061	0.0000
S-4, S-4, S-2	62				2.2686	0.0000
S-4, S-4, S-3	63				2.7252	0.0000
S-4, S-4, S-4	64				5.0230	0.0000
Risk-Neutral Value =						6.82433

If both parties to the arrears swap are default free, the arbitrage-free value of the interest rate swap is easy to calculate without using the cash flow table. If our swap counterparty receives the fixed rate side and pays the floating-rate side, the value of the swap is $8.8290 - 6.8243 = 2.0047$. Alternatively, we can calculate the net cash flows and insert them into the cash flow table. We calculate the risk-neutral value of these net cash flows and get the same value, 2.0047, as shown in Exhibit 26.3.

The equilibrium spread for which an arrears swap has an efficient market value of zero is easy to calculate. We can either vary the fixed rate on the swap or add (or subtract) a constant spread from the floating reference rate such that the market value of the swap is zero. This is easy to do in common spreadsheet software. We let

EXHIBIT 26.3 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.8347	−1.6388	−3.1106	0.0000
2		2.3089	1.5858	−2.5667	0.0000
3		1.9028	0.3911	−2.1710	0.0000
4		1.6389	−1.9179	−5.0375	0.0000
5			0.7504	−0.1041	0.0000
6			2.4016	0.4087	0.0000
7			2.1879	0.9709	0.0000
8			1.7541	−1.0628	0.0000
9			1.0783	−2.2210	0.0000
10			1.5851	1.0216	0.0000
11			2.1409	−0.1798	0.0000
12			0.1305	−2.5017	0.0000
13			0.7912	−2.4252	0.0000
14			1.2995	−1.8848	0.0000
15			1.8568	−1.4917	0.0000
16			−0.1593	−4.3396	0.0000
17				−1.1517	0.0000
18				2.0579	0.0000
19				0.8687	0.0000
20				−1.4295	0.0000
21				0.7964	0.0000
22				2.4469	0.0000
23				2.2333	0.0000
24				1.7996	0.0000
25				0.3769	0.0000
26				3.3734	0.0000
27				2.7407	0.0000
28				0.8321	0.0000
29				1.2236	0.0000
30				1.7297	0.0000
31				2.2847	0.0000
32				0.2772	0.0000
33				0.7196	0.0000
34				1.2282	0.0000
35				1.7859	0.0000
36				−0.2315	0.0000
37				0.7196	0.0000
38				1.2282	0.0000
39				1.7859	0.0000
40				−0.2315	0.0000
41				−0.6354	0.0000
42				2.3907	0.0000
43				1.7517	0.0000

(Continued)

EXHIBIT 26.3 (Continued)

Row Number	Current Time				
	0	1	2	3	4
44				-0.1757	0.0000
45				-1.5070	0.0000
46				1.7136	0.0000
47				0.5203	0.0000
48				-1.7858	0.0000
49				-0.5013	0.0000
50				2.6883	0.0000
51				1.5066	0.0000
52				-0.7773	0.0000
53				0.4169	0.0000
54				0.9270	0.0000
55				1.4864	0.0000
56				-0.5370	0.0000
57				0.4169	0.0000
58				0.9270	0.0000
59				1.4864	0.0000
60				-0.5370	0.0000
61				-1.5061	0.0000
62				0.7314	0.0000
63				0.2748	0.0000
64				-2.0230	0.0000
Risk-Neutral Value =					2.0047

the fixed rate vary and solve for the fixed rate, which produces zero squared error. The fixed rate that produces a zero value for the swap is 2.3188 percent, a dollar coupon of \$2.3188 instead of \$3. When we drop the new net cash flows into the cash flow table and multiply by the probability-weighted cash flows, we get the confirmation in Exhibit 26.4 that the risk-neutral value is 0.

This change in coupon reduces the value of the fixed rate cash flows to 6.8243, the same as the (unchanged) value of the floating rate cash flows. Alternatively, we could have increased the value of the floating rate cash flows by a present value amount of 2.0047 (generated by adding a fixed amount of $\$3 - 2.3188$ to each floating rate cash flow), while leaving the fixed rate coupon unchanged.

DIGITAL OPTION

Another category is the “digital” category of derivatives that pays either 1 or 0 depending on the level of a random variable. Given our focus on fixed income derivatives, consider the value of a derivative security that pays \$1 at time T_0 if a short-term reference rate (such as LIBOR) is greater than or equal to a critical level K . We demonstrated the valuation of such instruments in the examples of Chapters 6 through 9.

EXHIBIT 26.4 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.1536	−2.3200	−3.7918	0.0000
2		1.6278	0.9046	−3.2479	0.0000
3		1.2216	−0.2901	−2.8522	0.0000
4		0.9578	−2.5991	−5.7186	0.0000
5			0.0692	−0.7852	0.0000
6			1.7205	−0.2725	0.0000
7			1.5067	0.2897	0.0000
8			1.0729	−1.7440	0.0000
9			0.3971	−2.9022	0.0000
10			0.9040	0.3404	0.0000
11			1.4597	−0.8609	0.0000
12			−0.5507	−3.1828	0.0000
13			0.1100	−3.1064	0.0000
14			0.6183	−2.5660	0.0000
15			1.1756	−2.1729	0.0000
16			−0.8404	−5.0208	0.0000
17				−1.8329	0.0000
18				1.3767	0.0000
19				0.1876	0.0000
20				−2.1107	0.0000
21				0.1152	0.0000
22				1.7657	0.0000
23				1.5521	0.0000
24				1.1185	0.0000
25				−0.3043	0.0000
26				2.6922	0.0000
27				2.0595	0.0000
28				0.1510	0.0000
29				0.5424	0.0000
30				1.0486	0.0000
31				1.6035	0.0000
32				−0.4040	0.0000
33				0.0384	0.0000
34				0.5470	0.0000
35				1.1048	0.0000
36				−0.9127	0.0000
37				0.0384	0.0000
38				0.5470	0.0000
39				1.1048	0.0000
40				−0.9127	0.0000
41				−1.3166	0.0000
42				1.7095	0.0000

(Continued)

EXHIBIT 26.4 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
43				1.0705	0.0000
44				−0.8568	0.0000
45				−2.1882	0.0000
46				1.0324	0.0000
47				−0.1608	0.0000
48				−2.4670	0.0000
49				−1.1825	0.0000
50				2.0071	0.0000
51				0.8254	0.0000
52				−1.4585	0.0000
53				−0.2643	0.0000
54				0.2459	0.0000
55				0.8052	0.0000
56				−1.2182	0.0000
57				−0.2643	0.0000
58				0.2459	0.0000
59				0.8052	0.0000
60				−1.2182	0.0000
61				−2.1873	0.0000
62				0.0502	0.0000
63				−0.4064	0.0000
64				−2.7042	0.0000
Risk-Neutral Value =					0.0000

DIGITAL RANGE NOTES

What if the security only pays a principal amount of \$100 and coupon rate of \$1 if the short-term rate is between critical levels K_1 and K_2 ? Sticking with our prior example, let's assume that the short rate reference rate is measured "in arrears" and that the bond has a three-year maturity. We assume the critical levels for the reference rate are 2 percent and 4 percent. The short-term spot rate, measured on an arrears basis, is only between 2 percent and 4 percent in the shaded scenarios given in Exhibit 26.5.

The cash flow on these digital range notes will be a principal payment of \$100 at time $T = 3$ and \$1 only in the shaded scenarios. The total cash flow on the digital range notes is dropped into the cash flow tables and multiplied by the appropriate probability-weighted discount factors to derive the bond value of 97.2304, as shown in Exhibit 26.6.

RANGE FLOATER

Another derivative floating rate structure is the so-called "range floater." Consider a four-period security that pays the reference rate when the rate is less than a critical level K , and 0 otherwise. Cash flow is determined in the normal method, that is, at

EXHIBIT 26.5 Time Period

State	Row Number	0	1	2	3	4
S-1, S-1, S-1	1		2.1653	4.6388	6.1106	0.0000
S-1, S-1, S-2	2		0.6911	1.4142	5.5667	0.0000
S-1, S-1, S-3	3		1.0972	2.6089	5.1710	0.0000
S-1, S-1, S-4	4		1.3611	4.9179	8.0375	0.0000
S-1, S-2, S-1	5			2.2496	3.1041	0.0000
S-1, S-2, S-2	6			0.5984	2.5913	0.0000
S-1, S-2, S-3	7			0.8121	2.0291	0.0000
S-1, S-2, S-4	8			1.2459	4.0628	0.0000
S-1, S-3, S-1	9			1.9217	5.2210	0.0000
S-1, S-3, S-2	10			1.4149	1.9784	0.0000
S-1, S-3, S-3	11			0.8591	3.1798	0.0000
S-1, S-3, S-4	12			2.8695	5.5017	0.0000
S-1, S-4, S-1	13			2.2088	5.4252	0.0000
S-1, S-4, S-2	14			1.7005	4.8848	0.0000
S-1, S-4, S-3	15			1.1432	4.4917	0.0000
S-1, S-4, S-4	16			3.1593	7.3396	0.0000
S-2, S-1, S-1	17				4.1517	0.0000
S-2, S-1, S-2	18				0.9421	0.0000
S-2, S-1, S-3	19				2.1313	0.0000
S-2, S-1, S-4	20				4.4295	0.0000
S-2, S-2, S-1	21				2.2036	0.0000
S-2, S-2, S-2	22				0.5531	0.0000
S-2, S-2, S-3	23				0.7667	0.0000
S-2, S-2, S-4	24				1.2004	0.0000
S-2, S-3, S-1	25				2.6231	0.0000
S-2, S-3, S-2	26				-0.3734	0.0000
S-2, S-3, S-3	27				0.2593	0.0000
S-2, S-3, S-4	28				2.1679	0.0000
S-2, S-4, S-1	29				1.7764	0.0000
S-2, S-4, S-2	30				1.2703	0.0000
S-2, S-4, S-3	31				0.7153	0.0000
S-2, S-4, S-4	32				2.7228	0.0000
S-3, S-1, S-1	33				2.2804	0.0000
S-3, S-1, S-2	34				1.7718	0.0000
S-3, S-1, S-3	35				1.2141	0.0000
S-3, S-1, S-4	36				3.2315	0.0000
S-3, S-2, S-1	37				2.2804	0.0000
S-3, S-2, S-2	38				1.7718	0.0000
S-3, S-2, S-3	39				1.2141	0.0000
S-3, S-2, S-4	40				3.2315	0.0000
S-3, S-3, S-1	41				3.6354	0.0000
S-3, S-3, S-2	42				0.6093	0.0000
S-3, S-3, S-3	43				1.2483	0.0000
S-3, S-3, S-4	44				3.1757	0.0000
S-3, S-4, S-1	45				4.5070	0.0000

(Continued)

EXHIBIT 26.5 (Continued)

State	Row Number	0	1	2	3	4
S-3, S-4, S-2	46				1.2864	0.0000
S-3, S-4, S-3	47				2.4797	0.0000
S-3, S-4, S-4	48				4.7858	0.0000
S-4, S-1, S-1	49				3.5013	0.0000
S-4, S-1, S-2	50				0.3117	0.0000
S-4, S-1, S-3	51				1.4934	0.0000
S-4, S-1, S-4	52				3.7773	0.0000
S-4, S-2, S-1	53				2.5831	0.0000
S-4, S-2, S-2	54				2.0730	0.0000
S-4, S-2, S-3	55				1.5136	0.0000
S-4, S-2, S-4	56				3.5370	0.0000
S-4, S-3, S-1	57				2.5831	0.0000
S-4, S-3, S-2	58				2.0730	0.0000
S-4, S-3, S-3	59				1.5136	0.0000
S-4, S-3, S-4	60				3.5370	0.0000
S-4, S-4, S-1	61				4.5061	0.0000
S-4, S-4, S-2	62				2.2686	0.0000
S-4, S-4, S-3	63				2.7252	0.0000
S-4, S-4, S-4	64				5.0230	0.0000

EXHIBIT 26.6 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	1.0000	0.0000	100.0000	0.0000
2		0.0000	0.0000	100.0000	0.0000
3		0.0000	1.0000	100.0000	0.0000
4		0.0000	0.0000	100.0000	0.0000
5			1.0000	101.0000	0.0000
6			0.0000	101.0000	0.0000
7			0.0000	101.0000	0.0000
8			0.0000	100.0000	0.0000
9			0.0000	100.0000	0.0000
10			0.0000	100.0000	0.0000
11			0.0000	101.0000	0.0000
12			1.0000	100.0000	0.0000
13			1.0000	100.0000	0.0000
14			0.0000	100.0000	0.0000
15			0.0000	100.0000	0.0000
16			1.0000	100.0000	0.0000
17				100.0000	0.0000
18				100.0000	0.0000
19				101.0000	0.0000
20				100.0000	0.0000
21				101.0000	0.0000

EXHIBIT 26.6 (Continued)

Row Number	Current Time				
	0	1	2	3	4
22				100.0000	0.0000
23				100.0000	0.0000
24				100.0000	0.0000
25				101.0000	0.0000
26				100.0000	0.0000
27				100.0000	0.0000
28				101.0000	0.0000
29				100.0000	0.0000
30				100.0000	0.0000
31				100.0000	0.0000
32				101.0000	0.0000
33				101.0000	0.0000
34				100.0000	0.0000
35				100.0000	0.0000
36				101.0000	0.0000
37				101.0000	0.0000
38				100.0000	0.0000
39				100.0000	0.0000
40				101.0000	0.0000
41				101.0000	0.0000
42				100.0000	0.0000
43				100.0000	0.0000
44				101.0000	0.0000
45				100.0000	0.0000
46				100.0000	0.0000
47				101.0000	0.0000
48				100.0000	0.0000
49				101.0000	0.0000
50				100.0000	0.0000
51				100.0000	0.0000
52				101.0000	0.0000
53				101.0000	0.0000
54				101.0000	0.0000
55				100.0000	0.0000
56				101.0000	0.0000
57				101.0000	0.0000
58				101.0000	0.0000
59				100.0000	0.0000
60				101.0000	0.0000
61				100.0000	0.0000
62				101.0000	0.0000
63				101.0000	0.0000
64				100.0000	0.0000
Risk-Neutral Value =					97.2304

time T_0 based on the level of an underlying zero-coupon bond with maturity T_1 , with cash flowing at time $T = 1$. We can use the previous spot rates, adjusted for payment with the normal one-period lag after the interest rate is determined, to calculate the payments on this \$100 four-period note. Let's assume that K is 3 percent and any spot rate value over 3 percent results in a coupon payment of zero. Exhibit 26.7 shows the cash flows on the note at each node of the bushy tree. The risk-neutral value of this structure is 96.5770.

EXHIBIT 26.7 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	0.3003	2.1653	0.0000	100.0000
2		0.3003	2.1653	0.0000	100.0000
3		0.3003	2.1653	0.0000	100.0000
4		0.3003	2.1653	0.0000	100.0000
5			0.6911	1.4142	100.0000
6			0.6911	1.4142	102.5913
7			0.6911	1.4142	102.0291
8			0.6911	1.4142	100.0000
9			1.0972	2.6089	100.0000
10			1.0972	2.6089	101.9784
11			1.0972	2.6089	100.0000
12			1.0972	2.6089	100.0000
13			1.3611	0.0000	100.0000
14			1.3611	0.0000	100.0000
15			1.3611	0.0000	100.0000
16			1.3611	0.0000	100.0000
17				2.2496	100.0000
18				2.2496	100.9421
19				2.2496	102.1313
20				2.2496	100.0000
21				0.5984	102.2036
22				0.5984	100.5531
23				0.5984	100.7887
24				0.5984	101.2004
25				0.8121	102.6231
26				0.8121	99.6266
27				0.8121	100.2593
28				0.8121	102.1679
29				1.2459	101.7764
30				1.2459	101.2703
31				1.2459	100.7153
32				1.2459	102.7228
33				1.9217	102.2804

EXHIBIT 26.7 (Continued)

Row Number	Current Time				
	0	1	2	3	4
34				1.9217	101.7718
35				1.9217	101.2141
36				1.9217	100.0000
37				1.4149	102.2804
38				1.4149	101.7718
39				1.4149	101.2141
40				1.4149	100.0000
41				0.8591	100.0000
42				0.8591	100.6093
43				0.8591	101.2483
44				0.8591	100.0000
45				2.8695	100.0000
46				2.8695	101.2864
47				2.8695	102.4797
48				2.8695	100.0000
49				2.2088	100.0000
50				2.2088	100.3117
51				2.2088	101.4934
52				2.2088	100.0000
53				1.7005	102.5831
54				1.7005	102.0730
55				1.7005	101.5136
56				1.7005	100.0000
57				1.1432	102.5831
58				1.1432	102.0730
59				1.1432	101.5136
60				1.1432	100.0000
61				0.0000	100.0000
62				0.0000	102.2686
63				0.0000	102.7252
64				0.0000	100.0000
Risk-Neutral Value =					96.5770

OTHER DERIVATIVE SECURITIES

As we discussed in Chapter 25, literally any fixed income derivative structure can be analyzed in this framework due to the richness of the HJM framework. Basis swaps, which exchange one floating-rate structure for another; CMT constant maturity swaps, which pay the par coupon bond yield at each point in time versus a more traditional floating-rate index; amortizing interest rate swaps; and many other

structures fit into the same basis structure. The only task that is fairly mechanical is to correctly calculate the cash flows at each node of the HJM bushy tree or at each point in time for N Monte Carlo scenarios in the HJM framework. All of these numerical solutions have related interest rate sensitivity statistics. Hedging portfolios can be derived for each individual transaction or (more efficiently from an execution point of view) for the aggregated interest rate risk position of the organization. Shifts in the zero-coupon bond prices input at any of the maturities (in our example, that is at times 1, 2, 3, and 4) produce deltas in value with respect to that particular point on the yield curve. Shifts in interest rate volatility assumptions produce the multifactor vega equivalents of vega in the one-factor Black-Scholes options model. The class of securities that we have not yet valued is the class that involves an American option. We turn to that problem in the next chapter. Before doing so, however, we turn as usual to the impact of credit risk on the previous valuation framework.

CREDIT RISK AND EXOTIC DERIVATIVES STRUCTURES

For all of the instruments we have analyzed in Chapters 19 through 26, we face a common dilemma. If we assume that the default probability of our counterparty is independent of the factors that drive payoffs on the instrument, then we make the same (potentially huge) mistake that JPMorgan Chase made in the famous 1998 incident with SK Securities, which is discussed in depth by van Deventer and Imai (2003). This mistake was repeated by a wave of financial institutions in the 2006–2011 credit crisis when they ignored home price risk and its correlated impact on the value of RMBS, CDOs, and credit default swaps of both individual firms and individual securities. This “default risk independence” assumption is very seductive because as Jarrow and Turnbull show, for many structures the impact of credit risk on the value of a structure just becomes a simple multiplier of the original valuation formula. As dangerous as this simplification can be, it is a better approximation than doing nothing. And “doing nothing,” it’s sad to say, was one of the most common risk management techniques prior to the recent credit crisis.

As is the case with the other structures discussed in the last few chapters, the authors believe that such an “independence” assumption is particularly dangerous in the case of derivative instrument credit risk (even in the presence of AAA-rated special purpose vehicles) for two reasons. First, interest rates clearly have an impact, and potentially a substantial impact, on the default probabilities of financial institutions as we have learned time and time again in incidents ranging from the savings and loan crisis to the Asia crisis. This increased default risk comes not necessarily from interest rate mismatches, but from the increase in the default probabilities of almost every counterparty on the balance sheet of the financial institution. Second, there is a tremendous concentration of market share in the derivatives business, and the default of any one of those institutions has a double impact, as we saw with the failures of Lehman Brothers and Bear Stearns in 2008 and the effective failures of firms like AIG and Merrill Lynch. Just one default would adversely affect 10 to 20 percent of the derivatives transactions on the books of an institution (since market shares of the major derivatives players are typically in the 10 to 20 percent range) and it would cause an important market disruption. It would require a huge wave of derivatives to be analyzed and replaced with another counterparty, one whose

default probability is most likely highly correlated with the institution that has defaulted. One bankruptcy lawyer told us that “the few remaining staff of Lehman Brothers are being less than helpful in calculating the net exposure my client has with Lehman.” This macro factor risk is much broader than the impact of interest rates alone. Home prices (remember Japan and 2006–2011 in the United States and Europe), commercial real estate, exchange rates, oil prices, and a very wide array of other factors can have similar massive correlated impacts on risk.

For that reason, we think the valuation and simulation of interest rate exotic structures have to be analyzed using the same process outlined beginning in Chapters 19 and 20:

1. Default probabilities of all counterparties should be explicitly linked to key macroeconomic factors.
2. Interest rate and other macro factor simulation will impact both the payoffs due on the structure and the default probabilities and zero-coupon bond prices of each risky counterparty.
3. The true credit-adjusted value and cash flow can then be accurately derived.

Any other strategy is the financial equivalent of sticking your head in the sand and ignoring the problem, because you don’t want to face the reality of the problem. Fortunately for shareholders, depositors and deposit insurance funds of various nations, regulators from the Basel Committee on Banking Supervision, the Financial Services Authority in the United Kingdom, and the array of banking regulators in the United States are focusing on this issue with the intensity it deserves. We do the same in Chapters 36 through 41.

American Fixed Income Options

In Chapters 21 through 26, we have emphasized the valuation of securities where the option embedded in the security was exercisable at only one date. These European options, as we have shown, generally have explicit analytical solutions in the Vasicek family of term structure models. As we saw in Chapter 3, however, one-factor interest rate models are very inaccurate models of interest rate movements. In this chapter, we continue to emphasize numerical solutions using the example of the three-factor Heath, Jarrow, and Morton (HJM) model from Chapter 9. In previous chapters, we have frequently remarked that an HJM Monte Carlo simulation was the preferred valuation methodology, with an HJM bushy tree used only for expositional purposes. In this chapter, the situation is reversed. Using the bushy tree approach is the most accurate technique for valuing American fixed income options. Monte Carlo simulation, even using the HJM approach, is one notch lower in accuracy for reasons we explain. Please note, however, that, even in the case of a one-factor term structure model, numerical methods are necessary for the accurate valuation of American fixed income options. This is even truer when the issuer of the security has a positive probability of default. The interaction of the security's payoffs with the credit risk of the issuer has interesting and important implications.

The major advances in computer science since the first volume of this book was published have dramatically relaxed the computational constraints that formerly motivated the use of closed form solutions instead of highly accurate multifactor simulations or bushy trees. The increasing use of 64-bit operating systems, with their huge expansion in available memory, makes the techniques in this chapter available at the transaction level for the entire enterprise.

We note also that the Jarrow-Merton put option, the key integrated risk measure that we focus on in this volume, can be thought of as an American option. The bank deposit insurance provided by the Federal Deposit Insurance Corporation in the United States is effectively an American put option on the value of deposits provided to depositors and paid for by the banks themselves (if priced correctly) and taxpayers (if priced incorrectly). Bank deposit insurance is an American option, because the bank can default at any time, not just one point in time like a European option.

Much of the early work on lattice or bushy tree valuations of American options was done using the popular Hull and White (1990b, 1993b, 1994a) trinomial lattice. The lattice was initially developed for one-factor term structure models and then expanded to the two-factor case. As we saw in Chapter 3, between 5 and 10 factors are needed to accurately model even a relatively transparent yield curve like that for

U.S. Treasury securities. A one- or two-factor model is simply not a good approximation to actual yield curve movements, and a lattice based on these assumptions will not capture some of the yield curve movements (e.g., yield curve twists) that can have a big impact on the exercise of an American option, such as the prepayment option on a mortgage in the United States and Japan.

The reason for devoting the bulk of this and the two following chapters to American fixed income options is that almost all financial institutions have the majority of their balance sheets devoted to American fixed income options, explicitly or implicitly. A brief list of examples shows how critical this topic is. Typical American options on financial institution balance sheets include:

- The right to terminate a life insurance policy in return for receipt of its surrender value
- The right to resign as the customer of an investment management firm
- The right to prepay a mortgage loan
- The right to withdraw as a bank from the consumer deposit gathering business
- The right to withdraw almost any consumer bank deposit at any time either at par or upon the payment of an early withdrawal penalty
- The right of a corporate borrower to declare bankruptcy
- The right of a financial institution to pay dividends
- The right to exercise a foreign exchange option
- The right to exercise a standard swaption contract

The list of examples is almost endless. Needless to say, the topic at hand is a critical one.

AN OVERVIEW OF NUMERICAL TECHNIQUES FOR FIXED INCOME OPTION VALUATION

Financial market participants generally use one of six approaches to the valuation of American fixed income options:

- Analytical approximations
- Monte Carlo simulation
- Bushy trees
- Finite difference methods
- Binomial lattices (a special case of the bushy tree, because the tree “recombines”)
- Trinomial lattices (another special case of bushy tree)

Analytical approximations are not sufficiently accurate for us to devote time to here. We will discuss the alternatives to the bushy tree after first using the three-factor Heath, Jarrow, and Morton framework from Chapter 9 on two examples of fixed income options, a callable fixed rate bond and a prepayable amortizing loan. There is a wide variety of American fixed income options in everyday finance, but loans and bonds that are callable or prepayable (the terms have the equivalent economic meaning) are particularly numerous.

AN EXAMPLE OF VALUATION OF A CALLABLE BOND WITH A THREE-FACTOR HJM BUSHY TREE

In Chapter 9, we analyzed the value of a bond with \$100 in principal outstanding that pays 3 percent interest annually and matures at time $T = 4$. The cash flows on the bond and its valuation in Chapter 9, 104.7071, are reproduced here in Exhibit 27.1. The valuation is calculated by multiplying the cash flows in the cash flow table times the corresponding probability-weighted discount factors that were displayed in Chapter 9.

In Chapter 14, we assumed that we could observe prices of callable fixed rate coupon bonds, where the call option was European in nature: the bonds could be prepaid or “called” at only one point in time. In this example, we assume that the bonds could be called at any time. In the context of our four-step bushy tree from Chapter 9, “any time” effectively means the bonds can be called at time 0, time 1, time 2, or time 3, prior to maturity at time 4. If we need to evaluate the call option’s

EXHIBIT 27.1 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		3.0000	3.0000	3.0000	103.0000
3		3.0000	3.0000	3.0000	103.0000
4		3.0000	3.0000	3.0000	103.0000
5			3.0000	3.0000	103.0000
6			3.0000	3.0000	103.0000
7			3.0000	3.0000	103.0000
8			3.0000	3.0000	103.0000
9			3.0000	3.0000	103.0000
10			3.0000	3.0000	103.0000
11			3.0000	3.0000	103.0000
12			3.0000	3.0000	103.0000
13			3.0000	3.0000	103.0000
14			3.0000	3.0000	103.0000
15			3.0000	3.0000	103.0000
16			3.0000	3.0000	103.0000
17				3.0000	103.0000
18				3.0000	103.0000
19				3.0000	103.0000
20				3.0000	103.0000
21				3.0000	103.0000
22				3.0000	103.0000
23				3.0000	103.0000
24				3.0000	103.0000

EXHIBIT 27.1 (Continued)

Row Number	Current Time				
	0	1	2	3	4
25				3.0000	103.0000
26				3.0000	103.0000
27				3.0000	103.0000
28				3.0000	103.0000
29				3.0000	103.0000
30				3.0000	103.0000
31				3.0000	103.0000
32				3.0000	103.0000
33				3.0000	103.0000
34				3.0000	103.0000
35				3.0000	103.0000
36				3.0000	103.0000
37				3.0000	103.0000
38				3.0000	103.0000
39				3.0000	103.0000
40				3.0000	103.0000
41				3.0000	103.0000
42				3.0000	103.0000
43				3.0000	103.0000
44				3.0000	103.0000
45				3.0000	103.0000
46				3.0000	103.0000
47				3.0000	103.0000
48				3.0000	103.0000
49				3.0000	103.0000
50				3.0000	103.0000
51				3.0000	103.0000
52				3.0000	103.0000
53				3.0000	103.0000
54				3.0000	103.0000
55				3.0000	103.0000
56				3.0000	103.0000
57				3.0000	103.0000
58				3.0000	103.0000
59				3.0000	103.0000
60				3.0000	103.0000
61				3.0000	103.0000
62				3.0000	103.0000
63				3.0000	103.0000
64				3.0000	103.0000
Risk-Neutral Value =					104.7071

impact more granularly, we would simply increase the number and shorten the length of the time steps in the bushy tree. Please note that a *Bermudan option* is an option that can be exercised at many points in time but not at any time. In the context of this example, one could imagine that a bond could be called at time 2 or time 3 but not at time 0 or time 1. This is a special case of a Bermudan option. The bond would be described as “four years, noncall two.” We assume that the call option can be exercised at any of the nodes of the lattice except the maturity date (because calling at that time is identical with normal principal and interest payments at maturity). Bermudan option analysis is identical except that the option to call or prepay is turned off at some points in time.

How would a rational corporate treasurer or bond investor value the probability that this 3 percent, four-year bond would be prepaid? (We use “called” when discussing bonds, “prepaid” when discussing loans.) These probabilities need to be assessed in order to appropriately populate the cash flow table for HJM valuation. The value of the bond at time zero, and the decision about whether to call the bond at time zero, depends on what the bond issuer would do at the four nodes at time $T = 1$. At each of the four nodes at $T = 1$, there are a set of four more nodes at time $T = 2$, and the time 1 behavior depends on what the bond issuer would do at time 2. Finally, at each of the 16 nodes at time $T = 2$, there are four nodes at time $T = 3$ and the time 2 behavior depends on the actual decisions that the bond issuer would make at time $T = 3$, where there are 64 nodes. This is a dynamic programming problem that we need to analyze in a recursive fashion. We start at the end of the bond’s life and set our call strategy for the 64 nodes at time $T = 3$. Having done that, we step backward in time to the 16 nodes at time $T = 2$. Because we know what the time $T = 3$ actions are, we can deduce what to do on the 16 nodes at time $T = 2$. We then repeat the process and step backward in time to time $T = 1$ (four nodes) and time $T = 0$ (one node). This gives us the cash flows at each node. We insert these cash flows in the HJM cash flow table, multiply by the appropriate probability-weighted discount factors, and get the risk-neutral value. We do that now.

First we analyze whether or not to call the bonds at time $T = 3$ at each of the 64 nodes. We assume the decision is made immediately after time $T = 3$ interest of \$3 is paid. The analysis is simple. We multiple the zero-coupon bond price for a bond maturing at $T = 4$ as of time 3 for that state s_t times the cash flow of \$103 to get the net present value of the no-call strategy, $P(3, 4, s_t)103$. If this present value is more than 100, we call the bonds now and pay \$100 in principal at time $T = 3$ in addition to the \$3 in interest paid one second earlier. If the present value is \$100 or less, we leave the bonds outstanding and pay \$103 at time $T = 4$. Exhibit 27.2 analyzes the call or no-call decision at each node at time $T = 3$.

Based on these rational call strategies, the cash flows on a bond that could be called only at time $T = 3$ are given in Exhibit 27.3. The value of this bond, a “four-year noncall 3” bond, would be 104.2480, less than the noncall version of the bond.

Having done the call analysis for time $T = 3$, we now step backward in time to the 16 nodes at time $T = 2$. At each of those 16 nodes, we look at the 64 nodes at time $T = 3$ and study the net present value of the no-call strategy as of time $T = 2$. The net present value as of time $T = 2$ is a function of the discount factor, shift probability, and time $T = 3$ value if the bonds are not called at time $T = 2$. Recall the shift probabilities from Chapter 9, which we reproduce here as Exhibit 27.4.

EXHIBIT 27.2 Maturity of Cash Flow Received

Row Number	Current Time				Action?
	0	1	2	3	
1	0.0000	0.0000	0.0000	97.0685	No Call
2		0.0000	0.0000	97.5686	No Call
3		0.0000	0.0000	97.9357	No Call
4		0.0000	0.0000	95.3373	No Call
5			0.0000	99.8991	No Call
6			0.0000	100.3984	Call
7			0.0000	100.9516	Call
8			0.0000	98.9787	No Call
9			0.0000	97.8892	No Call
10			0.0000	101.0018	Call
11			0.0000	99.8258	No Call
12			0.0000	97.6288	No Call
13			0.0000	97.6996	No Call
14			0.0000	98.2030	No Call
15			0.0000	98.5724	No Call
16			0.0000	95.9571	No Call
17				98.8942	No Call
18				102.0387	Call
19				100.8506	Call
20				98.6311	No Call
21				100.7792	Call
22				102.4334	Call
23				102.2163	Call
24				101.7783	Call
25				100.3673	Call
26				103.3860	Call
27				102.7336	Call
28				100.8145	Call
29				101.2022	Call
30				101.7080	Call
31				102.2685	Call
32				100.2698	Call
33				100.7035	Call
34				101.2068	Call
35				101.7645	Call
36				99.7757	No Call
37				100.7035	Call
38				101.2068	Call
39				101.7645	Call
40				99.7757	No Call
41				99.3869	No Call
42				102.3762	Call

(Continued)

EXHIBIT 27.2 *(Continued)*

Row Number	Current Time				Action?
	0	1	2	3	
43				101.7301	Call
44				99.8297	No Call
45				98.5580	No Call
46				101.6918	Call
47				100.5077	Call
48				98.2958	No Call
49				99.5157	No Call
50				102.6800	Call
51				101.4844	Call
52				99.2509	No Call
53				100.4064	Call
54				100.9082	Call
55				101.4642	Call
56				99.4813	No Call
57				100.4064	Call
58				100.9082	Call
59				101.4642	Call
60				99.4813	No Call
61				98.5589	No Call
62				100.7152	Call
63				100.2675	Call
64				98.0738	No Call

EXHIBIT 27.3 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		3.0000	3.0000	3.0000	103.0000
3		3.0000	3.0000	3.0000	103.0000
4		3.0000	3.0000	3.0000	103.0000
5			3.0000	3.0000	103.0000
6			3.0000	103.0000	0.0000
7			3.0000	103.0000	0.0000
8			3.0000	3.0000	103.0000
9			3.0000	3.0000	103.0000
10			3.0000	103.0000	0.0000
11			3.0000	3.0000	103.0000
12			3.0000	3.0000	103.0000
13			3.0000	3.0000	103.0000
14			3.0000	3.0000	103.0000

EXHIBIT 27.3 (Continued)

Row Number	Current Time				
	0	1	2	3	4
15			3.0000	3.0000	103.0000
16			3.0000	3.0000	103.0000
17				3.0000	103.0000
18				103.0000	0.0000
19				103.0000	0.0000
20				3.0000	103.0000
21				103.0000	0.0000
22				103.0000	0.0000
23				103.0000	0.0000
24				103.0000	0.0000
25				103.0000	0.0000
26				103.0000	0.0000
27				103.0000	0.0000
28				103.0000	0.0000
29				103.0000	0.0000
30				103.0000	0.0000
31				103.0000	0.0000
32				103.0000	0.0000
33				103.0000	0.0000
34				103.0000	0.0000
35				103.0000	0.0000
36				3.0000	103.0000
37				103.0000	0.0000
38				103.0000	0.0000
39				103.0000	0.0000
40				3.0000	103.0000
41				3.0000	103.0000
42				103.0000	0.0000
43				103.0000	0.0000
44				3.0000	103.0000
45				3.0000	103.0000
46				103.0000	0.0000
47				103.0000	0.0000
48				3.0000	103.0000
49				3.0000	103.0000
50				103.0000	0.0000
51				103.0000	0.0000
52				3.0000	103.0000
53				103.0000	0.0000
54				103.0000	0.0000
55				103.0000	0.0000

(Continued)

EXHIBIT 27.3 (Continued)

Row Number	Current Time				
	0	1	2	3	4
56				3.0000	103.0000
57				103.0000	0.0000
58				103.0000	0.0000
59				103.0000	0.0000
60				3.0000	103.0000
61				3.0000	103.0000
62				103.0000	0.0000
63				103.0000	0.0000
64				3.0000	103.0000
Risk-Neutral Value =					104.2480

EXHIBIT 27.4 Shift Probabilities

Scenario	Row Number	Current Time			
		0	1	2	3
S-1, S-1, S-1	1	100%	12.50%	12.50%	12.50%
S-1, S-1, S-2	2		12.50%	12.50%	12.50%
S-1, S-1, S-3	3		25.00%	25.00%	25.00%
S-1, S-1, S-4	4		50.00%	50.00%	50.00%
S-1, S-2, S-1	5			12.50%	12.50%
S-1, S-2, S-2	6			12.50%	12.50%
S-1, S-2, S-3	7			25.00%	25.00%
S-1, S-2, S-4	8			50.00%	50.00%
S-1, S-3, S-1	9			12.50%	12.50%
S-1, S-3, S-2	10			12.50%	12.50%
S-1, S-3, S-3	11			25.00%	25.00%
S-1, S-3, S-4	12			50.00%	50.00%
S-1, S-4, S-1	13			12.50%	12.50%
S-1, S-4, S-2	14			12.50%	12.50%
S-1, S-4, S-3	15			25.00%	25.00%
S-1, S-4, S-4	16			50.00%	50.00%
S-2, S-1, S-1	17				12.50%
S-2, S-1, S-2	18				12.50%
S-2, S-1, S-3	19				25.00%
S-2, S-1, S-4	20				50.00%
S-2, S-2, S-1	21				12.50%
S-2, S-2, S-2	22				12.50%
S-2, S-2, S-3	23				25.00%

EXHIBIT 27.4 (Continued)

Scenario	Row Number	Current Time			
		0	1	2	3
S-2, S-2, S-4	24				50.00%
S-2, S-3, S-1	25				12.50%
S-2, S-3, S-2	26				12.50%
S-2, S-3, S-3	27				25.00%
S-2, S-3, S-4	28				50.00%
S-2, S-4, S-1	29				12.50%
S-2, S-4, S-2	30				12.50%
S-2, S-4, S-3	31				25.00%
S-2, S-4, S-4	32				50.00%
S-3, S-1, S-1	33				12.50%
S-3, S-1, S-2	34				12.50%
S-3, S-1, S-3	35				25.00%
S-3, S-1, S-4	36				50.00%
S-3, S-2, S-1	37				12.50%
S-3, S-2, S-2	38				12.50%
S-3, S-2, S-3	39				25.00%
S-3, S-2, S-4	40				50.00%
S-3, S-3, S-1	41				12.50%
S-3, S-3, S-2	42				12.50%
S-3, S-3, S-3	43				25.00%
S-3, S-3, S-4	44				50.00%
S-3, S-4, S-1	45				12.50%
S-3, S-4, S-2	46				12.50%
S-3, S-4, S-3	47				25.00%
S-3, S-4, S-4	48				50.00%
S-4, S-1, S-1	49				12.50%
S-4, S-1, S-2	50				12.50%
S-4, S-1, S-3	51				25.00%
S-4, S-1, S-4	52				50.00%
S-4, S-2, S-1	53				12.50%
S-4, S-2, S-2	54				12.50%
S-4, S-2, S-3	55				25.00%
S-4, S-2, S-4	56				50.00%
S-4, S-3, S-1	57				12.50%
S-4, S-3, S-2	58				12.50%
S-4, S-3, S-3	59				25.00%
S-4, S-3, S-4	60				50.00%
S-4, S-4, S-1	61				12.50%
S-4, S-4, S-2	62				12.50%
S-4, S-4, S-3	63				25.00%
S-4, S-4, S-4	64				50.00%

The discount factors that are relevant are the one-period, zero-coupon bond prices as of time $T = 2$ for maturity at $T = 3$ for each of the 16 nodes or states, $P(2, 3, s_t)$. Those discount factors are shown in column 2 of Exhibit 27.5, which we derived in Chapter 9.

If the bonds are called at time $T = 3$, the time $T = 3$ value is the payoff of the bonds (\$100) plus the time $T = 3$ interest payment (\$3), for a total of \$103. If the bonds are *not* called at time $T = 3$, the time $T = 3$ value is the sum of the \$3 interest payment made at time $T = 3$ plus the time $T = 3$ risk-neutral value of the final payment of \$103. We gave these risk-neutral values as of time $T = 3$ in Exhibit 27.8.

For each of the 16 nodes at time $T = 2$, we now calculate the risk-neutral value of the no-call strategy at time $T = 2$. This analysis is identical to the valuation of a one-period, zero-coupon bond that could have alternative payoffs, depending on the state. Exhibit 27.6 shows the call/no-call analysis for each of the 16 nodes. If the net present value of not calling at time $T = 2$ is greater than 100, the bonds will be called. If the net present value of not calling at time $T = 2$ is less than or equal to 100, the bonds will not be called.

Take the first node at time $T = 2$ in Exhibit 27.6. Using the evolution of states from Chapter 9, the first node is a Shift 1 state to get to time $T = 1$ followed by another Shift 1 to get to time $T = 2$. Each of the first four rows in the columns labeled “probability” and “Time 3 value” represent the shift values and Time 3 values (derived as above) for Shifts 1, 2, 3, and 4 to reach time $T = 3$. We multiply the probability of each Time 3 value by the probability of the state and then discount by the appropriate risk-neutral discount factor (in the state of node 1 at time $T = 2$, this discount factor is 0.9557) to get the time 2 value of 95.0720. Since this time 2 value is 100 or less, the decision at node 1 for time $T = 2$ is “no call.” The same analysis is repeated for all 16 nodes.

EXHIBIT 27.5 Three-Year, Zero-Coupon Bond Price

Row Number	Current Time		
	0	1	2
1	0.9619	0.9425	0.9557
2		0.9816	0.9861
3		0.9692	0.9746
4		0.9640	0.9531
5			0.9780
6			0.9941
7			0.9919
8			0.9877
9			0.9811
10			0.9860
11			0.9915
12			0.9721
13			0.9784
14			0.9833
15			0.9887
16			0.9694

EXHIBIT 27.6 Call/No-Call Analysis

Time 2 Node	Time 2 Value	Action?	Discount Factor	Time 3 Node	Probability	Time 3 Value
1	95.0720	No Call	0.9557	1	0.1250	100.0685
			0.9557	2	0.1250	100.5686
			0.9557	3	0.2500	100.9357
			0.9557	4	0.5000	98.3373
2	101.0477	Call	0.9861	5	0.1250	102.8991
			0.9861	6	0.1250	103.0000
			0.9861	7	0.2500	103.0000
			0.9861	8	0.5000	101.9787
3	98.9261	No Call	0.9746	9	0.1250	100.8892
			0.9746	10	0.1250	103.0000
			0.9746	11	0.2500	102.8258
			0.9746	12	0.5000	100.6288
4	95.4169	No Call	0.9531	13	0.1250	100.6996
			0.9531	14	0.1250	101.2030
			0.9531	15	0.2500	101.5724
			0.9531	16	0.5000	98.9571
5	99.9293	No Call	0.9780	17	0.1250	101.8942
			0.9780	18	0.1250	103.0000
			0.9780	19	0.2500	103.0000
			0.9780	20	0.5000	101.6311
6	102.3874	Call	0.9941	21	0.1250	103.0000
			0.9941	22	0.1250	103.0000
			0.9941	23	0.2500	103.0000
			0.9941	24	0.5000	103.0000
7	102.1703	Call	0.9919	25	0.1250	103.0000
			0.9919	26	0.1250	103.0000
			0.9919	27	0.2500	103.0000
			0.9919	28	0.5000	103.0000
8	101.7325	Call	0.9877	29	0.1250	103.0000
			0.9877	30	0.1250	103.0000
			0.9877	31	0.2500	103.0000
			0.9877	32	0.5000	103.0000
9	100.9479	Call	0.9811	33	0.1250	103.0000
			0.9811	34	0.1250	103.0000
			0.9811	35	0.2500	103.0000
			0.9811	36	0.5000	102.7757
10	101.4524	Call	0.9860	37	0.1250	103.0000
			0.9860	38	0.1250	103.0000
			0.9860	39	0.2500	103.0000
			0.9860	40	0.5000	102.7757
11	101.9623	Call	0.9915	41	0.1250	102.3869
			0.9915	42	0.1250	103.0000

(Continued)

EXHIBIT 27.6 (Continued)

Time 2 Node	Time 2 Value	Action?	Discount Factor	Time 3 Node	Probability	Time 3 Value
12	99.1233	No Call	0.9915	43	0.2500	103.0000
			0.9915	44	0.5000	102.8297
			0.9721	45	0.1250	101.5580
			0.9721	46	0.1250	103.0000
13	100.3484	Call	0.9721	47	0.2500	103.0000
			0.9721	48	0.5000	101.2958
			0.9784	49	0.1250	102.5157
			0.9784	50	0.1250	103.0000
14	101.0227	Call	0.9784	51	0.2500	103.0000
			0.9784	52	0.5000	102.2509
			0.9833	53	0.1250	103.0000
			0.9833	54	0.1250	103.0000
15	101.5794	Call	0.9833	55	0.2500	103.0000
			0.9833	56	0.5000	102.4813
			0.9887	57	0.1250	103.0000
			0.9887	58	0.1250	103.0000
16	98.7374	No Call	0.9887	59	0.2500	103.0000
			0.9887	60	0.5000	102.4813
			0.9694	61	0.1250	101.5589
			0.9694	62	0.1250	103.0000
			0.9694	63	0.2500	103.0000
			0.9694	64	0.5000	101.0738

We now revise the cash flow table to recognize rational call behavior at either time $T = 2$ or time $T = 3$. If the bonds are called at node k at time $T = 2$, all states that stem from node k at times $T = 3$ and $T = 4$ will have cash flows of zero. The cash flow at node k at time $T = 3$ will have a cash flow equal to the sum of the \$3 interest plus the repayment of the \$100 principal, for a total of \$103. The revised cash flow table is shown in Exhibit 27.7. When these cash flows are multiplied times the appropriate probability-weighted discount factors, we get a risk-neutral value for this “four-year noncall 2” bond equal to 103.5473.

We now step back one more time step to determine our optimal call/no-call strategy from the perspective of the four nodes at time $T = 1$. The analysis is identical to the exercise that we went through for the 16 nodes at time $T = 2$. The probabilities for Shifts 1, 2, 3, and 4 are unchanged. The discount factors for a one-period, zero-coupon bond as of time $T = 1$ to mature at time $T = 2$ depends on the state, $P(1, 2, s_t)$. The values for these one-period, zero-coupon bonds were given in Chapter 9. We reproduce them here in Exhibit 27.8.

We now calculate the probability-weighted Time 1 values of the possible outcomes of the time $T = 2$ nodes. Those calculations parallel the calculation of the Time 2 values. If the time 1 risk-neutral value of no call is greater than 100, we call the bonds. If the risk-neutral value of no call is 100 or less, the decision is no call. Exhibit 27.9 shows the call/no-call strategy at time $T = 1$.

EXHIBIT 27.7 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		3.0000	103.0000	3.0000	103.0000
3		3.0000	3.0000	3.0000	103.0000
4		3.0000	3.0000	3.0000	103.0000
5			3.0000	0.0000	0.0000
6			103.0000	0.0000	0.0000
7			103.0000	0.0000	0.0000
8			103.0000	0.0000	0.0000
9			103.0000	3.0000	103.0000
10			103.0000	103.0000	0.0000
11			103.0000	3.0000	103.0000
12			3.0000	3.0000	103.0000
13			103.0000	3.0000	103.0000
14			103.0000	3.0000	103.0000
15			103.0000	3.0000	103.0000
16			3.0000	3.0000	103.0000
17				3.0000	103.0000
18				103.0000	0.0000
19				103.0000	0.0000
20				3.0000	103.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0000	0.0000

(Continued)

EXHIBIT 27.7 (Continued)

Row Number	Current Time				
	0	1	2	3	4
43				0.0000	0.0000
44				0.0000	0.0000
45				3.0000	103.0000
46				103.0000	0.0000
47				103.0000	0.0000
48				3.0000	103.0000
49				0.0000	0.0000
50				0.0000	0.0000
51				0.0000	0.0000
52				0.0000	0.0000
53				0.0000	0.0000
54				0.0000	0.0000
55				0.0000	0.0000
56				0.0000	0.0000
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				3.0000	103.0000
62				103.0000	0.0000
63				103.0000	0.0000
64				3.0000	103.0000
Risk-Neutral Value =					103.5473

EXHIBIT 27.8 Two-Year, Zero-Coupon Bond Price

Row Number	Current Time		
	0	1	2
1	0.9841	0.9788	1
2		0.9931	1
3		0.9891	1
4		0.9866	1

The results show that the bonds are called in three out of the four nodes at time $T = 1$. We now revise our call-adjusted cash flow table so that there are no cash flows on any branches of the bushy tree that follow a “call” decision at time $T = 1$. The results of this adjustment are shown in Exhibit 27.10. The risk-neutral valuation for this “four-year noncall 1” bond is 102.4060.

Now, what happens if we step back one more period and allow the bond to be called at time 0 (which would be the case if the bond had been issued some periods before)? Of

EXHIBIT 27.9 Call/No-Call Strategy at $T = 1$

Time 1 Node	Time 1 Value	Action?	Discount Factor	Time 2 Node	Probability	Time 2 Value
1	97.7083	No Call	0.9788	1	0.1250	98.0720
			0.9788	2	0.1250	103.0000
			0.9788	3	0.2500	101.9261
			0.9788	4	0.5000	98.4169
2	102.2843	Call	0.9931	5	0.1250	102.9293
			0.9931	6	0.1250	103.0000
			0.9931	7	0.2500	103.0000
			0.9931	8	0.5000	103.0000
3	101.4485	Call	0.9891	9	0.1250	103.0000
			0.9891	10	0.1250	103.0000
			0.9891	11	0.2500	103.0000
			0.9891	12	0.5000	102.1233
4	100.9941	Call	0.9866	13	0.1250	103.0000
			0.9866	14	0.1250	103.0000
			0.9866	15	0.2500	103.0000
			0.9866	16	0.5000	101.7374

EXHIBIT 27.10 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		103.0000	103.0000	3.0000	103.0000
3		103.0000	3.0000	3.0000	103.0000
4		103.0000	3.0000	3.0000	103.0000
5			0.0000	0.0000	0.0000
6			0.0000	0.0000	0.0000
7			0.0000	0.0000	0.0000
8			0.0000	0.0000	0.0000
9			0.0000	3.0000	103.0000
10			0.0000	103.0000	0.0000
11			0.0000	3.0000	103.0000
12			0.0000	3.0000	103.0000
13			0.0000	3.0000	103.0000
14			0.0000	3.0000	103.0000
15			0.0000	3.0000	103.0000
16			0.0000	3.0000	103.0000
17				0.0000	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000

(Continued)

EXHIBIT 27.10 (Continued)

Row Number	Current Time				
	0	1	2	3	4
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				0.0000	0.0000
46				0.0000	0.0000
47				0.0000	0.0000
48				0.0000	0.0000
49				0.0000	0.0000
50				0.0000	0.0000
51				0.0000	0.0000
52				0.0000	0.0000
53				0.0000	0.0000
54				0.0000	0.0000
55				0.0000	0.0000
56				0.0000	0.0000
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				0.0000	0.0000
62				0.0000	0.0000
63				0.0000	0.0000
64				0.0000	0.0000
Risk-Neutral Value =					102.4060

course, the answer is obvious; since the risk-neutral value if not called is 102.4060, it is cheaper to pay the bonds off now at 100. Using the format of our prior analysis, we replicate this pricing and reach the call conclusion as shown in Exhibit 27.11.

The cash flow table for this strategy is not very interesting. At time zero, the bonds pay 100 and all other cash flows are zero. Risk-neutral value is, of course, 100. While this may not be visually interesting, it is analytically interesting because bonds get called all the time and we have just demonstrated the analysis one would go through to make an accurate decision in this regard using the three-factor HJM model of Chapter 9.

WHAT IS THE PAR COUPON ON A CALLABLE BOND?

Given that a 3 percent callable bond is “in the money” on the option to call the bond, one might ask, “What coupon level produces a mark-to-market value of par on the assumption that the bonds are a new issue and will not be called at time zero?” The answer can be obtained by using the solver function in common spreadsheet software to iterate the coupon and option-adjusted valuation in the HJM framework. The option-adjusted par coupon is 1.8647 percent. The cash flows that result from the optimal call strategy are given in Exhibit 27.12. All nodes that show a cash flow equal to 101.8647 represent a call decision. All nodes that show a cash flow equal to 1.8647 represent a no-call decision.

This calculation has an enormous set of applications above and beyond the pricing of new issue bonds or swaptions at fair value. It is used for transfer pricing and performance measurement for assets or liabilities which themselves have an embedded call option. All of the real-world instruments listed in the introduction to this chapter require a similar kind of American fixed income options analysis. Because of the power of the HJM framework, our probability-weighted discount factors do not change just because the instrument involves an American fixed income option. Only the cash flows used to populate the cash flow table change.

AN EXAMPLE OF VALUATION OF A RATIONALLY PREPAID AMORTIZING LOAN

By transaction count, a huge proportion of the assets held by major financial institutions around the world that involve an American fixed income option are consumer

EXHIBIT 27.11 Call Conclusion

Time 0 Node	Time 0 Value	Action?	Discount Factor	Time 1 Node	Probability	Time 1 Value
1	102.4060	Call	0.9970	1	0.1250	100.7083
			0.9970	2	0.1250	103.0000
			0.9970	3	0.2500	103.0000
			0.9970	4	0.5000	103.0000

EXHIBIT 27.12 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	1.8647	1.8647	1.8647	101.8647
2		101.8647	1.8647	1.8647	101.8647
3		1.8647	1.8647	1.8647	101.8647
4		1.8647	1.8647	1.8647	101.8647
5			0.0000	1.8647	101.8647
6			0.0000	1.8647	101.8647
7			0.0000	1.8647	101.8647
8			0.0000	1.8647	101.8647
9			1.8647	1.8647	101.8647
10			1.8647	1.8647	101.8647
11			101.8647	1.8647	101.8647
12			1.8647	1.8647	101.8647
13			1.8647	1.8647	101.8647
14			1.8647	1.8647	101.8647
15			1.8647	1.8647	101.8647
16			1.8647	1.8647	101.8647
17				0.0000	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				1.8647	101.8647
34				101.8647	0.0000
35				101.8647	0.0000
36				1.8647	101.8647
37				1.8647	101.8647
38				101.8647	0.0000
39				101.8647	0.0000
40				1.8647	101.8647
41				0.0000	0.0000

EXHIBIT 27.12 (Continued)

Row Number	Current Time				
	0	1	2	3	4
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				1.8647	101.8647
46				101.8647	0.0000
47				1.8647	101.8647
48				1.8647	101.8647
49				1.8647	101.8647
50				101.8647	0.0000
51				101.8647	0.0000
52				1.8647	101.8647
53				1.8647	101.8647
54				1.8647	101.8647
55				101.8647	0.0000
56				1.8647	101.8647
57				1.8647	101.8647
58				1.8647	101.8647
59				101.8647	0.0000
60				1.8647	101.8647
61				1.8647	101.8647
62				1.8647	101.8647
63				1.8647	0.0000
64				1.8647	103.0000
Risk-Neutral Value =					100.0000

loans of some sort. Home mortgage loans, second trust deed/home equity loans, and auto loans commonly have an amortizing structure. We demonstrate the slight differences in the cash flow tables that result from an amortizing structure in the next two chapters. The process of valuation, however, is nearly identical to the one used in this chapter. We now turn to alternative techniques for valuation American fixed income options.

Monte Carlo Simulation

Monte Carlo simulation is justifiably popular in financial markets and in this volume, but it must be used with great care. Adams and van Deventer (1993) have discussed the reasons for such caution in detail. In general, if highly accurate closed form/analytical solutions exist, Monte Carlo simulation is not the fastest way to get an accurate answer. In the case of American fixed income options, where an accurate analytical solution is not available for a large number of interest rate risk factors, Monte Carlo simulation has to be used with care. We explain why in this section.

In general, however, Monte Carlo simulation's limitations can be summarized as follows:

- As Hull (1993) notes, “one limitation of the Monte Carlo simulation approach is that it can be used only for European-style derivative securities” (334). To correctly value an American option, one measures value (as we did previously) by working backward from maturity to calculate value, assuming at each decision point that the option holder does the rational thing. Monte Carlo works by projecting one interest rate path at a time, so there is not enough information at any point on that path to correctly analyze whether or not an American option holder should prepay. For this reason, Monte Carlo simulation almost always requires the user to specify a decision rule regarding what the holder of the option should do in any given interest rate scenario, often in the form of a prepayment table or prepayment function in the case of a prepayment option. We analyze the implications of these approaches, which are very common, in the next two chapters. The use of a prepayment table or function is putting the cart before the horse, since the user has to guess how the option will be exercised before the user knows what the option is worth. Rather than going through this error-filled exercise, we think most users would get better results by just guessing the value of the option directly as we have done so far in this chapter in the HJM framework!
- The calculation speed of Monte Carlo techniques is slow for problems where there is a small number of random variables. One example would be the case of a single-factor term structure model with a few time steps. A bushy tree or Hull and White lattice would normally give an accurate and fast answer in this (Model Risk Alert) unrealistic circumstance. Monte Carlo does have a speed advantage for problems with a large number of random variables, and it is this case that is typical for a complete analysis of integrated interest rate risk and credit risk. In that circumstance, the HJM framework needs to be used to generate the Monte Carlo scenarios.
- Monte Carlo simulation by definition does not use all possible scenarios for valuation. There are an infinite number of interest rate scenarios, and Monte Carlo simulation, due to its speed problems, inevitably requires the user to use too few scenarios in the interests of time. In Chapters 36 to 41, we devote some time to this issue of the number of scenarios necessary to achieve statistical significance of a risk measure like value at risk. Some horrible mistakes have been made by some large institutions that we discussed in Chapter 10 and again at length in later chapters.
- As a result of using less than all possible scenarios, Monte Carlo simulations have sampling error, which results from “throwing the dice” too few times. Many users of Monte Carlo simulation are under the mistaken impression that the beautiful probability distribution that is displayed on their computer screen reflects the true uncertainty about the value of a security as measured by Monte Carlo. Nothing could be farther from the truth. If you call a major securities firm to get a bid on a mortgage security, you get a bid in the form of one number, not a probability distribution. The beautiful probability distribution reflects the inaccuracy or sampling error of the technique itself, and this kind of graph reflects a weakness of Monte Carlo, not a strength. This kind of sampling error can lead to very serious problems when calculating hedge amounts. In most of

the circumstances we discuss in the following chapters, there is no realistic alternative to a Monte Carlo simulation and sampling error, and like the warning label on prescription medications, is a risk that needs to be dealt with on the basis of complete information.

- To reduce sampling error to a meaningful level, a sophisticated analyst can calculate the number of simulations required for an answer with sampling error small enough to allow decision making to be based on the Monte Carlo results.¹
- Sampling error becomes even more important when basing hedges on the results of Monte Carlo simulation. The authors feel that the primary purpose of the calculations in this book are to define action, rather than simply to describe the amount of risk an institution currently has. Knowing how much risk you have without knowing what to do about it is nearly useless from a credit risk, market risk, liquidity risk and an interest rate risk perspective. Consider a hedger who values his portfolio using only 200 simulations via Monte Carlo (one of the three largest banks in North America that has used fewer scenarios than this as recently as 2010 for enterprise-wide risk management). The result shows a value of 100 and a sampling error standard deviation of, say, 2. This means that there is roughly a 65 percent probability that the true value lies between 98 and 102. In order to determine the proper amount of the hedge, the analyst shifts rates up by 10 basis points and repeats the analysis, getting a value of 99 and a sampling error standard deviation of 2 again. The analyst concludes that the “delta” of his portfolio is $100 - 99 = 1$ and wants to base his hedge on this result. This is fine, as far as it goes, but the delta has sampling error also. The sampling error on the delta is a function of the sampling error on the two simulation runs, and it is calculated as follows:

$$\sigma_{\text{Hedge } \Delta} = \sqrt{\sigma_{\text{Run 1}}^2 + \sigma_{\text{Run 2}}^2}$$

In the example given, the sampling error of the hedge delta of 1.00 works out to 2.828. What does it mean for the precision of the hedge? It means that there is a 36.2 percent chance that the hedge amount is not only the wrong magnitude *but the wrong sign!* This is a career-ending error that has happened often enough on Wall Street that it’s become a familiar story.²

- The delta from a Monte Carlo simulation can only be derived from doing the calculation twice (or preferably three times for securities with high convexity), further aggravating its speed disadvantages. For simulations with a large number of macro factors driving total risk, there is a delta for each risk factor as well.
- Monte Carlo simulations must be done on the basis of the risk-neutral distributions of all random variables for accurate valuation, an adjustment that many users fail to make.

On the plus side, Monte Carlo simulation has a number of advantages that should not be overlooked:

- It is sometimes the only alternative where the cash flow is path dependent, as imprecise as it might be in that case. A bushy tree can also handle path dependence

accurately and should not be overlooked as an alternative. Longstaff and Schwartz (2001) make a creative suggestion on how to enhance the accuracy of Monte Carlo simulation when a bushy tree approach is not available.

- Monte Carlo has speed advantages for problems with a large number of variables. Generally, this means three or more variables. Recent regulatory stress tests employ 20 to 40 macroeconomic variables, so Monte Carlo is almost an automatic choice for stress testing. This is typically the case for the integrated interest rate risk, market risk, liquidity risk, and credit risk analysis that is the central focus of this book.
- Monte Carlo is relatively simple to implement, although a shockingly small number of enterprise risk management software firms have both built such a system and successfully commercialized it.

Conclusions

What are our conclusions about Monte Carlo simulation? First of all, it is a tool that all users should have access to. Risk management software packages that don't clearly display the sampling error of both value and hedges derived from Monte Carlo simulation calculations should be used with extreme caution. With this caveat, Monte Carlo simulation is essential to the complete integrated analysis of credit risk, market risk, liquidity risk, and interest rate risk. Finally, Monte Carlo needs to be supplemented with the other techniques in this chapter in order to be used with efficiency and accuracy. For example, in a multiperiod Monte Carlo simulation that requires valuation at each point in time, closed form solutions can be used for this valuation when they are available, and the other techniques discussed below can be used when they are not available. The authors believe this "hybrid Monte Carlo" approach is firmly established as the best-practice approach for integrated risk management, subject to the caveats above.

FINITE DIFFERENCE METHODS

Finite difference methods are more complex but more general solution methods commonly used in engineering applications. Finite difference methods provide a direct general solution of the partial differential equation (that defines the price of a callable security), unlike lattice methods, which model the evolution of the random variables. These methods fall into two main groups, *explicit* and *implicit* finite difference methods. The explicit finite difference methods are equivalent to the lattice methods. Implicit methods are more robust in the types of problems that they can handle, but implicit methods are computationally more difficult. Both methods usually use a grid-based calculation method, rather than the lattice approach, to arrive at numerical solutions.

Finite difference methods can be used to solve a wide range of derivative product valuation problems, and they are not restricted to a small number of stochastic processes (i.e., normal or lognormal) that describe how the random variable moves. These methods can be used to value both American- and European-style options. Like lattice methods, the valuation by finite difference methods is performed by stepping backward through time. This is a one-stage process under the finite

difference method; there is no need to model the evolution of the random variable before the valuation can be performed.

The authors believe that the finite difference method, in its grid rather than a lattice implementation, is a very useful tool for most institutions. But its use to date in the financial services industry has been on single instrument valuation in a derivatives trading context, rather than a tool that has been successfully embedded in an enterprise-wide risk management system.

BINOMIAL LATTICES

A *binomial lattice* (also called a *binomial tree*) is a discrete time model for describing the movement of a random variable whose movements at each node on the tree can be reduced to an up or down movement with a known probability. Unlike the one-factor bushy tree used in Chapter 6 in the HJM context, the binomial lattice is usually constructed by constraining the nature of movements of the underlying (usually single) random variable. The model is usually specified so that an upward movement followed by a downward movement gives the same value as a downward movement followed by an upward movement; this means there will be three possible values of the random variable at the end of the second time interval and $k + 1$ possible values at the end of time interval k . This recombining nature of the lattice is what makes it different from the bushy trees used in Chapters 6 through 9. In the most realistic risk applications, constraining random variable movements to recombine reduces accuracy substantially.

Jarrow (1996) uses the binomial tree concept for modeling interest rate movements, and Jarrow and Turnbull (1996) show how the same kind of tree (with appropriate modifications) can be used for modeling stock price movements. The binomial lattice can be used to value American-style options as well as European-style options.

TRINOMIAL LATTICES

In a number of papers with important theoretical and practical implications, Hull and White (1990, 1993b, 1994a) developed the trinomial lattice valuation technique, which has quickly established itself as a popular valuation technique, not only for the Vasicek model and its extended version (the Hull and White model), but also for a number of other single-factor term structure models that are Markov in nature. A *trinomial lattice* (also called a *trinomial tree*) is a discrete time model for describing the movement of a random variable whose movements at each node on the tree can be reduced to one of three possibilities: an up movement, a down movement, or no change. The model is specified so that an upward movement followed by a downward movement and a downward movement followed by an upward movement give the same value as two movements where no change occurs; this means there will be three possible values of the random variable at the end of the first time interval and $k + 2$ possible values at the end of time interval k . This recombining feature maximizes the efficiency of the calculation.

The additional outcome at each node provides an additional degree of freedom that enhances the power of the model; trinomial lattices can be used to model almost any Markov stochastic process, including ones in which the parameters are functions of time and the random variable itself. In modeling interest rate–related derivatives, it allows rates to be modeled so that the current term structure of interest rates and term structure of rate volatility are matched exactly by the modeling process. This offsets the more complicated calculations of the values at the ends of each branch. The trinomial lattice can be used to value American as well as European-style options.

In recent years, the Hull and White lattice approach has been overshadowed by the more general HJM approach, which deals with a large number of interest rate risk factors much more easily, as we have demonstrated in Chapters 6 through 9.

HJM VALUATION OF AMERICAN FIXED INCOME OPTIONS WHEN DEFAULT RISK IS PRESENT

How does the possibility of default affect the call option on an ABC Company 10-year bond that is callable at any time after five years? The credit model and the option are not separable because they impact each other in the following ways:

- The value of the call option is reduced by the possibility that ABC company defaults before it is able to exercise the call option.
- The value of the call option is reduced by the possibility that ABC's default probability is so high when exercising the call is rational that it cannot raise the cash to retire the bonds by calling them.
- The potential losses from default are reduced by the possibility that the bonds have been called before default can occur.

This complex interaction is typical of multipayment callable instruments that range from home mortgages to auto loans to bonds issued by corporations and governmental agencies. How is the HJM bushy tree or Monte Carlo modified to handle this interaction, particularly in the key case (our constant focus) where interest rates and other macroeconomic factors are drivers of default probability levels? We discussed the appropriate adjustments in Chapters 17, 19, and 20 but we review them again here.

Jarrow and Turnbull (1995) describe this process in detail. At each node on the bushy tree, the default probability of the counterparty will take on a different value because the default probabilities vary by time and level of interest rates. Also, at each node on the lattice, the company can either default or not default. If the company is in default, the payoff on the security is its defaulted value.³

Jarrow (1999) goes on to show in detail how to construct lattice and bushy tree valuation techniques for the Jarrow model described in Chapter 16, where the default intensity is driven by both interest rates and other macroeconomic factors.

In this chapter, we have assumed that the American option is exercised “rationally,” with no considerations other than the obvious financial considerations driving exercise. We now turn to “irrational” exercise of American options and discuss the implications for integrated enterprise risk management.

NOTES

1. Sophisticated users of the enterprise risk management system constructed by the authors and described on www.kamakuraco.com routinely do simulations with hundreds of thousands or millions of scenarios.
2. A major New York bank reported a loss of more than \$100 million after discovering the Monte Carlo simulation routine it was using to value its portfolio was producing a “gamma” (or second derivative) with the wrong sign; the loss was the magnitude of the mark-to-market error discovered after using a more sophisticated technique to value the same portfolio.
3. In the Jarrow-Turnbull (1995) case, this is principal times recovery rate; in the Jarrow (1999, 2001) and Duffie-Singleton (1999) models, this is the recovery rate times the value of the security just an instant before default.

Irrational Exercise of Fixed Income Options

As we continue to progress toward the mark to market of the entire balance sheet of a financial institution—and in a way that integrates interest rate risk, market risk, liquidity risk, and credit risk—we have to deal with this reality: The vast majority (by transaction count) of financial institutions' investments are extensions of credit to individuals or small businesses whose creditworthiness is inseparable from the creditworthiness of the proprietor. Moreover, in Chapter 29, we deal with securitized pools of assets where the same problems exist.

One reason why the analysis of consumer-related financial products is so complex is the so-called “irrationality” of consumers. Perhaps the most prominent example from a Wall Street perspective is the mortgage market, which is at the heart of the 2006–2011 credit crisis. Mortgages are often prepaid when current mortgage rates are higher than the borrower's mortgage rate, and many borrowers fail to prepay even when rates have fallen far below the rate on their loan. In the credit crisis, this was usually due to the fact that the house price had fallen so low that it was worth less than the principal on the loan. Besides these common examples of irrationality, partial prepayments are common, something that is impossible in a traditional Black-Scholes context or in the context of the fixed income options analysis we explored in Chapter 21. In both of those cases, options exercise is “all or nothing,” so many observers classify partial prepayments by consumers as irrational.

As a working assumption, the hypothesis that borrowers are not very intelligent runs contrary to the assumptions behind all developments in modern financial theory over the 40 years since the Black-Scholes options formula was first published. The perception of irrationality is simply a shorthand description for the fact that lenders and academic researchers do not have enough data to see through the individual loan or loan pool data to understand why the individual borrower's behavior is more rational than it appears at first glance. We believe even more strongly that the underlying assumption of rationality is the right one since one of us has failed to refinance his mortgage loan in the past, even though it must seem irrational from the bank's perspective. The bank was simply unaware that the first edition of this book was more important than a few yen earned from refinancing a Japanese mortgage in the midst of writing the book. Once the book was done, the “irrational” author suddenly regained his rationality from the bank's perspective.

The need to deal with hard-to-explain consumer behavior is essential for accurate risk management in the insurance, investment management, and banking

industries. Insurance companies must deal with the reality that some traditional whole life insurance policyholders will cancel their policies in return for the surrender value when interest rates rise, putting the policy back to the insurance company. Investment managers who suffer poor performance know that some customers of the firm will put their ownership in a mutual fund back to the investment management company. Bankers who make home equity loans know that some of these loans will be called (as in Chapter 27) or prepaid by the borrower. Bankers also know that some time-deposit customers will put the deposit back to the bank and willingly pay an early withdrawal penalty when rates rise in order to earn more on their funds. Without dealing with the correct degree of irrational consumer behavior behind these products, institutions can neither price nor hedge their risk properly. This chapter deals explicitly with a method for doing so.

In the last section of the chapter, we explore the strong synergies between our approach to irrational behavior and the credit-adjusted valuation and simulation we have been discussing in earlier chapters.

ANALYSIS OF IRRATIONALITY: CRITERIA FOR A POWERFUL EXPLANATION

The irrationality problem is so all-pervasive that we deal with the problem in general terms in this chapter. We deal with alternative approaches to irrationality in Chapter 29. This approach has two virtues. First, it will help us deal with the problem from a fresh perspective, unburdened by conventional wisdom. Second, armed with a general solution, we can compare this general solution to the conventional wisdom and judge the relative strengths of the two approaches without prejudice. In Chapter 29, armed with the tools of this chapter, we analyze all the traditional practices of mortgage-related products in light of both this chapter and the truths revealed (or confirmed) by the 2006–2011 credit crisis. In this chapter, however, we need to establish principles that will solve problems in insurance, investment management, and in banking. They have to apply equally well to call options, put options, and securities that are callable but have rate caps and floors as well. A few basic principles come to mind:

1. A general approach to irrational behavior should take advantage of, to as great an extent as possible, the advances in finance in the past 40 years.

All derivative pricing theory is based on the premise of no-arbitrage and rational behavior. Behavioral finance has revealed that, in many cases, “rational behavior” leads to exceptions that turn out to be rational after all, if only the academic researchers had made more realistic assumptions in the first place. If all consumers were totally irrational, we would not need to price the options embedded in retail finance products, because the exercise of those options would be truly random and uncorrelated with economic events. Portfolio theory would allow us to argue that we can diversify away this random behavior and ignore the option all together. The reality, as we shall see in Chapter 29, however, is that consumer behavior is highly correlated with economic variables and behavior is not totally irrational. It is this mixture of rational and irrational behavior (which, in fact, is largely due to rational acts based on unobservable information), that makes the problem both important and difficult.

We need, then, an approach that allows us to scale the degree of irrationality to fit the particular product, company, and market at hand. We need an approach where consumer action is partially rational so we can use the first 27 chapters of this book.

2. History is not necessarily a good guide to the future, so we want to be able to derive the level of irrationality implied by observable market prices as a supplement to the analysis of historical behavior.

Wall Street firms have spent millions of dollars examining historical data for clues to consumer behavior as reflected in historical prepayment activity. As we shall see in Chapter 29, this work has been at best partly satisfactory. For financial institutions that need to make an immediate decision about the proper hedge for a security, whose market price and local rate sensitivity (i.e., the value change for small changes in interest rates) is observable, the historical approach can be too slow, too inaccurate, and too expensive. The parallels in the debate about the use of historical volatility versus implied volatility in the Black-Scholes options model are very strong. We want an approach that gives us the implied level of irrationality. Most bankers who price retail banking products that contain embedded options come to the conclusion that most products are unprofitable if the options risk is fully hedged. Bankers implicitly recognize that consumers are partially rational and that a partial hedge is necessary. How can we use market data to calculate exactly how much rationality is embedded in a given product?

3. Any model of irrationality must be able to explain observable phenomenon in pools of consumer-related securities: path-dependent prepayment behavior, decreased interest rate sensitivity over time (burnout), and a lower propensity to exercise embedded options as the financial product nears its expiration.
4. The model must allow for very rapid calculation time and maximum use of analytical solutions for security valuation.

One approach that meets these criteria is the transactions cost approach, which we explain in the next section.

THE TRANSACTIONS COST APPROACH

Over the past 20 years, the rational approach to irrational behavior has steadily gained prominence among both academics (see McConnell and Singh 1994 and Stanton 1995 for early examples) and practitioners as the best method to model consumer behavior. The rational, or transactions cost, approach holds that consumers are rational but exercise their options subject to transactions costs. These transaction costs can be both explicit financial costs (the points from refinancing a mortgage or the penalty for early withdrawal of bank deposits) and more subtle costs that reflect the fact that consumers are maximizing a utility function with more arguments in it than the present value of rational action in one specific aspect of their lives. For example, someone with a life insurance policy that rationally should be canceled and reinitiated at current market pricing may have a serious disease that would cause the applicant to be rejected at the health check on the new policy. The disease is a transactions cost that at least partially blocks rational action. J. Thurston

Howell III, the wealthy castaway in the old American sitcom *Gilligan's Island*, may fail to refinance the mortgage on his ski chalet in Zermatt because his chartered yacht has run aground on some deserted desert island without a Bloomberg screen to keep him abreast of the benefits of refinancing. In our case, the principal amount of our bank time deposits is so small (since other investments are so superior in risk-adjusted yield), that the benefits of early withdrawal of our time deposits, while positive, are smaller than the cost of the time it takes to go to the bank and negotiate the transaction. The cost of our time is a transactions cost. Finally, in the 2006–2011 credit crisis, the impact of home prices on both the desire and the ability to prepay a mortgage is a hugely important factor in whether the mortgage repays or not. The “underwater” amount by which the loan balance exceeds the current home price value can be thought of as a large transaction cost as well.

The objective of the transactions cost approach is to firmly divorce rational from irrational behavior. Accordingly, the transactions cost function can be any time-dependent function that is not dependent on the level of interest rates. Determinants of the level of transactions costs could be any of the following:

- Time to maturity on the security
- The level of the coupon on the security
- The fixed dollar cost of exercising the option embedded in the security
- The dollar opportunity cost of the time it takes to exercise the embedded option
- The value of the collateral on the loan, relative to the loan balance
- The month of the year, recognizing that the transactions cost of refinancing a home in Alaska in January is higher than the cost in August
- The date of origination of the security
- The length of time the security has been outstanding without the embedded option being exercised

There are an endless number of factors that can go into the transactions cost function. In short, however, as of current time t , all of these factors are nonrandom functions of time. We call this transactions cost function X and make it, without loss of generality, a function of the remaining time to maturity on the security:

$$\text{Transactions cost} = X(t, T) = X(\tau)$$

We can model any degree of irrationality using this function. If the consumer is a retired Salomon Brothers partner living in Greenwich with a Bloomberg screen in his living room and a healthy bank account, $X = 0$. The consumer is totally rational. If the consumer is working hard on a book on risk management, X is infinity, the consumer is totally irrational, and no embedded options will be exercised. Any level of irrationality in between these two extremes can be modeled in the same way. We illustrate the approach for European options in the next section.

IRRATIONAL EXERCISE OF EUROPEAN OPTIONS

In Chapter 21, we analyzed the value of a European option on a zero-coupon bond using both the three-factor Heath, Jarrow, and Morton (HJM) framework and

Jamshidian's valuation formula from the one-factor Vasicek term structure model. How would the analysis of a European option in Chapter 21 differ from the results presented in that chapter if there were barriers to options exercise due to transactions costs? If the option at issue is a European option, the time-dependent nature of the transactions cost X is irrelevant with regard to a single European option. Only the level of X that will prevail at the exercise time T_1 matters, so we drop the time-dependent notation. The only difference between the rationally exercised option on a zero-coupon bond P in Chapter 21 and the irrationally exercised option is easy to state. The "irrational" holder of a fixed income option exercises it only when the short rate of interest on date T_1 reaches a critical level s^* (in the Vasicek context) or state s_t in the HJM context such that:

$$P(s^*, T_1, T_2) = K + X$$

When X is zero, the analysis produces the same answer as the assumption that the option is exercised rationally. If X is infinity, the option will never be exercised, and the option is worthless from the perspective of the holder. We illustrate the power of this formulation with a specific example.

THE IRRATIONAL EXERCISE OF AMERICAN OPTIONS

In Chapter 27, we reviewed the fundamentals of the pricing of American fixed income options. An American call option was valued on the premise of rational behavior throughout the life of the call, such that the holder of the call always acted in such a way as to maximize the value of the call. In the case of a callable security, the holder of the call always acts in such a way as to minimize the value of the security. Working backward, one period from maturity, the holder of the call is assumed to prepay if there is even a one cent advantage of prepayment at par when compared to the present value of leaving the security outstanding at least one more period.

In the case of an irrationally exercised American call option, at every point (working backward from maturity) the holder of an option to call their bond or loan with principal amount B will exercise that call option only if:

$$\text{Value if prepayment option unexercised} > B(\tau) + X(\tau)$$

That is, exercise will take place at that point in time and at that state of interest rates only if the present value of the security—if the call option is left unexercised—exceeds the principal amount B by more than the transactions cost X prevailing at that instant in time. The fact that X varies over time will result in a very rich array of realistic behavior in securities modeled in this way.

A Worked Example Using an Amortizing Loan with Rational and Irrational Prepayment Behavior

In this example, we modify the American options analysis of Chapter 27 using the three-factor HJM valuation framework from Chapter 9. In Chapter 27, we analyzed a bullet bond with all principal (\$100) due at the maturity of the bond at time $T = 4$. In this section, we assume instead that the loan has a more common structure for

consumer credits, an amortizing structure with a level payment that ultimately retires all principal at maturity at $T = 4$. If the loan has a principal amount at time 0 and a 2.00 percent coupon rate, the reader can verify that the loan will have annual payments of \$26.2624, divided into interest and principal payments due at each annual time period (we use annual, not monthly, periods for expositional convenience) as shown in Exhibit 28.1.

If there are infinitely high transactions costs, the loan will never be prepaid and we will have cash flows of \$26.2624 at times $T = 1, 2, 3$, and 4 and in every state or at each node on the HJM bushy tree. This is shown in Exhibit 28.2, where the valuation comes from dropping these cash flows into the cash flow table that is then multiplied by the same probability-weighted discount factors that we have employed since Chapter 9. Exhibit 28.2 shows a risk-neutral value of 101.7369.

In a no-arbitrage world, this value should equal the product of our annual cash flow \$26.2624 and the sum of the zero-coupon bond prices for maturity at times 1, 2, 3, and 4 used as inputs to the HJM bushy tree: 0.9970057865, 0.9841101497, 0.9618922376, and 0.9308550992. (The reader can verify that value is indeed 101.7369, consistent with Chapter 9 and the formulas of Chapter 4.) How do the cash flows vary if the borrower under this loan is completely rational with zero transactions costs? We use the same analysis we employed in Chapter 27 to answer this question. We start at the last options exercise point, time $T = 3$, and analyze the optimal options exercise policy. We prepay or call the loan whenever the risk-neutral value of the payment \$26.2624 due at time $T = 4$ is more than the principal that remains on the loan (\$25.7474 as shown in Exhibit 28.1) after making the payment of \$26.2624 due at time $T = 3$. Exhibit 28.3 gives the rational prepayment behavior for each of the 64 nodes on the HJM bushy tree at time $T = 3$.

The cash flows that result from this call strategy at time $T = 3$ (ignoring our ability to call or prepay the loan at times $T = 1$ and 2) are shown in Exhibit 28.4.

Following the procedures of Chapter 27, we then step back to time $T = 2$ and again reexamine the optimal call or prepayment strategy. When we do that, we may decide to prepay the loan on a branch of the bushy tree for which prepaying at time $T = 3$ was also rational. In this case, once the prepayment is made at time $T = 2$, the

EXHIBIT 28.1 Scheduled Amortization

Scheduled Amortization

Coupon Rate:	2.0000%
Time 0 Book Value	100
Periods to Maturity	4
Periodic Payment Amount	26.2624

Period Number	Initial Principal	Amount Paid	Interest	Principal	Ending Principal
0	100.0000				100.0000
1	100.0000	26.2624	2.0000	24.2624	75.7376
2	75.7376	26.2624	1.5148	24.7476	50.9900
3	50.9900	26.2624	1.0198	25.2426	25.7474
4	25.7474	26.2624	0.5149	25.7474	0.0000

EXHIBIT 28.2 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	26.2624	26.2624	26.2624	26.2624
2		26.2624	26.2624	26.2624	26.2624
3		26.2624	26.2624	26.2624	26.2624
4		26.2624	26.2624	26.2624	26.2624
5			26.2624	26.2624	26.2624
6			26.2624	26.2624	26.2624
7			26.2624	26.2624	26.2624
8			26.2624	26.2624	26.2624
9			26.2624	26.2624	26.2624
10			26.2624	26.2624	26.2624
11			26.2624	26.2624	26.2624
12			26.2624	26.2624	26.2624
13			26.2624	26.2624	26.2624
14			26.2624	26.2624	26.2624
15			26.2624	26.2624	26.2624
16			26.2624	26.2624	26.2624
17				26.2624	26.2624
18				26.2624	26.2624
19				26.2624	26.2624
20				26.2624	26.2624
21				26.2624	26.2624
22				26.2624	26.2624
23				26.2624	26.2624
24				26.2624	26.2624
25				26.2624	26.2624
26				26.2624	26.2624
27				26.2624	26.2624
28				26.2624	26.2624
29				26.2624	26.2624
30				26.2624	26.2624
31				26.2624	26.2624
32				26.2624	26.2624
33				26.2624	26.2624
34				26.2624	26.2624
35				26.2624	26.2624
36				26.2624	26.2624
37				26.2624	26.2624
38				26.2624	26.2624
39				26.2624	26.2624
40				26.2624	26.2624
41				26.2624	26.2624
42				26.2624	26.2624
43				26.2624	26.2624

EXHIBIT 28.2 (Continued)

Row Number	Current Time				
	0	1	2	3	4
44				26.2624	26.2624
45				26.2624	26.2624
46				26.2624	26.2624
47				26.2624	26.2624
48				26.2624	26.2624
49				26.2624	26.2624
50				26.2624	26.2624
51				26.2624	26.2624
52				26.2624	26.2624
53				26.2624	26.2624
54				26.2624	26.2624
55				26.2624	26.2624
56				26.2624	26.2624
57				26.2624	26.2624
58				26.2624	26.2624
59				26.2624	26.2624
60				26.2624	26.2624
61				26.2624	26.2624
62				26.2624	26.2624
63				26.2624	26.2624
64				26.2624	26.2624
Risk-Neutral Value =					101.7369

EXHIBIT 28.3 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	Action?
1	0.0000	0.0000	0.0000	24.7500	No Call
2		0.0000	0.0000	24.8775	No Call
3		0.0000	0.0000	24.9711	No Call
4		0.0000	0.0000	24.3086	No Call
5			0.0000	25.4717	No Call
6			0.0000	25.5990	No Call
7			0.0000	25.7401	No Call
8			0.0000	25.2370	No Call
9			0.0000	24.9593	No Call
10			0.0000	25.7529	Call

(Continued)

EXHIBIT 28.3 *(Continued)*

Row Number	Current Time				Action?
	0	1	2	3	
11			0.0000	25.4530	No Call
12			0.0000	24.8929	No Call
13			0.0000	24.9109	No Call
14			0.0000	25.0393	No Call
15			0.0000	25.1335	No Call
16			0.0000	24.4666	No Call
17				25.2155	No Call
18				26.0173	Call
19				25.7143	No Call
20				25.1484	No Call
21				25.6961	No Call
22				26.1179	Call
23				26.0625	Call
24				25.9509	Call
25				25.5911	No Call
26				26.3608	Call
27				26.1944	Call
28				25.7051	No Call
29				25.8040	Call
30				25.9330	Call
31				26.0759	Call
32				25.5663	No Call
33				25.6768	No Call
34				25.8052	Call
35				25.9474	Call
36				25.4403	No Call
37				25.6768	No Call
38				25.8052	Call
39				25.9474	Call
40				25.4403	No Call
41				25.3411	No Call
42				26.1033	Call
43				25.9386	Call
44				25.4540	No Call
45				25.1298	No Call
46				25.9288	Call
47				25.6269	No Call
48				25.0629	No Call
49				25.3740	No Call
50				26.1808	Call
51				25.8759	Call
52				25.3065	No Call
53				25.6011	No Call

EXHIBIT 28.3 (Continued)

Row Number	Current Time				Action?
	0	1	2	3	
54				25.7290	No Call
55				25.8708	Call
56				25.3652	No Call
57				25.6011	No Call
58				25.7290	No Call
59				25.8708	Call
60				25.3652	No Call
61				25.1300	No Call
62				25.6798	No Call
63				25.5657	No Call
64				25.0063	No Call

EXHIBIT 28.4 Maturity of Cash Flow Received

Row Number	Current Time			
	0	1	2	3
1		26.2624	26.2624	26.2624
2		26.2624	26.2624	26.2624
3		26.2624	26.2624	26.2624
4		26.2624	26.2624	26.2624
5			26.2624	26.2624
6			26.2624	26.2624
7			26.2624	26.2624
8			26.2624	26.2624
9			26.2624	26.2624
10			26.2624	52.0098
11			26.2624	26.2624
12			26.2624	26.2624
13			26.2624	26.2624
14			26.2624	26.2624
15			26.2624	26.2624
16			26.2624	26.2624
17				26.2624
18				52.0098
19				26.2624
20				26.2624
21				26.2624
22				52.0098

(Continued)

EXHIBIT 28.4 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
23				52.0098	0.0000
24				52.0098	0.0000
25				26.2624	26.2624
26				52.0098	0.0000
27				52.0098	0.0000
28				26.2624	26.2624
29				52.0098	0.0000
30				52.0098	0.0000
31				52.0098	0.0000
32				26.2624	26.2624
33				26.2624	26.2624
34				52.0098	0.0000
35				52.0098	0.0000
36				26.2624	26.2624
37				26.2624	26.2624
38				52.0098	0.0000
39				52.0098	0.0000
40				26.2624	26.2624
41				26.2624	26.2624
42				52.0098	0.0000
43				52.0098	0.0000
44				26.2624	26.2624
45				26.2624	26.2624
46				52.0098	0.0000
47				26.2624	26.2624
48				26.2624	26.2624
49				26.2624	26.2624
50				52.0098	0.0000
51				52.0098	0.0000
52				26.2624	26.2624
53				26.2624	26.2624
54				26.2624	26.2624
55				52.0098	0.0000
56				26.2624	26.2624
57				26.2624	26.2624
58				26.2624	26.2624
59				52.0098	0.0000
60				26.2624	26.2624
61				26.2624	26.2624
62				26.2624	26.2624
63				26.2624	26.2624
64				26.2624	26.2624

cash flows at times $T = 3$ and 4 on this branch of the bushy tree are set to zero. We set the strategy for time $T = 2$ and then step back to time $T = 1$ and repeat the process. When we do that, we get the cash flows and valuation shown in Exhibit 28.5 for zero transactions costs and a 100 percent rational borrower.

If the borrower is not allowed to prepay at time 0, the loan has a value of 101.4234. If the borrower is allowed to prepay at time zero, the borrower (if he is totally rational) will prepay so that

EXHIBIT 28.5 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	26.2624	26.2624	26.2624	26.2624
2		102.0000	26.2624	26.2624	26.2624
3		102.0000	26.2624	26.2624	26.2624
4		26.2624	26.2624	26.2624	26.2624
5			0.0000	26.2624	26.2624
6			0.0000	26.2624	26.2624
7			0.0000	26.2624	26.2624
8			0.0000	26.2624	26.2624
9			0.0000	26.2624	26.2624
10			0.0000	52.0098	0.0000
11			0.0000	26.2624	26.2624
12			0.0000	26.2624	26.2624
13			26.2624	26.2624	26.2624
14			26.2624	26.2624	26.2624
15			77.2524	26.2624	26.2624
16			26.2624	26.2624	26.2624
17				0.0000	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000

(Continued)

EXHIBIT 28.5 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				0.0000	0.0000
46				0.0000	0.0000
47				0.0000	0.0000
48				0.0000	0.0000
49				26.2624	26.2624
50				52.0098	0.0000
51				52.0098	0.0000
52				26.2624	26.2624
53				26.2624	26.2624
54				26.2624	26.2624
55				52.0098	0.0000
56				26.2624	26.2624
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				26.2624	26.2624
62				26.2624	26.2624
63				26.2624	26.2624
64				26.2624	26.2624
Risk-Neutral Value =					101.4234

$$\begin{aligned}
 \text{Value} &= \text{Minimum}(\text{par value, risk-neutral value if not prepaid}) \\
 &= \text{Minimum}(100, 101.4234) \\
 &= 100
 \end{aligned}$$

Now, let's allow our client to be rational but subject to transactions costs. If prepayment is allowed at time zero, we know that the loan will be prepaid and have a time zero value of 100 for any transactions cost that is $101.4234 - 100 = 1.4234$ or less. Let's take a more interesting case. Let's assume that the loan was made one second earlier and that time $T = 1$ is the borrower's first opportunity to prepay. Let's assume we can see the loan trading in the secondary market at a net present value (price plus accrued interest, consistent with our comments in Chapter 4) of 101.56. What

constant transactions cost (relevant to exercise at times $T = 1, 2$, and 3) is implied by this market price? We use iteration with respect to the transaction cost in common spreadsheet software to find the answer. Note that the typical solver function isn't very helpful in this case because our valuation answer is in some sense a step function, rather than a smooth continuous function, due to the size of the time steps and the small number of branches on the bushy tree relative to what would be used in practice. We find that a transactions cost of \$0.92 produces a value of 101.56 and generates the cash flows (and implied prepayment strategies) shown in Exhibit 28.6.

We now turn to the implications of the transactions cost approach to modeling irrational behavior.

EXHIBIT 28.6 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	26.2624	26.2624	26.2624	26.2624
2		102.0000	26.2624	26.2624	26.2624
3		26.2624	26.2624	26.2624	26.2624
4		26.2624	26.2624	26.2624	26.2624
5			0.0000	26.2624	26.2624
6			0.0000	26.2624	26.2624
7			0.0000	26.2624	26.2624
8			0.0000	26.2624	26.2624
9			26.2624	26.2624	26.2624
10			26.2624	26.2624	26.2624
11			26.2624	26.2624	26.2624
12			26.2624	26.2624	26.2624
13			26.2624	26.2624	26.2624
14			26.2624	26.2624	26.2624
15			26.2624	26.2624	26.2624
16			26.2624	26.2624	26.2624
17				0.0000	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000

(Continued)

EXHIBIT 28.6 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
33				26.2624	26.2624
34				26.2624	26.2624
35				26.2624	26.2624
36				26.2624	26.2624
37				26.2624	26.2624
38				26.2624	26.2624
39				26.2624	26.2624
40				26.2624	26.2624
41				26.2624	26.2624
42				26.2624	26.2624
43				26.2624	26.2624
44				26.2624	26.2624
45				26.2624	26.2624
46				26.2624	26.2624
47				26.2624	26.2624
48				26.2624	26.2624
49				26.2624	26.2624
50				26.2624	26.2624
51				26.2624	26.2624
52				26.2624	26.2624
53				26.2624	26.2624
54				26.2624	26.2624
55				26.2624	26.2624
56				26.2624	26.2624
57				26.2624	26.2624
58				26.2624	26.2624
59				26.2624	26.2624
60				26.2624	26.2624
61				26.2624	26.2624
62				26.2624	26.2624
63				26.2624	26.2624
64				26.2624	26.2624
Risk-Neutral Value =					101.5589

IMPLIED IRRATIONALITY AND HEDGING

When irrationality is modeled in this way, the level of irrationality can be implied from observable market prices just like implied volatility in the Black-Scholes model, achieving one of our primary objectives for a model of irrationality. Almost no observable security reflects the behavior of one consumer, so the best replication of the market price movement of a given security reflects a combination of irrational

behavior by many consumers. Using standard HJM procedures, we know the value of the loan as time changes from time 0 to one of the four shift states at time $T = 1$. From these value changes, we can solve for the hedging position that generates zero risk; that is, a return equal to the return on the money market fund (an investment for one period at the risk-free rate).

This powerful insight of the HJM framework gives us an alternative to historical research on prepayment behavior—we can imply the degree of rationality of prepayment behavior from observable prices. We have the tools for efficient hedging of securities that contain irrationally exercised embedded options, and we can achieve very fast calculation times with much greater accuracy than we could using the legacy prepayment tools discussed in the next chapter, historical prepayment functions, or prepayment tables. We close this chapter by exploring the links between our credit risk analysis of Chapters 16 and 17 and prepayment analysis.

CREDIT RISK AND IRRATIONAL PREPAYMENT BEHAVIOR

One of the realities of consumer lending is that consumers appear irrational in the three ways discussed in the previous example:

- They sometimes fail to prepay *when they should*.
- They sometimes prepay *when they shouldn't*.
- They often partially prepay instead of fully prepaying *as they should*.

In each of these cases, the word “should” stems from an analysis of the consumer’s behavior that is based only on what the lender knows—the terms of the loan and not much else. Many of these allegedly irrational acts stem from the complex interaction of the long-term creditworthiness and financial well-being of the borrower and the loan itself. We can give rational explanations for each of these types of irrational behavior:

- They sometimes fail to prepay *when they should*.
 - *Case A*: The price of the home may be less than the principal amount of the loan and the consumer may not have enough liquid financial assets to bridge this gap. This is extremely common in the wake of the 2006–2011 credit crisis and home price collapse.
 - *Case B*: The consumer may be in financial distress and unable to qualify for a new loan.
 - *Case C*: The consumer may be on a one-year sabbatical at the University of Tahiti.
 - *Case D*: The loan may be so close to maturity that the benefits of refinancing are not worth the costs of going through the process.
 - *Case E*: The consumer’s personal wealth may now be so great that they have no intention of borrowing any more money and prepayment may simply be delayed until the consumer has enough liquid assets to prepay with cash.

There are many other similar examples we could give to further illustrate this point.

- They sometimes prepay *when they shouldn't*.

Even in this circumstance, there are many explanations that are quite rational for this superficially irrational response, many of which again are due to the creditworthiness and net worth of the consumer:

- *Case A:* The borrower may have sold the house (due to the famous trio of death, divorce, or relocation).
- *Case B:* The borrower may have just received very large payments from structuring CDOs for Goldman Sachs in the credit crisis; prepaying the mortgage is in fact a better risk-adjusted investment than U.S. Treasuries since the credit risk of the now very wealthy consumer is the same as the Treasury: zero. The mortgage is prepaid because it has a higher yield than Treasuries and the lender is not reflecting the borrower's credit risk correctly.
- *Case C:* The borrower may have been a problem borrower in the past and prepays the loan simply because the borrower can no longer tolerate a relationship with a financial institution that has pursued overly aggressive tactics to get the borrower current on the loan; this is a common refrain from consumers these days.
- They often partially prepay instead of fully prepaying *as they should*.

This phenomenon is inconsistent with the “all-or-nothing” implications of fixed income options explained in Chapters 21 and 27, but it is a very common practice. It has a lot in common with the partial drawdowns of commercial lines of credit and with revolving charge card balances that we discuss in later chapters.

 - *Case A:* The consumer's credit quality has improved considerably since the mortgage was granted and the true market mortgage rate for the consumer is now less than the rate on the loan, but the bank doesn't know this. The consumer signals this by prepaying in whatever amounts fit his total financial plan.
 - *Case B:* Market interest rates are now lower but the consumer's total financial plan suggests that occasional partial prepayments are better than refinancing the loan. There are a number of reasons why this may be the case. Most revolve around the fact that partial prepayments are less costly (with “cost” broadly defined) to the consumer than one lump-sum prepayment.

We explore the implications of these links between credit risk, market risk, liquidity risk, and interest rate risk and prepayment behavior for the Jarrow-Merton put option as a measure of integrated risk in detail using U.S. mortgage and consumer loan data in the next chapter.

Mortgage-Backed Securities and Asset-Backed Securities

We are now in the “home stretch” as we prepare for the comprehensive Jarrow-Merton put option as a measure of integrated credit risk, market risk, liquidity risk, and interest rate risk. Our task in this chapter is to apply the analysis of option exercise subject to transaction costs from the previous chapter to portfolios or pools of loans. The authors feel strongly that this should be done on a transaction-level basis, loan by loan, rather than analyzing the pool or portfolio as if that pool was a single transaction. If there is one lesson from the 2006–2011 credit crisis, it is that failure to do the analysis of securitized loans at the transaction level is a recipe for disaster. The complete disappearance of the collateralized debt obligation (CDO) market, outlined in Chapter 20, came about as more and more investors realized that the complexity of that structure was intentional; it was a convenient way to discourage investors from looking closely at the underlying collateral. It facilitated the unending attempts by Wall Street to sell investors lumps of coal packaged as “early stage diamonds.”¹

The credit crisis came about because the process of securitization helped to obscure the true nature of the underlying mortgages that were the ultimate collateral in mortgage-backed securities that were in turn, after being split into tranches, used as collateral in CDOs. These mortgages presented loans whose principal was greater than the value of the underlying house and whose borrowers were the prototypical NINJAs (“no income, no job, and no assets”). Investors ignored these huge red flags because they did not want to spend the time or money to investigate what they were buying at the transaction level. If ever the reader feels the same way, go home immediately or call in sick to work. No structured product should ever be purchased without transaction-level due diligence. If the reader’s institution cannot afford the time or money to do this, it should not be buying structured products.

From a computer science and financial theory perspective, there is no reason for allowing unnecessary data aggregation from the structuring process to obscure our analysis of the risk-adjusted return on the mortgages and consumer loans that make up the ultimate collateral in a structured product. We illustrate the reasons for this view during the course of the chapter and illustrate how the traditional approaches to prepayment are, in fact, off track because of attempts to value a pool of loans as if they were a single very large loan. Even if one ignores the impact of home prices on defaults and prepayments, the traditional methods for measuring the interest rate risk of pools of loans were misguided attempts to shortcut a loan-by-loan analysis.

We use the concept of transactions costs in Chapter 28 to illustrate the pitfalls of common Wall Street shortcuts, such as prepayment rates, prepayment tables, prepayment functions—and “involuntary prepayment” known to normal human beings as default. We also introduce a more modern approach to joint modeling of default and prepayment on pools of loans, multinomial logit. Finally, we close with some comments on the ultimate securitized product, the valuation of mortgage-servicing rights.

TRANSACTIONS COSTS, PREPAYMENTS, DEFAULT, AND MULTINOMIAL LOGIT

In Chapter 28, we analyzed rational prepayment, in which the decision rule was a simple question and its answer: Which strategy minimizes the net present value of my borrowing costs, prepaying now or deciding one period from now? In Chapter 28, the probability of prepaying now was assumed to be 100 percent if that option lowered the present value of costs. In this section, we add another tool to the transactions cost technique for analyzing prepayments and defaults.

Traditional asset and liability management (ALM) has ignored mortgage defaults and focused on interest rate-driven and mortgage age-driven prepayment. The events of the 2006–2011 credit crisis have made it more obvious that prepayment and default are intimately linked and that home prices are a critical driver of both probabilities. This section explains how mortgage prepayment and default are modeled on an integrated basis using multinomial logit. This technique can be applied both to a single mortgage loan and (more likely) to a pool of mortgage loans on a loan-by-loan basis.

The events of the 2006–2011 period have made it clear that the sole focus on a single macroeconomic factor, interest rates, has been insufficiently robust to prevent the effective failures of institutions such as Lehman Brothers, Washington Mutual, New Century, IndyMac, Countrywide, FNMA (Fannie Mae), and FHLMC (Freddie Mac). The more modern approach to risk management is a fully integrated approach to credit risk, market risk, liquidity risk, and interest rate risk, the central theme of this book. In the context of traditional interest rate risk management, default was almost always ignored by most practitioners for two reasons:

- Credit risk was in a “separate risk silo” managed by others.
- Most silo-focused interest rate risk assessment software was incapable of modeling default even if the analyst had wanted to do so.

Both of these constraints, one political and one technical, have now been removed at institutions at the forefront of best practice risk management. One of the tasks triggered by this new approach to integrated risk is that the traditional assumptions of ALM have to be modernized, one by one. In no area is this truer than in mortgage prepayment and default analysis. The techniques traditionally used for mortgage prepayment analysis can be listed briefly as follows:

- Constant prepayment rate (CPR), either expressed as a conditional prepayment rate (annualized basis) or as a single monthly mortality rate. This assumption

ignores the impact of interest rates and home prices on prepayment. It is also derived from behavior of mortgage pools rather than individual loan data.

- Prepayment rate as a linear function of mortgage age, interest rates, and mortgage attributes. This approach also ignores home prices as a driving factor and is derived from pooled behavior rather than individual loan behavior.
- Prepayment tables where the prepayment rate is driven by a table of N explanatory factors, typically the same factors as the linear prepayment function. Again, these tables are typically derived from pools of mortgages and the related prepayment rate on the pool.
- Rational prepayment models, based on the all-or-nothing valuation of the mortgage on an American option basis. Again, default and home prices are often ignored. Chapter 28 explained this approach in detail, with default risk incorporated if the rational prepayment analysis uses a valuation yield curve that reflects the probability of default.
- Rational prepayment subject to transaction costs, which recognizes the true transaction costs of refinancing and other implicit costs that cause the prepayment on mortgages to be less frequent than the fully rational approach would predict. We introduced this approach in the latter half of Chapter 28 as well.

In all of these approaches, the amount of prepayment that occurs, given market conditions, is never random. Given the inputs to each of these techniques, the prepayment amount is known with certainty.

Best practice approaches differ in three key dimensions:

- Prepayment is recognized as being random, just like default, although the probability of prepayment can be estimated just like a default probability.
- Prepayment and default are being analyzed at the individual loan level. The fact that those who analyzed mortgage collateral at the loan level fared much better than investors who did not (i.e., asset-backed CDO investors and the rating agencies) during the credit crisis has not been lost on senior management.
- Default and prepayment are increasingly being modeled as mutually exclusive events that are modeled together, not separately.

The first step forward in recognizing home prices and other loan-specific inputs in the prepay/don't prepay decision and the default/no-default outcome was the use of logistic regression. Logistic regression is commonly used to model events that have only binary 0/1 outcomes, like we saw in Chapter 16 on reduced form credit models. Many analysts have applied this technique to mortgage prepayment and default, estimating separate logistic regressions for these two events. The impact of common risk factors like home prices, interest rates, and mortgage age can be explicitly taken into account on a loan-by-loan basis.

This logistic regression-based approach has many extremely attractive attributes:

- Default and prepayment are recognized as random events, not events that are predicted with certainty.
- Previously ignored factors such as home prices can be incorporated into both prepayment and default modeling on a consistent basis.

- Loss given default can be recognized as related to default probability by modeling home prices forward and using those home prices as inputs to the default and prepayment probability functions. In a default scenario, the simulated value of the home allows one to derive loss given default after applying appropriate liquidation costs and time lags.

Emerging best practice is to use multinomial logit to deal with two practical problems that arise when independent logistic regressions are used:

- With independent logistic regressions, the modeler is uncertain about which probability to apply first in a given period and a given scenario: the simulation of prepay/don't prepay or the simulation of default/don't default.
- Similarly, it's possible that the probabilities of prepayment and default could add up to more than 100 percent.

In order to deal with these issues, a number of advanced techniques can be employed. One of the most popular techniques is multinomial logit, which allows N potential outcomes instead of the binary 0/1 outcomes assumed in normal logistic regression. Here are a number of potential specifications of the outcomes from a multinomial logistic regression:

Model A Outcomes

- Default
- Prepayment
- Neither of the above

Model B Outcomes

- Default
- Prepayment for interest rate reasons, remaining in the same house
- Prepayment because of relocation
- None of the above

Model C Outcomes

- Default for purely nonmedical financial reasons
- Default triggered by medical expenses
- Prepayment because of divorce
- Prepayment because of interest rate reasons
- Prepayment because of job relocation
- Prepayment because of voluntary relocation
- None of the above

These multiple outcomes offer the analyst and risk manager a much more precise assessment of the impact of interest rates and home prices on the value of a pool of mortgage loans. The multinomial logit coefficients, given the explanatory variables chosen by the analyst, are derived either using joint maximum likelihood estimation or by estimating a logistic regression for each outcome and making appropriate adjustments.

For those institutions that place a priority on correctly understanding risk-adjusted shareholder value creation and business strategy, the logistic and multinomial logistic approaches are an essential step forward. In the next few sections, we compare this best practice approach with legacy approaches to modeling prepayment. With few exceptions, these legacy approaches ignore default risk.

LEGACY PREPAYMENT ANALYSIS OF MORTGAGE-BACKED SECURITIES

The valuation and analysis of mortgages, both fixed and floating rate, are some of the most complex analytical problems in U.S. financial markets. As mentioned previously, much of this difficulty stems from the confusion between analysis of a single mortgage and analysis of a portfolio or pool of mortgages. Douglas Breeden (1994), one of the leading U.S. researchers in finance, former chief executive officer of mortgage fund manager Smith Breeden, and now Dean of the Duke University Business School, commented about the problem as follows:

[M]ortgages are viewed as far too complicated to value precisely and rigorously, even with the Black-Scholes model and the many improvements developed in the eighteen subsequent years.

As we saw in Chapter 28, and as we can confirm using multinomial logit, accurate valuation of mortgage-backed securities and other pools of consumer loans is a complex interaction of credit risk, interest rate levels, and prepayment behavior. The purpose of this section is to take a second look at mortgage-backed securities in the light of the implications of that chapter: that the transactions cost approach promises dramatic improvements in the speed and accuracy of credit-adjusted valuation and risk analysis for consumer loans on both an individual loan level and a portfolio level. Before looking at the implications of that approach, we look first at traditional industry practice.

Legacy Approaches: Prepayment Speeds and the Valuation of Mortgages

Prior to the 2006–2011 credit crisis, Wall Street analysis of mortgage-backed securities relied almost exclusively on prepayment models as a means of analyzing the embedded options in mortgage-backed securities. The prepayment models, in combination with Monte Carlo simulation, produce another number as output—the option-adjusted spread (OAS) on a mortgage-backed security. In the next few sections, we pose a number of questions and attempt to answer them:

- Do prepayment models have an options component?
- Are prepayment speeds predictable enough to use as an input for the option-adjusted valuation of mortgages?
- Is the truth about prepayment speeds obvious enough that sophisticated market participants can reach a consensus about their levels?

Most prepayment models have an objective similar to our objectives throughout this book:

1. Security valuation on a credit-adjusted and option-adjusted basis
2. Measurement of interest rate sensitivity
3. Accurate hedges
4. Guidance on “rich/cheap” analysis for security selection

Prepayment speed models come in five basic varieties:

1. Models in which the prepayment speed is a function of time, but not rate levels. Examples are the lifetime prepayment speed models of the Public Securities Association (PSA), conditional (or constant) prepayment rate (CPR), and single monthly mortality (SMM).
2. Tables in which different prepayment speeds are assigned depending on the remaining maturity of the mortgage, the coupon level, and the current level of interest rates. We call this the *prepayment table approach*.
3. Prepayment functions, which are typically derived from historical data on prepayments and include more inputs than a prepayment table to provide the accurate prepayment speed. Over a long period of time, linear prepayment functions have been the most common functional form used.
4. Logistic regression models, like those we used in Chapter 16, to fit default probability formulas to historical default databases. These logistic regression models are relatively new in prepayment analysis, but offer great promise for understanding the probability of prepayment based on historical behavior of individual borrowers.
5. Multinomial logistic regression, which we described previously as the best practice approach for joint modeling of default and prepayment.

Analytically, the constant prepayment speed models are the simplest. Prepayment tables, which are one step up in complexity, assume that we can accurately forecast prepayment speeds. Prepayment functions assume the future will be like the historical data set that produced the estimates used in the prepayment function. The same is true for logistic regression models to some extent. How well do these models work?

Constant Prepayment Speeds Are Simply a Principal Amortization Assumption

The most important point to make about the PSA, CPR, and SMM analysis is that they have no options component at all. They represent an assumption about principal amortization and make it quite clear that the rate of principal amortization does not depend on the level of interest rates. As a result, there is no rational consumer response to lower rates at all reflected in these three models. A higher prepayment speed lowers the level of principal outstanding. A higher prepayment speed also has the impact of increasing cash flow on the mortgage in the early years and decreasing it in later years, effectively shortening its duration. These cash flows are not stochastic and do not depend on the level of interest rates.

The single-speed approach to prepayment analysis does not work very well. One way many market participants have sought to improve the accuracy of prepayment analysis is by creating a prepayment table that specifies what prepayment rates will be at different times in different interest rate environments. Typically, a prepayment table will contain at least two dimensions (and usually more), one for each payment and another for various levels of refinancing advantage, the spread between the rate on the mortgage loan being analyzed and the current coupon rate on new mortgages. If there are five interest rate tiers and 360 payments, 1,800 input numbers will be required as inputs to the valuation and hedging analysis.

Breeden (1994) takes a very interesting look at Wall Street's ability to forecast prepayment speeds, and we use data from his study liberally in this section. First, he shows that Wall Street typically lacks consensus on the best constant lifetime prepayment speed for a given security. He then shows that the Wall Street estimates of prepayment speeds on a 10 percent FNMA mortgage-backed security maturing in 2018 ranged from 20 percent to 40 percent as of December 1992.

Breeden also shows that the conditional prepayment rates as a function of the coupon on the mortgage-backed security and the refinancing incentive have been very unstable and have been inconsistent with rationality in many cases. Prepayment speeds for high-coupon mortgages have often been lower than prepayment speeds for low-coupon mortgages (a phenomenon called "burnout," shorthand for the fact that the most interest rate-sensitive borrowers in a pool of loans leave the pool first via refinancing, reducing the interest rate sensitivity of the pool over time). In fact, it is easier to guess the value of an embedded prepayment option directly than it is to guess the prepayment rates needed as inputs to a formula to calculate option values.

LEGACY APPROACHES: OPTION-ADJUSTED SPREAD

Given the difficulties with prepayment models in general, what are the implications for the option-adjusted spread quotations common on Wall Street? First, we should summarize the OAS procedures, for OAS refers more to the procedures used than it does to the true concept of the risky spread on mortgages. Typically, analysts do the following:

- They assume a prepayment function or prepayment table is true.
- They assume that Monte Carlo simulation is a good method for valuing an American option to prepay, ignoring the concerns voiced in Chapter 27.
- They solve for the spread over a risk-free interest rate, which provides a calculated value for the mortgage-backed security that matches the observable value.

In addition to the comments in Chapter 27 about the fundamental problems with the use of Monte Carlo simulation in the valuation of American option-related securities, OAS analysis typically incorporates these errors:

1. The refinancing incentive is not measured on a matched maturity basis.

Within the prepayment table, the refinancing incentive in the current market is measured in order to choose the correct prepayment rate from the prepayment table or prepayment function. This is measured as the difference between the rate

on the mortgage being analyzed and the rate on a new 30-year fixed rate mortgage. This sounds good at first glance, but a closer look gives cause for serious concern. If there are five years remaining on the mortgage, we are comparing its five-year rate with the 30-year loan and assuming that refinancing is done on the basis of an “apples-to-oranges” comparison. In a swap trading unit or an investment management firm, anyone who said 30-year bonds are better than five-year bonds because their yield is higher would lose their job. This sort of comparison, however, is embedded in almost all OAS analysis. Why? Historical data on mortgage rates includes only rates at new issue maturities (typically 15 and 30 years in the United States) and the analysts had no other data to measure refinancing spread. The reason for the choice of measure is clear, but it is an important source of inaccuracy.

2. The refinancing incentive is measured on the basis of a 30-year mortgage rate, which does not include a premium for the call provisions of the loan.

Within the Monte Carlo simulation routine, there is no capability to calculate how much the premium should be on a new mortgage that is callable compared to a hypothetical noncall mortgage. Why? You have a “chicken-and-egg problem”—using Monte Carlo and a prepayment table, we need to know what the option is worth in order to calculate what the option is worth! All of the refinancing incentives measured within the Monte Carlo routine are normally done assuming the mortgage is noncallable, introducing another source of error.

3. The prepayment table should be path-dependent, but it is not.

If rates start at 6 percent, jump to 12 percent for 10 years, and then fall to 4 percent, the prepayment function and prepayment tables will produce the same prepayment speed prediction as if rates started at 6 percent, fell to 3 percent for 10 years, and then jumped to 4 percent. Clearly, that should not be the case. Monte Carlo is often relied upon on the rationale that prepayments are path dependent. We discuss that comment in more detail below, but it is clear Monte Carlo replaces a path-dependence problem in the prepayment rates themselves with a path-dependence problem in the prepayment table.

4. The OAS will match the market value of the observable security, but not its interest rate risk.

Why? For the same reason that the single prepayment speed model did not work. We have two equations (value and rate sensitivity) we want to solve simultaneously, but only OAS to adjust in order to do this. We can match both only by accident.

OAS is the “plug” or balancing number that offsets all of the errors inherent in Monte Carlo analysis and the four problems above. It is not surprising, given the lack of consensus on Wall Street about prepayment rates, that median broker estimates of OAS show very strange patterns when plotted by coupon level. Breeden (1994) cites data for the third quarter of 1989 and the fourth quarter of 1992, which show investors could double their spread theoretical OAS by simply buying an MBS with a different coupon level, even though this discrepancy is public information widely disseminated by Wall Street firms. In light of the 2006–2011 credit crisis and the drop in home prices, traditional prepayment models have gone seriously “off the tracks” and a repeat of the Breeden study would show that even more serious errors exist now.

The fact that this difference in OAS was not arbitrated away is due to the market's view that the OAS numbers were not accurate and that the spread differential was not real. Wall Street did not bid up the price of the MBS with the widest reported OAS relative to the MBS with the lowest MBS. Clearly, OAS numbers are taken with a grain of salt by market participants. Gregg Patruno (1994) of Goldman, Sachs & Co. summarizes a common view, "The standard measure of mortgage relative value . . . OAS does not account for the complete nature of prepayment risk. Rather, it adjusts for the optionality due to interest rate fluctuations under the presumption that prepayment rates are a known, permanent function of interest rates" (53).

Patruno also suggests both the reason for the problems with OAS and a potential solution to the problem: "Homeowners are willing to refinance if the financial incentive is high enough to meet their requirements. We use the measure of refinancing incentive actually considered by typical homeowners and their bankers—not an abstract interest differential or ad hoc statistical artifact, but the real dollar savings expected on an after tax basis, taking into account an appropriate mix of available mortgage rates and points" (47). In other words, the only way to effectively model consumer prepayment behavior is to consider the rational value of the option to prepay. Patruno concludes that "the levels of current mortgage rates and transactions costs determine what fraction of the homeowners in the distribution should be considered willing to refinance" (47). The credit crisis also made it clear that the value of the home, relative to the principal on the loan, had a major impact on the ability of a borrower to prepay.

For these reasons, the authors believe that the transactions cost approach of Chapter 28, as analyzed by McConnell and Singh (1994) and Stanton (1995), supplemented by multinomial logistic regression for prepayment and default analysis, represents best practice.

IMPLICATIONS FOR OAV SPREAD, CMOs, AND ARMs

The primary reason for the lack of OAS consensus among Wall Street firms is that all of the assumptions being made are assumptions about the borrower. It is no wonder that consensus is impossible. Instead, we recommend the option-adjusted value (OAV) spread. Let us consider a corporate treasurer who has to choose between issuing a seven-year noncall bond or a 10-year bond that is noncallable for seven years. We modeled this structure in Chapter 27. He makes the decision in a straightforward way. First, he asks for the rate on a comparable Treasury seven-year noncall bond and calculates his spread over Treasuries. Second, he calculates (since no observable Treasury is a "10 years noncall seven" bond) what the coupon would be on a 10-year noncall seven Treasury bond that trades at par value, using the techniques of Chapter 27. His spread on the 10-year noncall seven bond issue is the difference between his coupon and the Treasury's coupon. He issues the bond with the lowest spread.

Collateralized mortgage obligations (CMOs) and adjustable rate mortgages (ARM) can be analyzed using the HJM approach, like the example in Chapter 9, as a variation on our results in Chapter 28.

LOGISTIC REGRESSION, CREDIT RISK, AND PREPAYMENT

Logistic regression is another powerful tool that we are familiar with from its use in modeling default probabilities for retail, small business, and large corporations using a historical database of defaults. In a similar way, we can use logistic regression to model the probability of prepayment. With this technique we can combine variables like the following, as explanatory variables for prepayment:

- The net present value benefit of prepaying as of the current observation (which captures the interest rate benefits of prepayment)
- The current value of the house relative to the principal balance of the outstanding loan
- The age of the loan
- The months remaining before maturity
- The month (to capture seasonality)
- Direct transactions costs of refinancing (points, fees, etc.)
- The default probability of the borrower, as captured directly by internal models or by some kind of internal credit rating
- Any other relevant consumer credit bureau approach

The coefficients of this logistic regression essentially let us derive the transactions costs that we implied previously. They also automate the adjustments for seasonality that we would also have to make in the approach of the previous section. The result is an extremely powerful integration of credit risk, interest rate risk, and prepayment risk that significantly advances the accuracy of our ability to model the Jarrow-Merton put option, our integrated measure of risk for the full balance sheet. Multinomial logistic regression is best practice because it gives integrated probabilities for default and prepayment. For those making the transition from legacy techniques, however, the logistic regression approach is a nice intermediate step on the way to a full multinomial logistic implementation.

MORTGAGE-SERVICING RIGHTS: THE ULTIMATE STRUCTURED PRODUCT²

One of the lessons of the 2006–2011 credit crisis of was that the conventional wisdom can be a dangerous thing. The conventional wisdom in early 2006 said that “home prices in the United States do not go down” and that the copula method was an accurate method for valuing tranches of collateralized debt obligations. Another complex security valuation issue is becoming increasingly important: the valuation of mortgage-servicing rights. Exhibit 29.1 shows the reported valuations for mortgage-servicing rights for eight large U.S. financial services firms between December 31, 2009, and September 30, 2011. The aggregate reported value of mortgage-servicing rights has declined from \$62.5 billion to \$35.6 billion in 21 months. Even at the lower valuation, the very large absolute value attributed to mortgage-servicing rights and the tremendous volatility of the valuation in the current market makes an examination of the conventional wisdom essential from both a risk management and corporate governance point of view.

EXHIBIT 29.1 Kamakura Corporation: Reported Valuations of Mortgage-Servicing Rights (USD millions)

Mortgage Servicer	December 31, 2009	December 31, 2010	September 30, 2011
Bank of America	19,465	14,900	7,880
Wells Fargo	17,123	15,886	13,769
JP Morgan Chase	15,531	13,649	7,833
Citigroup	6,530	4,554	2,852
US Bancorp	1,749	1,837	1,466
Sun Trust	936	1,439	1,033
BB&T	832	830	573
First Horizon	303	207	151
Total	62,469	53,302	35,557

Source: Kamakura Corporation, SEC filings.

One of the things that these figures make obvious is that the mortgage-servicing business is an enormously concentrated business with highly skewed market shares. Another obvious point is that the percentage of total market capitalization for these firms that is attributable to mortgage-servicing rights is very large. Mortgage-servicing rights are classified as a Level 3 asset for accounting purposes because there are very few observable traded prices and because many of the inputs to a reasonable valuation are also unobservable and the result, in the end, of educated guesses. If there is model risk in reported mortgage-servicing rights valuation, equity market participants need to be informed of these model risks, and those hedging mortgage-servicing rights need to take great care because of potential errors in both the amount and *direction* of the hedging needed.

AN INTRODUCTION TO THE VALUATION OF MORTGAGE-SERVICING RIGHTS

As is common in the risk management business, the head of a mortgage-servicing rights hedging program once said “I don’t care about accuracy, I just want to do the same MSR calculation as everybody else.” Therein lies the problem. What if “everybody else” is just seven firms? How can they be sure that there are no errors in this complex valuation calculation? One approach is to survey all of the other participants in the market, which is commonly done, to insure consistency. Unfortunately, the assumptions that home prices do not go down and that copulas are appropriate for CDO valuation were consistent modeling assumptions in U.S. financial markets in 2006. The consistency alone is not sufficient to establish accuracy. The recent controversy about the manipulation of the U.S. dollar LIBOR market during the credit crisis, when the U.S. dollar LIBOR panel consisted of 16 firms, illustrates another risk: the risk that the eight firms who dominate the mortgage-servicing business have an economic incentive to agree on a valuation methodology that maximizes the value of this large asset. For all of these reasons, an independent approach of maximum accuracy in valuing mortgage-servicing rights is essential.

Starting from the ground up, one can approach the valuation of mortgage-servicing rights as the valuation of a fixed income (broadly defined) security subject to default risk and prepayment risk. Mortgage-servicing rights, like all of the other securities valued in prior chapters, can be valued in the HJM framework.³

In the next section, we follow Lin, Chu, and Prather (2006) and decompose mortgage-servicing rights into their component cash flows and compare the “conventional wisdom” and best practice.

COMPARING BEST PRACTICE AND COMMON PRACTICE IN VALUING AND HEDGING MORTGAGE-SERVICING RIGHTS

There are so many differences between best practice and common practice in valuing and hedging mortgage-servicing rights that we will have to be brief to keep this section of reasonable length. We address these differences point by point.

Valuation Yield Curve for Cash Flows

Best practice discounts cash flows using a series of yield curves, depending on the nature of the cash flows, particularly when some cash flows are best valued in nominal dollars and some are best valued in real dollars. Common practice discounts all cash flows at the same yield curve, typically the U.S. dollar LIBOR-swap curve. The use of this curve is usually justified with the phrase “all mortgage market participants use the swap curve for hedging.” Unfortunately, that has nothing to do with whether the curve is appropriate for valuation. As we shall see, the right valuation curve depends on the characteristics of the mortgage-servicing right (MSR) cash flow component’s credit risk, prepayment risk, and inflation sensitivity (real or nominal). The LIBOR-swap curve has a number of serious problems when employed for valuation. We list many of those problems here:

1. As the fines of \$450 million levied on Barclays in 2012 for LIBOR manipulation confirm, the LIBOR curve has been manipulated and is no longer a legitimate tool for hedging.
2. Market participants commonly assume that fixed rate mortgage yields are a simple spread over an interpolated swap yield at two key maturities, such as the 2-year and 10-year interpolated swap rate at an effective maturity of six years. This assumption is false. Over the period from January 8, 2004, to November 17, 2011, using weekly data, the appeal of using swap rates for valuation is deceptively attractive, as shown by the high correlations of swap yields (all over 86 percent) with the 30-year conventional fixed rate mortgage rate. The true relationship between swap and Treasury yields and fixed rate mortgage yields is more complex, as we discuss next.
3. Using conventional 30-year fixed rate conventional mortgage reported on the Federal Reserve H15 statistical release, the interpolated 2- and 10-year swap rates explain only 82.8 percent of the weekly variation in fixed rate mortgage yields with a standard error of 30.47 basis points.

By contrast, a simple linear regression using multiple points (13) on both the swap curve and the U.S. Treasury curve explains 98.17 percent of the weekly

variation in fixed rate mortgage yields with a standard error only one-third as big as the conventional wisdom, 9.9 basis points. The conventional wisdom, to benchmark valuation off of two points of the LIBOR-swap curve or off of the entire swap curve alone, has no basis in the facts. The U.S. Treasury curve plays at least as great a role in the movements of the 30-year fixed rate mortgage yield as does the LIBOR-swap curve and there is much less risk of manipulation by market participants.

4. The spread between the 30-year swap yield and the matched maturity Treasury yield has been consistently negative since October 24, 2008. The reason for this negative spread is the difference in credit exposures in the two yield curves. A default by the U.S. government, which is no longer an unthinkable possibility, puts potentially the entire principal amount at risk. If there is a default by a swap counterparty, the amount at risk is, at a maximum, the change in the present value of the swap position since its value at origination (which would be zero for a swap done “at market” at the time). With the increase in the perceived risk of the U.S. government, the spread between the swap yield at 30 years and the 30-year U.S. Treasury yield has been negative on 749 business days from February 9, 2006, to November 17, 2011. Over the same period, the 10-year spread between the swap curve and the U.S. Treasury curve has been negative for 67 days, with the first negative spread reported on January 21, 2009.

It is unthinkable for market participants to value mortgage-servicing rights by discounting at yields below U.S. Treasury yields given the current relative default risk of mortgages and the U.S. Treasury. There is no legitimate argument that justifies such a valuation choice. The graph in Exhibit 29.2 from December 2, 2011, plots the zero-coupon yields, using advanced yield curve smoothing techniques, for the swap curve versus the U.S. Treasury curve. The negative spread at 30 years contaminates the spread at much shorter maturities (on this date, maturities 13 years and over), so the choice of the swap curve as a valuation yield curve is a dramatic error. The short end of the curve shows a hump in discount rates for the LIBOR-swap curve because the credit risk of one-, three-, and six-month LIBOR is incompatible with the credit risk of the swap yields used for one-year and longer maturities, as discussed above.

Simulation of Random Movements in Yields

The conventional wisdom for simulating the random interest rate environment when valuing mortgage-servicing rights is embedded in legacy asset and liability management systems. For the most part, these systems rely only on one-factor models of the term structure. As commonly implemented, these one-factor models imply that changes in yields can only (1) all rise together, (2) all fall together, or (3) all remain unchanged. Unfortunately, this implication of one-factor term structure models is false as we discussed at length in Chapter 3. With respect to the simulation of random interest rates, the conventional wisdom is dramatically wrong and is in violation of the Basel Committee on Banking Supervision’s December 31, 2010, Revised Basel II Framework for Market Risk, which requires on page 12 that

For material exposures to interest rate movements in the major currencies and markets, banks must model the yield curve using a minimum of six risk factors.




EXHIBIT 29.2 U.S. Treasury and USD Interest Rate Swap Zero-Coupon Yields Derived from the Federal Reserve H15 Statistical Release Using Maximum Smoothness Forward Rate Smoothing

Sources: Kamakura Corporation; Federal Reserve.

Because the conventional wisdom assumes away the yield curve twists that are very heavy influences on mortgage prepayments and defaults, the resulting valuations could potentially be grossly incorrect.

The Role of Home Prices in Defaults and Prepayments

The conventional wisdom in the servicing of a whole loan portfolio is to assume a constant default rate that is not influenced by home prices or interest rates. This assumption is false, and that is well recognized by many. They are constrained by legacy asset and liability systems to use a valuation assumption that they know to be false. Default rates are very sensitive to home prices. A recent study by Kamakura Corporation of the mortgage default rate reported by the U.S. Mortgage Bankers Association found that the change in home prices is statistically significant in predicting defaults, with a t -score equivalent well above the t -score level of 2 that is normally the benchmark for judging statistical significance. Modern enterprise risk systems allow mortgage-servicing rights valuation with default rates that are determined by the simulated random movements in home prices, interest rates, and other key factors. The erroneous assumption of constant default rates is avoided.

Similarly, it is very common in the valuation of mortgage-servicing rights to assume a constant, linear, or tabular prepayment function in which home prices play no role. This assumption is also false. Recent publications by Calhoun and Deng

(2002) and Campbell and Cocco (2011) are among many research papers that document the very significant role of home prices in prepayment as well.

Other Sources of Cash Flow Related to Mortgage-Servicing Rights

The total cash flow from mortgage-servicing rights can be broken into components as follows:

- Net servicing fee, equal to the gross servicing fee less guarantee fee, if any
- Net cost to service
- Net cash flow triggered by delinquency and default
- Float on taxes and insurance
- Float on principal and interest
- Float and costs from prepayments

The magnitude and very existence of these cash flows depends on the probabilities of default, partial prepayment, and prepayment in full, which have been misspecified in the conventional approach to valuing mortgage-servicing rights as explained previously.

There are many other significant errors that can be made in valuing individual components of cash flow. The cost of service, for example, might be assumed to be constant, but that assumption is false. The cost of service is affected both by inflation and by the economies of scale in servicing relative to other competitors. The float on taxes and insurance could be calculated on the assumption that taxes and insurance are constant, but that assumption is also false. Both taxes and the cost of insurance are strongly linked to the value of the house. One might assume, as a better approximation, that home prices are correlated with inflation. Unfortunately, looking at correlation between the Case-Shiller home price indices and the consumer price index shows a wild variation in the linkages.⁴

The float on principal and interest and the float and costs from prepayments are complex derivatives that depend not just on interest rates but also on home prices. The conventional wisdom ignores the home price linkage. Stated more precisely, analysts recognize the linkages but they are constrained by legacy systems that are unable to take advantage of this economic reality. A modern state-of-the-art enterprise risk management system avoids this valuation error.

Incorrect Hedging of Mortgage-Servicing Rights

The conventional wisdom notes that mortgage-servicing rights are to be reported at market value under generally accepted accounting principles. The corollary to that fact, however, is wrong—the conventional wisdom concludes that, since these market values are random, mortgage-servicing rights need to be hedged. That conclusion is potentially wildly incorrect, for a simple reason. Market participants have simply ignored the “factory” that services mortgages (other than assuming the constant dollar cost above). This factory is not marked to market on the balance sheet under GAAP, but, in fact, the value of the operations center (buildings, software, hardware, staff, and so on) rises and falls with the value of mortgage-servicing rights in a highly correlated way. By ignoring this link and partial or full offset, those who hedge may

in fact be *increasing risk*, not reducing it. Take the example of a scenario like that of 2006 to 2011, when U.S. home prices have fallen sharply. This reduces the value of mortgage-servicing rights (because default rates have increased), but it also has reduced the cost of being a participant in the operational processing of mortgages. For example, if all of the mortgages being serviced defaulted, what would happen? The present value of labor costs of mortgage servicing would essentially go to zero, because all staff would be laid off. The present value of computer hardware needs would go to zero, because the hardware used would be sold or moved to other businesses. The present value cost of software used would go to zero (in the case of annual rental pricing, since the contract would be terminated) or be sharply reduced (in the case of perpetual license pricing, since the maintenance fee would be terminated). If the mortgage-servicing business were terminated, the value of the related buildings and land would go to their value of the next best use.

In short, there is a natural offset of some or all of the variation in the value of mortgage-servicing rights. The conventional wisdom, however, has completely forgotten this simply because the accountants treat mortgage-servicing rights and the physical and intellectual property of the operations center differently. This is a very common source of risk management disasters, in which risk managers pay attention to the accountants and ignore the true economics of what they are trying to hedge.

CONCLUSION

Among mortgage-backed securities broadly defined, the valuation of mortgage-servicing rights is both the most complex and the best illustration of why best practices need to be employed in valuing mortgage-backed securities in general. The conventional wisdom in the valuation of mortgage-servicing is notoriously sticky because there are so few market participants in the business and because reported values are so large. There is no one inside the mortgage-servicing industry with a vested interest in pointing out that the emperor's new clothes may be a bit threadbare. Only the accountants and regulators who would have to deal with any markdowns in value have downside risk. The shareholders and taxpayers who would ultimately have to pay were one of the large servicing banks to fail are not given sufficient information to make their own judgments about the value of mortgage-servicing rights. As pointed out earlier, tens of billions of dollars in value relies on a series of assumptions that are known to be false.

We now turn to another major modeling issue in enterprise risk management and the correct valuation of the Jarrow-Merton put option on the firm: the valuation of nonmaturity deposits.

NOTES

1. Two key publications are essential reading for investors who own or are considering an investment in securitized transactions:

- United States Senate, Special Committee on Investigations, *Wall Street and the Financial Crisis: Anatomy of a Financial Collapse* ["The Levin Report"], Majority

and Minority Staff Report, April 13, 2011 (discussed in Chapter 18 and pictured in Exhibit 18.9).

- Michael Lewis, *The Big Short: Inside the Doomsday Machine* (New York: Thorndike Press, 2010).

Both of these exceptionally detailed books describe exactly how the process of securitization is designed to hide the flaws of the ultimate collateral in a structured product.

2. The authors would like to thank our colleague Mark Slattery for his helpful comments on this section.
3. Che-chun Lin, Ting-heng Chu, and Larry J. Prather, “Valuation of Mortgage Servicing Rights with Foreclosure Delay and Forbearance Allowed,” *Review of Quantitative Finance and Accounting*, 26 (2006): 41--54 provides an introduction to this approach.
4. Documented in “Risk Management Strategies for Individual Investors, Part 4: Home Prices and Inflation,” Kamakura blog, June 1, 2009, www.kamakuraco.com.

Nonmaturity Deposits

In this chapter, we turn our attention to the liability side of the balance sheet of financial institutions, in particular to deposits of commercial banks. This is particularly appropriate in light of massive deposit runoffs in the credit 2006–2011 crisis, which we document later in this chapter and in Chapter 37 on liquidity risk. This focus on deposits is also critical to correct valuation of the Jarrow-Merton put option as the best comprehensive measure of integrated credit risk and interest rate risk in the commercial banking sector. Why? Because we need to value a put option that includes both assets and liabilities of the financial institution because some of the liabilities may be providing a hedge of the assets. One of the advantages of the reduced form credit modeling technology of Chapter 16 is that it can handle the complex liability structure that is typical of large financial institutions, whereas the Merton model of Chapter 18 assumes a single zero-coupon bond as a liability. In fact, it is the multiperiod cash flows of nonmaturity deposits that explained much of the drama of the credit crisis.

The bulk of bank deposits and insurance liabilities are made up of securities with explicit maturities, although they are often putable (back to the bank or insurance company) by the consumer who supplied the funds, whether they come in the form of a time deposit or life insurance policy. These liabilities can be analyzed with the techniques of Chapters 27 and 28, even if the consumer's right to put the security back to the financial institution is exercised irrationally. A substantial portion, however, of the liabilities of major banks worldwide consists of nonmaturity deposits. Similarly, the funds supplied to mutual fund managers of fixed income funds also have no specific maturities. This chapter is relevant to both types of liabilities, but we will refer to them as nonmaturity deposits within a banking industry context from here on. These deposits have no specific maturity, and individual depositors can freely add or subtract balances as they wish. Most importantly, there is no way to distinguish new clients from old clients because the same interest rates are paid to both, unlike a certificate of deposit where the rates are changing constantly for new clients. The interest rate on nonmaturity deposits is usually, but not always, a function of open market interest rates. Similarly, the level of deposit balances in aggregate often moves in sympathy with open market interest rates. With the strong trend in mark-to-market-based risk management in banking, all bankers are faced with the difficult task of calculating the mark-to-market value of these nonmaturity deposits. This chapter provides an introduction to that difficult topic.

In September 1993, the Federal Reserve Board (1993) published proposed regulations for interest rate risk and capital adequacy. In its regulations, the Federal Reserve Board noted “the inherent difficulties in determining the appropriate treatment of non-maturity deposits.” In this chapter, we use the Heath, Jarrow, and Morton (HJM) interest rate framework of Chapters 6 through 9 to measure the mark-to-market value of deposits. We do this in a way that is theoretically rigorous and yet easy to apply in practice. Most of the arguments that are applied here to the valuation of nonmaturity deposits can be applied equally well to the valuation of charge card loans, where again the balance as well as the rate on the loan fluctuates in response to market rate movements.¹

THE VALUE OF THE DEPOSIT FRANCHISE

When we measure the value of deposits, what is it that we are trying to measure? Here “value” means the present value of future costs. If the bank had issued a bond to finance its assets, the value of the bond is its present value—the present value of future payments that the bank must make on the security that it has issued.

In the case of nonmaturity deposits, the value of the deposit could also mean the premium that a third party would be willing to pay above the face value of the deposits to assume ownership of the bank’s deposit franchise.² In this case, value means the net present value benefit of owning the deposit franchise—the value of cash provided by depositors less the present value cost of the deposit franchise.

In what follows, we use the word “value” to refer to the present value of the cost of the deposit franchise, exactly in keeping with the analogy of the bank as bond issuer. The net present value benefit of having the deposit franchise can be calculated from the valuation formulas presented below by subtracting the values we calculate from the face value of the deposit. For example, the present value cost of \$1,000 face value of savings deposits might be \$650. The net present value benefit of the deposit franchise is $\$1,000 - \$650$ or \$350. All references to value below refer to the present value of the cost of the franchise, \$650.

We also initially assume that the deposit franchise is guaranteed by a riskless third party, which is the case (in theory) in the United States and many other countries. This allows us to avoid dealing with the complexities of bankruptcy risk. We deal with that explicit risk later in the chapter.

Some bankers prefer to focus on the value of specific nonmaturity deposits owned by current clients; that is, valuing only deposit accounts on the books of the bank today and ignoring future business. This orientation has a logical basis because we don’t value the franchise of any other business line on the bank balance sheet.

There is a good reason, however, for looking at the franchise for valuation. The pricing of nonmaturity deposits is exactly the same for new customers as it is for old customers. The bank cannot control the flow of new demand deposit accounts or savings deposit accounts by changing pricing to attract new clients without immediately making the same price change on existing accounts. This is not the case for auto loans or consumer certificates of deposit or any other account on the balance sheet of the bank except for the charge card business. For this reason, we think the focus on the franchise is the right one, but we recognize that some will disagree.

TOTAL CASH FLOW OF NONMATURITY DEPOSITS

The first step in valuing nonmaturity deposits is to isolate the total cash flows from the deposit franchise. Total costs consist of the following elements:

- Total deposit cost cash flows in a given time period
- = Interest paid on the deposit
- + Noninterest expense of servicing the deposit
- Noninterest revenue from the deposit franchise
- *Net increase in deposit balances*

Many readers will be surprised to see the net increase in deposit balances subtracted from costs in the equation above. The net change in deposit balances on many deposit products is the single most important cash flow in the valuation process, and yet many analysts ignore this important determinant of value.

In the valuation of any security, total cash flow is the basic building block of valuation. For most securities, the maturity date is fixed and cash flow from principal stems from a predetermined payment schedule associated with that security. In the case of the nonmaturity deposit, however, the security is a perpetual one and principal is never “returned.” Instead, changes in deposit balances are simply another source of cash flow. A simple example illustrates the point.

Let us consider a deposit category with balances of \$1,000 where the sum of interest expense and processing costs is 2 percent of balances annually. Let’s assume that this cash flow occurs at the end of every year. Let’s also assume that the balances on this account grow at 2 percent a year, and that balances also increase in a one-time jump at the end of each year. What is the net cash flow associated with this account? On December 31 of each year, 2 percent of the deposit amount flows out in the form of interest expense and processing costs, and it flows back in on the same day in the form of increased deposit balances. The net cash flow from the deposit franchise is zero every year forever.

The present value cost of this deposit is zero because there is never net cash inflow or outflow. Since the present value cost is zero, the net present value or premium that would be paid to gain this deposit franchise is \$1,000, calculated as above: The deposit balance \$1,000 minus the present value of deposit costs (\$0) is \$1,000.

This example is deceptively simple: it is consistent with common sense, and yet it will lead us to some surprising conclusions about the value of nonmaturity deposits. We will analyze the value of nonmaturity deposits in what follows using realistic assumptions about random deposit rate and balance behavior. This is covered in more detail in Jarrow and van Deventer (1996, 1998) and Janosi, Jarrow, and Zullo (1999), which provide explicit closed form solutions for a commonly used one-factor term structure model. Step-by-step examples using a one-factor term structure model can also be found in the first edition of this book. In this volume, however, consistent with our conclusions in Chapter 3, we believe that a one-factor model is not accurate enough to meet the current twenty-first century best practice standard. In this chapter, we continue to use the three-factor HJM framework of Chapter 9.

As in Chapters 19 through 29, we use the probability-weighted discount factors from Chapter 9 to value nonmaturity deposits. Our task in this section is to calculate the cash flows on each node of the bushy tree (or in each scenario of an HJM-based Monte Carlo simulation) that stem from the nonmaturity deposit cash flows.

SPECIFYING THE RATE AND BALANCE MOVEMENT FORMULAS

Step 1 in the valuation process is to specify how nonmaturity deposit rates, noninterest income, and noninterest expense vary with interest rates and other macroeconomic factors. Step 2 is to specify a similar formula for the variation in deposit balances. In this example, we ignore the credit risk of the institution. Later in this chapter, we explicitly incorporate the risk of the bank in light of actual experiences of “too big to fail” banks during the 2006–2011 credit crisis.

Model Risk Alert

It is very important for a bank that has never had credit problems previously to realize that its own data is “too good” for fitting the rate and balance relationships. Accurate relationships can only be based on data from both “good” and “bad” banks; if we do not fit the relationships in this way, they cannot possibly capture the behavior of nonmaturity deposits at a distressed bank in an accurate way.

Let’s assume we compile such a data set and we derive a relationship for the sum of deposit interest plus noninterest expense less noninterest revenue for the bank in question. Assume that the effective deposit rate is as follows:

$$\text{Effective deposit rate} = 0.10\% + 0.2(\text{Lagged one-period spot rate})$$

We are implicitly assuming that the effective deposit interest at time $T = k$ is the effective rate at time $T = k$ (which depends on the spot rate at time $T = k - 1$). We base this calculation on the one-period spot rates from Chapter 9 (Exhibit 30.1).

Using the effective deposit rate formula above, we calculate the effective deposit rate at each node of the bushy tree. The rates we derive are shown in Exhibit 30.2.

In order to calculate the dollars of effective interest paid, we need to model deposit balances. With all of the caveats above still in mind, we derive a deposit balance equation that has the following terms:

$$\begin{aligned} \text{Deposit balance} = & 1.005(\text{Lagged deposit balance}) \\ & - 200(\text{Lagged one-period spot rate}) \end{aligned}$$

This gives the deposit balances for each node of the bushy tree in Exhibit 30.3, assuming our time 0 initial balance is \$1,000.

The effective dollar amount of deposit interest (which includes noninterest expense less noninterest income on deposits) is the product of the rates in Exhibit 30.2 and the balances in Exhibit 30.3 for each node of the tree. This dollar cash outflow is shown in Exhibit 30.4 as a positive number, just like interest on a bond.

EXHIBIT 30.1 Spot Rate Process

State	Row Number	0	1	2	3
S-1, S-1, S-1	1	0.3003%	2.1653%	4.6388%	6.1106%
S-1, S-1, S-2	2		0.6911%	1.4142%	5.5667%
S-1, S-1, S-3	3		1.0972%	2.6089%	5.1710%
S-1, S-1, S-4	4		1.3611%	4.9179%	8.0375%
S-1, S-2, S-1	5			2.2496%	3.1041%
S-1, S-2, S-2	6			0.5984%	2.5913%
S-1, S-2, S-3	7			0.8121%	2.0291%
S-1, S-2, S-4	8			1.2459%	4.0628%
S-1, S-3, S-1	9			1.9217%	5.2210%
S-1, S-3, S-2	10			1.4149%	1.9784%
S-1, S-3, S-3	11			0.8591%	3.1798%
S-1, S-3, S-4	12			2.8695%	5.5017%
S-1, S-4, S-1	13			2.2088%	5.4252%
S-1, S-4, S-2	14			1.7005%	4.8848%
S-1, S-4, S-3	15			1.1432%	4.4917%
S-1, S-4, S-4	16			3.1593%	7.3396%
S-2, S-1, S-1	17				4.1517%
S-2, S-1, S-2	18				0.9421%
S-2, S-1, S-3	19				2.1313%
S-2, S-1, S-4	20				4.4295%
S-2, S-2, S-1	21				2.2036%
S-2, S-2, S-2	22				0.5531%
S-2, S-2, S-3	23				0.7667%
S-2, S-2, S-4	24				1.2004%
S-2, S-3, S-1	25				2.6231%
S-2, S-3, S-2	26				−0.3734%
S-2, S-3, S-3	27				0.2593%
S-2, S-3, S-4	28				2.1679%
S-2, S-4, S-1	29				1.7764%
S-2, S-4, S-2	30				1.2703%
S-2, S-4, S-3	31				0.7153%
S-2, S-4, S-4	32				2.7228%
S-3, S-1, S-1	33				2.2804%
S-3, S-1, S-2	34				1.7718%
S-3, S-1, S-3	35				1.2141%
S-3, S-1, S-4	36				3.2315%
S-3, S-2, S-1	37				2.2804%
S-3, S-2, S-2	38				1.7718%
S-3, S-2, S-3	39				1.2141%
S-3, S-2, S-4	40				3.2315%
S-3, S-3, S-1	41				3.6354%
S-3, S-3, S-2	42				0.6093%
S-3, S-3, S-3	43				1.2483%
S-3, S-3, S-4	44				3.1757%

EXHIBIT 30.1 (Continued)

State	Row Number	0	1	2	3
S-3, S-4, S-1	45				4.5070%
S-3, S-4, S-2	46				1.2864%
S-3, S-4, S-3	47				2.4797%
S-3, S-4, S-4	48				4.7858%
S-4, S-1, S-1	49				3.5013%
S-4, S-1, S-2	50				0.3117%
S-4, S-1, S-3	51				1.4934%
S-4, S-1, S-4	52				3.7773%
S-4, S-2, S-1	53				2.5831%
S-4, S-2, S-2	54				2.0730%
S-4, S-2, S-3	55				1.5136%
S-4, S-2, S-4	56				3.5370%
S-4, S-3, S-1	57				2.5831%
S-4, S-3, S-2	58				2.0730%
S-4, S-3, S-3	59				1.5136%
S-4, S-3, S-4	60				3.5370%
S-4, S-4, S-1	61				4.5061%
S-4, S-4, S-2	62				2.2686%
S-4, S-4, S-3	63				2.7252%
S-4, S-4, S-4	64				5.0230%

EXHIBIT 30.2 Effective Deposit Rate Paid

State	Row Number	0	1	2	3	4
S-1, S-1, S-1	1	0.1601%	0.5331%	1.0278%	1.3221%	
S-1, S-1, S-2	2	0.1601%	0.5331%	1.0278%	1.2133%	
S-1, S-1, S-3	3	0.1601%	0.5331%	1.0278%	1.1342%	
S-1, S-1, S-4	4	0.1601%	0.5331%	1.0278%	1.7075%	
S-1, S-2, S-1	5		0.2382%	0.3828%	0.7208%	
S-1, S-2, S-2	6		0.2382%	0.3828%	0.6183%	
S-1, S-2, S-3	7		0.2382%	0.3828%	0.5058%	
S-1, S-2, S-4	8		0.2382%	0.3828%	0.9126%	
S-1, S-3, S-1	9		0.3194%	0.6218%	1.1442%	
S-1, S-3, S-2	10		0.3194%	0.6218%	0.4957%	
S-1, S-3, S-3	11		0.3194%	0.6218%	0.7360%	
S-1, S-3, S-4	12		0.3194%	0.6218%	1.2003%	
S-1, S-4, S-1	13		0.3722%	1.0836%	1.1850%	
S-1, S-4, S-2	14		0.3722%	1.0836%	1.0770%	
S-1, S-4, S-3	15		0.3722%	1.0836%	0.9983%	
S-1, S-4, S-4	16		0.3722%	1.0836%	1.5679%	
S-2, S-1, S-1	17			0.5499%	0.9303%	
S-2, S-1, S-2	18			0.5499%	0.2884%	
S-2, S-1, S-3	19			0.5499%	0.5263%	

(Continued)

EXHIBIT 30.2 *(Continued)*

State	Row Number	0	1	2	3	4
S-2, S-1, S-4	20				0.5499%	0.9859%
S-2, S-2, S-1	21				0.2197%	0.5407%
S-2, S-2, S-2	22				0.2197%	0.2106%
S-2, S-2, S-3	23				0.2197%	0.2533%
S-2, S-2, S-4	24				0.2197%	0.3401%
S-2, S-3, S-1	25				0.2624%	0.6246%
S-2, S-3, S-2	26				0.2624%	0.0253%
S-2, S-3, S-3	27				0.2624%	0.1519%
S-2, S-3, S-4	28				0.2624%	0.5336%
S-2, S-4, S-1	29				0.3492%	0.4553%
S-2, S-4, S-2	30				0.3492%	0.3541%
S-2, S-4, S-3	31				0.3492%	0.2431%
S-2, S-4, S-4	32				0.3492%	0.6446%
S-3, S-1, S-1	33				0.4843%	0.5561%
S-3, S-1, S-2	34				0.4843%	0.4544%
S-3, S-1, S-3	35				0.4843%	0.3428%
S-3, S-1, S-4	36				0.4843%	0.7463%
S-3, S-2, S-1	37				0.3830%	0.5561%
S-3, S-2, S-2	38				0.3830%	0.4544%
S-3, S-2, S-3	39				0.3830%	0.3428%
S-3, S-2, S-4	40				0.3830%	0.7463%
S-3, S-3, S-1	41				0.2718%	0.8271%
S-3, S-3, S-2	42				0.2718%	0.2219%
S-3, S-3, S-3	43				0.2718%	0.3497%
S-3, S-3, S-4	44				0.2718%	0.7351%
S-3, S-4, S-1	45				0.6739%	1.0014%
S-3, S-4, S-2	46				0.6739%	0.3573%
S-3, S-4, S-3	47				0.6739%	0.5959%
S-3, S-4, S-4	48				0.6739%	1.0572%
S-4, S-1, S-1	49				0.5418%	0.8003%
S-4, S-1, S-2	50				0.5418%	0.1623%
S-4, S-1, S-3	51				0.5418%	0.3987%
S-4, S-1, S-4	52				0.5418%	0.8555%
S-4, S-2, S-1	53				0.4401%	0.6166%
S-4, S-2, S-2	54				0.4401%	0.5146%
S-4, S-2, S-3	55				0.4401%	0.4027%
S-4, S-2, S-4	56				0.4401%	0.8074%
S-4, S-3, S-1	57				0.3286%	0.6166%
S-4, S-3, S-2	58				0.3286%	0.5146%
S-4, S-3, S-3	59				0.3286%	0.4027%
S-4, S-3, S-4	60				0.3286%	0.8074%
S-4, S-4, S-1	61				0.7319%	1.0012%
S-4, S-4, S-2	62				0.7319%	0.5537%
S-4, S-4, S-3	63				0.7319%	0.6450%
S-4, S-4, S-4	64				0.7319%	1.1046%

EXHIBIT 30.3 Deposit Balance

State	Row Number	0	1	2	3	4
S-1, S-1, S-1	1	1000	1004.399359	1005.090811	1000.838591	993.6215439
S-1, S-1, S-2	2		1004.399359	1005.090811	1000.838591	994.7093495
S-1, S-1, S-3	3		1004.399359	1005.090811	1000.838591	995.5007146
S-1, S-1, S-4	4		1004.399359	1005.090811	1000.838591	989.7678473
S-1, S-2, S-1	5			1008.039251	1007.287888	1006.116177
S-1, S-2, S-2	6			1008.039251	1007.287888	1007.141678
S-1, S-2, S-3	7			1008.039251	1007.287888	1008.266076
S-1, S-2, S-4	8			1008.039251	1007.287888	1004.19867
S-1, S-3, S-1	9			1007.226895	1004.898457	999.4809481
S-1, S-3, S-2	10			1007.226895	1004.898457	1005.966125
S-1, S-3, S-3	11			1007.226895	1004.898457	1003.563401
S-1, S-3, S-4	12			1007.226895	1004.898457	998.9196377
S-1, S-4, S-1	13			1006.699225	1000.280387	994.431393
S-1, S-4, S-2	14			1006.699225	1000.280387	995.5121719
S-1, S-4, S-3	15			1006.699225	1000.280387	996.2984252
S-1, S-4, S-4	16			1006.699225	1000.280387	990.6025894
S-2, S-1, S-1	17				1008.580276	1005.319685
S-2, S-1, S-2	18				1008.580276	1011.73896
S-2, S-1, S-3	19				1008.580276	1009.360652
S-2, S-1, S-4	20				1008.580276	1004.764079
S-2, S-2, S-1	21				1011.882739	1012.534979
S-2, S-2, S-2	22				1011.882739	1015.835956
S-2, S-2, S-3	23				1011.882739	1015.408698
S-2, S-2, S-4	24				1011.882739	1014.541409
S-2, S-3, S-1	25				1011.455288	1011.266338
S-2, S-3, S-2	26				1011.455288	1017.25936
S-2, S-3, S-3	27				1011.455288	1015.993923
S-2, S-3, S-4	28				1011.455288	1012.176821
S-2, S-4, S-1	29				1010.587609	1012.08771
S-2, S-4, S-2	30				1010.587609	1013.100006
S-2, S-4, S-3	31				1010.587609	1014.209924
S-2, S-4, S-4	32				1010.587609	1010.194894
S-3, S-1, S-1	33				1008.419563	1008.900784
S-3, S-1, S-2	34				1008.419563	1009.918093
S-3, S-1, S-3	35				1008.419563	1011.033508
S-3, S-1, S-4	36				1008.419563	1006.998595
S-3, S-2, S-1	37				1009.433304	1009.919594
S-3, S-2, S-2	38				1009.433304	1010.936903
S-3, S-2, S-3	39				1009.433304	1012.052318
S-3, S-2, S-4	40				1009.433304	1008.017405
S-3, S-3, S-1	41				1010.544807	1008.32672
S-3, S-3, S-2	42				1010.544807	1014.378858
S-3, S-3, S-3	43				1010.544807	1013.100938

(Continued)

EXHIBIT 30.3 (Continued)

State	Row Number	0	1	2	3	4
S-3, S-3, S-4	44				1010.544807	1009.246184
S-3, S-4, S-1	45				1006.524045	1002.542593
S-3, S-4, S-2	46				1006.524045	1008.983766
S-3, S-4, S-3	47				1006.524045	1006.597345
S-3, S-4, S-4	48				1006.524045	1001.985091
S-4, S-1, S-1	49				1007.315059	1005.349076
S-4, S-1, S-2	50				1007.315059	1011.72826
S-4, S-1, S-3	51				1007.315059	1009.364806
S-4, S-1, S-4	52				1007.315059	1004.79694
S-4, S-2, S-1	53				1008.331656	1008.207056
S-4, S-2, S-2	54				1008.331656	1009.227375
S-4, S-2, S-3	55				1008.331656	1010.346091
S-4, S-2, S-4	56				1008.331656	1006.299237
S-4, S-3, S-1	57				1009.44629	1009.327263
S-4, S-3, S-2	58				1009.44629	1010.347583
S-4, S-3, S-3	59				1009.44629	1011.466299
S-4, S-3, S-4	60				1009.44629	1007.419444
S-4, S-4, S-1	61				1005.414202	1001.42911
S-4, S-4, S-2	62				1005.414202	1005.904074
S-4, S-4, S-3	63				1005.414202	1004.990898
S-4, S-4, S-4	64				1005.414202	1000.39531

EXHIBIT 30.4 Dollars of Effective Deposit Interest Paid

State	Row Number	0	1	2	3	4
S-1, S-1, S-1	1	0	1.607682952	5.357681289	10.28629259	13.13690932
S-1, S-1, S-2	2		1.607682952	5.357681289	10.28629259	12.06924105
S-1, S-1, S-3	3		1.607682952	5.357681289	10.28629259	11.2910385
S-1, S-1, S-4	4		1.607682952	5.357681289	10.28629259	16.90022365
S-1, S-2, S-1	5			2.401254888	3.856278353	7.252236034
S-1, S-2, S-2	6			2.401254888	3.856278353	6.226803129
S-1, S-2, S-3	7			2.401254888	3.856278353	5.100063266
S-1, S-2, S-4	8			2.401254888	3.856278353	9.163972455
S-1, S-3, S-1	9			3.217546831	6.248265922	11.43606209
S-1, S-3, S-2	10			3.217546831	6.248265922	4.986396862
S-1, S-3, S-3	11			3.217546831	6.248265922	7.385773262
S-1, S-3, S-4	12			3.217546831	6.248265922	11.99034362
S-1, S-4, S-1	13			3.747066285	10.83891692	11.78440529
S-1, S-4, S-2	14			3.747066285	10.83891692	10.72128438
S-1, S-4, S-3	15			3.747066285	10.83891692	9.946409118
S-1, S-4, S-4	16			3.747066285	10.83891692	15.53185521
S-2, S-1, S-1	17				5.546356067	9.352983719
S-2, S-1, S-2	18				5.546356067	2.91807482

EXHIBIT 30.4 (Continued)

State	Row Number	0	1	2	3	4
S-2, S-1, S-3	19				5.546356067	5.311785732
S-2, S-1, S-4	20				5.546356067	9.906068025
S-2, S-2, S-1	21				2.222811019	5.474953142
S-2, S-2, S-2	22				2.222811019	2.139550082
S-2, S-2, S-3	23				2.222811019	2.572492288
S-2, S-2, S-4	24				2.222811019	3.450195764
S-2, S-3, S-1	25				2.654219501	6.316598889
S-2, S-3, S-2	26				2.654219501	0.257574701
S-2, S-3, S-3	27				2.654219501	1.542930847
S-2, S-3, S-4	28				2.654219501	5.400716108
S-2, S-4, S-1	29				3.528808672	4.607870266
S-2, S-4, S-2	30				3.528808672	3.586922162
S-2, S-4, S-3	31				3.528808672	2.465161455
S-2, S-4, S-4	32				3.528808672	6.511365296
S-3, S-1, S-1	33				4.884246203	5.610372958
S-3, S-1, S-2	34				4.884246203	4.588631364
S-3, S-1, S-3	35				4.884246203	3.465977256
S-3, S-1, S-4	36				4.884246203	7.515297078
S-3, S-2, S-1	37				3.865852116	5.616038434
S-3, S-2, S-2	38				3.865852116	4.593260396
S-3, S-2, S-3	39				3.865852116	3.469469892
S-3, S-2, S-4	40				3.865852116	7.522900523
S-3, S-3, S-1	41				2.746884995	8.339680704
S-3, S-3, S-2	42				2.746884995	2.250575495
S-3, S-3, S-3	43				2.746884995	3.542401908
S-3, S-3, S-4	44				2.746884995	7.419319643
S-3, S-4, S-1	45				6.782949956	10.03953363
S-3, S-4, S-2	46				6.782949956	3.604997158
S-3, S-4, S-3	47				6.782949956	5.99863577
S-3, S-4, S-4	48				6.782949956	10.59255922
S-4, S-1, S-1	49				5.457292121	8.045364743
S-4, S-1, S-2	50				5.457292121	1.642413629
S-4, S-1, S-3	51				5.457292121	4.024164678
S-4, S-1, S-4	52				5.457292121	8.595731266
S-4, S-2, S-1	53				4.437733033	6.216865446
S-4, S-2, S-2	54				4.437733033	5.193422523
S-4, S-2, S-3	55				4.437733033	4.068888916
S-4, S-2, S-4	56				4.437733033	8.124937832
S-4, S-3, S-1	57				3.317475236	6.223772934
S-4, S-3, S-2	58				3.317475236	5.199187042
S-4, S-3, S-3	59				3.317475236	4.07340024
S-4, S-3, S-4	60				3.317475236	8.133982473
S-4, S-4, S-1	61				7.35814311	10.026471
S-4, S-4, S-2	62				7.35814311	5.569890521
S-4, S-4, S-3	63				7.35814311	6.482568131
S-4, S-4, S-4	64				7.35814311	11.0503297

The last step in the cash flow calculation is to calculate the change in deposit balances. An increase in deposit balances is a *negative* number because that represents cash in. A decrease in deposit balances is shown as a positive number, because we are trying to value the cost of the deposit franchise and a reduction in balances is a reduction in principal, like a principal payment on a bond. Exhibit 30.5 shows the net cash outflow on the deposit franchise at each node of the tree, with the additional

EXHIBIT 30.5 Change in Deposit Balance (Net Cash Outflow)

State	Row Number	0	1	2	3	4
S-1, S-1, S-1	1	0	-4.399358842	-0.691452301	4.252219768	993.6215439
S-1, S-1, S-2	2		-4.399358842	-0.691452301	4.252219768	994.7093495
S-1, S-1, S-3	3		-4.399358842	-0.691452301	4.252219768	995.5007146
S-1, S-1, S-4	4		-4.399358842	-0.691452301	4.252219768	989.7678473
S-1, S-2, S-1	5			-3.639892243	-2.197076488	1006.116177
S-1, S-2, S-2	6			-3.639892243	-2.197076488	1007.141678
S-1, S-2, S-3	7			-3.639892243	-2.197076488	1008.266076
S-1, S-2, S-4	8			-3.639892243	-2.197076488	1004.19867
S-1, S-3, S-1	9			-2.827535996	0.1923542	999.4809481
S-1, S-3, S-2	10			-2.827535996	0.1923542	1005.966125
S-1, S-3, S-3	11			-2.827535996	0.1923542	1003.563401
S-1, S-3, S-4	12			-2.827535996	0.1923542	998.9196377
S-1, S-4, S-1	13			-2.2998659	4.810424629	994.431393
S-1, S-4, S-2	14			-2.2998659	4.810424629	995.5121719
S-1, S-4, S-3	15			-2.2998659	4.810424629	996.2984252
S-1, S-4, S-4	16			-2.2998659	4.810424629	990.6025894
S-2, S-1, S-1	17				-0.541024597	1005.319685
S-2, S-1, S-2	18				-0.541024597	1011.73896
S-2, S-1, S-3	19				-0.541024597	1009.360652
S-2, S-1, S-4	20				-0.541024597	1004.764079
S-2, S-2, S-1	21				-3.843488146	1012.534979
S-2, S-2, S-2	22				-3.843488146	1015.835956
S-2, S-2, S-3	23				-3.843488146	1015.408698
S-2, S-2, S-4	24				-3.843488146	1014.541409
S-2, S-3, S-1	25				-3.416037252	1011.266338
S-2, S-3, S-2	26				-3.416037252	1017.25936
S-2, S-3, S-3	27				-3.416037252	1015.993923
S-2, S-3, S-4	28				-3.416037252	1012.176821
S-2, S-4, S-1	29				-2.548357804	1012.08771
S-2, S-4, S-2	30				-2.548357804	1013.100006
S-2, S-4, S-3	31				-2.548357804	1014.209924
S-2, S-4, S-4	32				-2.548357804	1010.194894
S-3, S-1, S-1	33				-1.192668141	1008.900784
S-3, S-1, S-2	34				-1.192668141	1009.918093
S-3, S-1, S-3	35				-1.192668141	1011.033508

EXHIBIT 30.5 (Continued)

State	Row Number	0	1	2	3	4
S-3, S-1, S-4	36				-1.192668141	1006.998595
S-3, S-2, S-1	37				-2.206409321	1009.919594
S-3, S-2, S-2	38				-2.206409321	1010.936903
S-3, S-2, S-3	39				-2.206409321	1012.052318
S-3, S-2, S-4	40				-2.206409321	1008.017405
S-3, S-3, S-1	41				-3.317912606	1008.32672
S-3, S-3, S-2	42				-3.317912606	1014.378858
S-3, S-3, S-3	43				-3.317912606	1013.100938
S-3, S-3, S-4	44				-3.317912606	1009.246184
S-3, S-4, S-1	45				0.702850045	1002.542593
S-3, S-4, S-2	46				0.702850045	1008.983766
S-3, S-4, S-3	47				0.702850045	1006.597345
S-3, S-4, S-4	48				0.702850045	1001.985091
S-4, S-1, S-1	49				-0.615834518	1005.349076
S-4, S-1, S-2	50				-0.615834518	1011.72826
S-4, S-1, S-3	51				-0.615834518	1009.364806
S-4, S-1, S-4	52				-0.615834518	1004.79694
S-4, S-2, S-1	53				-1.632431249	1008.207056
S-4, S-2, S-2	54				-1.632431249	1009.227375
S-4, S-2, S-3	55				-1.632431249	1010.346091
S-4, S-2, S-4	56				-1.632431249	1006.299237
S-4, S-3, S-1	57				-2.747065465	1009.327263
S-4, S-3, S-2	58				-2.747065465	1010.347583
S-4, S-3, S-3	59				-2.747065465	1011.466299
S-4, S-3, S-4	60				-2.747065465	1007.419444
S-4, S-4, S-1	61				1.285023047	1001.42911
S-4, S-4, S-2	62				1.285023047	1005.904074
S-4, S-4, S-3	63				1.285023047	1004.990898
S-4, S-4, S-4	64				1.285023047	1000.39531

assumption that we terminate our valuation as of time $T = 4$ by having all deposits “mature” or run off.

The total cash outflow on the deposit franchise is the sum of the effective interest paid in Exhibit 30.4 and the net cash outflow on the deposit franchise shown in Exhibit 30.5. We add these numbers at each node of the bushy tree to get the cash flow cost of the deposit franchise at each node of the bushy tree. We drop those numbers into the cash flow table in Exhibit 30.6 and derive the risk-neutral value (cost) of this nonmaturity deposit franchise with an initial balance of \$1,000. That value is \$946.7281.

In a rational world, a bank would be willing to pay up to \$1,000 – 946.7281 = 53.2719 to buy this deposit franchise, a premium of 5.32719 percent above par value of the deposits.

EXHIBIT 30.6 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1		−2.7917	4.6662	14.5385	1,006.7585
2		−2.7917	4.6662	14.5385	1,006.7786
3		−2.7917	4.6662	14.5385	1,006.7918
4		−2.7917	4.6662	14.5385	1,006.6681
5			−1.2386	1.6592	1,013.3684
6			−1.2386	1.6592	1,013.3685
7			−1.2386	1.6592	1,013.3661
8			−1.2386	1.6592	1,013.3626
9			0.3900	6.4406	1,010.9170
10			0.3900	6.4406	1,010.9525
11			0.3900	6.4406	1,010.9492
12			0.3900	6.4406	1,010.9100
13			1.4472	15.6493	1,006.2158
14			1.4472	15.6493	1,006.2335
15			1.4472	15.6493	1,006.2448
16			1.4472	15.6493	1,006.1344
17				5.0053	1,014.6727
18				5.0053	1,014.6570
19				5.0053	1,014.6724
20				5.0053	1,014.6701
21				−1.6207	1,018.0099
22				−1.6207	1,017.9755
23				−1.6207	1,017.9812
24				−1.6207	1,017.9916
25				−0.7618	1,017.5829
26				−0.7618	1,017.5169
27				−0.7618	1,017.5369
28				−0.7618	1,017.5775
29				0.9805	1,016.6956
30				0.9805	1,016.6869
31				0.9805	1,016.6751
32				0.9805	1,016.7063
33				3.6916	1,014.5112
34				3.6916	1,014.5067
35				3.6916	1,014.4995
36				3.6916	1,014.5139
37				1.6594	1,015.5356
38				1.6594	1,015.5302
39				1.6594	1,015.5218
40				1.6594	1,015.5403
41				−0.5710	1,016.6664
42				−0.5710	1,016.6294

EXHIBIT 30.6 (Continued)

Row Number	Current Time				
	0	1	2	3	4
43				−0.5710	1,016.6433
44				−0.5710	1,016.6655
45				7.4858	1,012.5821
46				7.4858	1,012.5888
47				7.4858	1,012.5960
48				7.4858	1,012.5777
49				4.8415	1,013.3944
50				4.8415	1,013.3707
51				4.8415	1,013.3890
52				4.8415	1,013.3927
53				2.8053	1,014.4239
54				2.8053	1,014.4208
55				2.8053	1,014.4150
56				2.8053	1,014.4242
57				0.5704	1,015.5510
58				0.5704	1,015.5468
59				0.5704	1,015.5397
60				0.5704	1,015.5534
61				8.6432	1,011.4556
62				8.6432	1,011.4740
63				8.6432	1,011.4735
64				8.6432	1,011.4456
Risk-Neutral Value =					946.7281

This example shows that there is nothing mysterious or different about the valuation (and by extension, hedging) of nonmaturity deposits. Using the standard HJM methodology like our example of Chapter 9, we populate the bushy tree with the appropriate cash flows and valuation follows directly. The same process gives us information for market risk (value at risk) and liquidity risk (cash flow at risk) from the nonmaturity deposit franchise. This is the power of the HJM framework in action. Results are both more accurate and simpler to derive than the purely mathematical approaches we cited earlier in the chapter, which rely on one-factor term structure models.

THE IMPACT OF BANK CREDIT RISK ON DEPOSIT RATES AND BALANCES

We learned one more time in the 2006–2011 credit crisis that deposit insurance does not mean deposits will not run off if the bank has credit problems. Northern Rock PLC was seized by the Bank of England early in 2008 even after the Bank of England

declared its full support for Northern Rock late in 2007. Similarly, as we discuss in detail in Chapter 37, the U.S. Federal Reserve lent more than \$1 trillion to both U.S. and international financial institutions in the 2008–2009 period, even though almost all of them were beneficiaries of U.S. government deposit insurance. Exhibit 30.7, provided by Kamakura Corporation, shows that Wachovia Bank was forced to borrow \$29 billion on October 6, 2008, from the Federal Reserve and that it increased borrowings to \$36 billion on October 8 because its liquid asset holdings were insufficient to fully pay off deposits withdrawn by depositors of all types.

How do we take the impact of the bank's own risk into account in the analysis above? We make that modification in this section. In a sense, nonmaturity deposits have an embedded put option or credit default swap that is triggered when a bank's risk hits a high enough level. In our example above, we are assuming that all deposits run off at time $T = 4$. If we incorporate credit risk, it will affect deposit rates and balances at times 1, 2, or 3. One sample modification is to say the following: the deposit balance will follow the equation above except when the bank's one-period funding cost (the spot rate, not the deposit rate) is 4.00 percent or higher. In this circumstance, it is assumed that 70 percent of the bank's deposits run off immediately. This happens at 8 of the 64 nodes at time $T = 3$, as seen in Exhibit 30.8.

The risk-neutral cost of the nonmaturity deposit franchise rises slightly to \$948.2673, because the effective maturity of the nonmaturity deposits is shortened. While the impact on value is relatively small in this example, the impact on cash flow at risk at time $T = 3$ is very, very large. In Chapters 36 through 41, we will analyze




EXHIBIT 30.7 Wachovia Bank NA, Primary, Secondary, and Other Extensions of Credit by the Federal Reserve, 2008–2009

Sources: Kamakura Corporation; Federal Reserve.

EXHIBIT 30.8 Maturity of Cash Flow Received

Row Number	Current Time			
	0	1	2	3
1		-2.7917	4.6662	706.6626
2		-2.7917	4.6662	706.6626
3		-2.7917	4.6662	706.6626
4		-2.7917	4.6662	706.6626
5			-1.2386	1.6592
6			-1.2386	1.6592
7			-1.2386	1.6592
8			-1.2386	1.6592
9			0.3900	6.4406
10			0.3900	6.4406
11			0.3900	6.4406
12			0.3900	6.4406
13			1.4472	706.8309
14			1.4472	706.8309
15			1.4472	706.8309
16			1.4472	706.8309
17				5.0053
18				5.0053
19				5.0053
20				5.0053
21				-1.6207
22				-1.6207
23				-1.6207
24				-1.6207
25				-0.7618
26				-0.7618
27				-0.7618
28				-0.7618
29				0.9805
30				0.9805
31				0.9805
32				0.9805
33				3.6916
34				3.6916
35				3.6916
36				3.6916
37				1.6594
38				1.6594
39				1.6594
40				1.6594
41				-0.5710
42				-0.5710
43				-0.5710

(Continued)

EXHIBIT 30.8 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
44				-0.5710	1,016.6655
45				7.4858	1,012.5821
46				7.4858	1,012.5888
47				7.4858	1,012.5960
48				7.4858	1,012.5777
49				4.8415	1,013.3944
50				4.8415	1,013.3707
51				4.8415	1,013.3890
52				4.8415	1,013.3927
53				2.8053	1,014.4239
54				2.8053	1,014.4208
55				2.8053	1,014.4150
56				2.8053	1,014.4242
57				0.5704	1,015.5510
58				0.5704	1,015.5468
59				0.5704	1,015.5397
60				0.5704	1,015.5534
61				8.6432	1,011.4556
62				8.6432	1,011.4740
63				8.6432	1,011.4735
64				8.6432	1,011.4456
Risk-Neutral Value =					948.2673

the impact of deposit outflows on the bankruptcy risk of the organization in much more detail.

CASE STUDY: GERMAN THREE-MONTH NOTICE SAVINGS DEPOSITS³

In this section, we use the Kamakura Risk Manager enterprise-wide risk management system to value and analyze the interest rate sensitivity of three-month notice savings deposits in Germany. Instead of the three-factor HJM model we have illustrated earlier, we use a one-factor model consistent with the papers by Janosi and colleagues (1999) and Jarrow and van Deventer (1996, 1998). This data set is of interest because it is perhaps the longest deposit history available. There are 469 monthly observations on the deposit balance data running from December 1959 to December 1998. There are 361 monthly observations on the deposit rate running from June 1967 to February 1999. Note that these figures are national deposit figures, rather than the numbers for a particular bank. As such, they dramatically understate the risk of deposit outflows. With this caveat, we proceed.

The valuation yield curve derived in Kamakura Risk Manager was based on April 28, 1999, interest rates in Germany, including 1-month, 3-month, 6-month,

9-month, and 12-month money market rates and 2-year, 3-year, 5-year, 7-year, 10-year, 15-year, 20-year, and 30-year Deutsche Mark interest rate swap rates. The maximum smoothness forward rate smoothing technique of Chapter 5 was selected for yield curve smoothing. The speed of mean reversion α was set at 0.01 and the interest rate volatility σ in the extended Vasicek model was set at 0.008.

Kamakura Risk Manager was used to fit the deposit interest rate regression of the Jarrow and van Deventer (1998) model to the deposit rate history. The regression explains deposit rate variation as a function of the lagged deposit rate, a constant, and short-term interest rates in Germany. The resulting relationship explains 98.76 percent in monthly variation of the three-month notice savings deposit rate, and the relationship realistically says, all other things being equal, that a 100 basis point move in short-term open market rates in Germany leads to a two basis point move in the same month in the deposit rate. Of course, the long-term equilibrium level of the deposit rate is much higher than this because of the impact of the lagged deposit rate.

The deposit balance relationship was also fitted in Kamakura Risk Manager using linear regression. The relationship explains 99.97 percent of the variation in deposit balances and shows quite clearly that (all other things being equal) a rise in money market rates in Germany lowers savings deposit balances—disintermediation takes place as has been found in many countries.

This excellent fit comes in spite of the impact of the unification with East Germany during the life of the data set, which caused a jump in deposit balances.

Using representative assumptions on long-term deposit balance levels, three-month notice deposits in Germany were found to have a present value cost of about 67 percent of their nominal amount and a negative duration of 2.29 years. This means that, all other things being equal, when interest rates rise, the present value cost of the deposits actually declines because of the reduction in balances that comes about as noted above. This is identical to the experience of the U.S. savings and loan industry during its high interest rate crisis in the late 1970s and 1980s.

Negative duration was also found for parallel shifts in the German yield curve, much the same as found by Chase Manhattan Bank in its own analysis of the effective interest rate sensitivity of nonmaturity deposits in the United States in the 1980s and 1990s.⁴

THE REGULATORS' VIEW

The regulatory agencies around the world clearly have a strong interest in valuation of nonmaturity deposits. James O'Brien (2000) of the Board of Governors of the Federal Reserve takes a Monte Carlo simulation approach to the valuation of nonmaturity deposits with a strong similarity to the Jarrow and van Deventer model. The National Credit Union Administration of the United States in December 2001 commissioned a study of nonmaturity deposit evaluation techniques by National Economic Research Associates, with David Ellis and James Jordan (2001) as lead authors. We urge readers with a strong interest in this topic to review both papers in detail.⁵

CONCLUSION

The default probability of the financial institution offering nonmaturity deposits to consumers is a critical determinant of nonmaturity deposit balances and rates. This in turn has a big impact on the valuation and risk of both the firm as a whole and nonmaturity deposits in particular. We revisit these issues in great detail in Chapters 36 through 41.

NOTES

1. For more on both deposit and charge card valuation, see Jarrow and van Deventer (1996, 1998) and Janosi, Jarrow, and Zullo (1999).
2. The average premium paid by purchasers of failed bank deposit franchises was 2.32 percent in 1,225 auctions run through October 28, 1994, by the Resolution Trust Corporation in the United States. The highest premium paid was 25.33 percent. Most of the deposits auctioned (more than 70 percent on average) were time deposits, which would be expected to have very low premiums.
3. The authors would like to thank Robert Fiedler, formerly of Deutsche Bank, and Karin Gradischnig of Reuters for helpful comments on this section.
4. Private conversation with Dr. Robert Selvaggio, formerly of Chase Manhattan Bank.
5. The O'Brien (2000) paper is available at this link: www.federalreserve.gov/pubs/feds/2000/200053/200053pap.pdf. The Ellis and Jordan (2001) paper is available at the website of the NCUA: <http://www.ncua.gov/resources/documents/alm/almnmdeposits.pdf>.

Foreign Exchange Markets

In this chapter, we continue to progress toward our integrated measure of interest rate risk, market risk, liquidity risk, and credit risk using the Jarrow-Merton put option concept from Chapter 1. A high percentage of smaller financial institutions have the luxury of financial assets and liabilities that are almost all denominated in a single currency. In fact, due in part to the advent of the euro, the percentage of financial institutions in this situation is rising. Nonetheless, for most of the largest financial institutions, foreign currency risk is a critical element of total risk. As van Deventer and Imai (2003) showed with respect to the Korean Development Bank and many other issuers, credit risk and foreign exchange rates of the home country can move very strongly together.

The Heath, Jarrow, and Morton (HJM) term structure model approach of Chapters 6 through 9 again provides a sound foundation with which to approach foreign currency denominated securities and derivatives. This chapter is heavily based on the HJM approach to foreign exchange (FX) options valuation of Amin and Jarrow (1991) and their paper, “Pricing Foreign Currency Options under Stochastic Interest Rates.” The variety of FX-related securities is great, but we will limit ourselves to an introductory view of the HJM approach’s implications for FX analysis. In this chapter, we concentrate solely on foreign exchange forward contracts and European options on the spot foreign exchange rate. Foreign exchange futures contracts and options on FX futures involve a combination of the techniques in this chapter and those in Chapters 22 and 23, but we leave such analysis as advanced exercises for interested readers.¹

The realism of the HJM term structure model framework is essential for modeling of foreign currency derivatives in order to minimize the model error inherent in the legacy versions of FX options formulas, which assumed that interest rates in both countries are constant.

SETTING THE STAGE: ASSUMPTIONS FOR THE DOMESTIC AND FOREIGN ECONOMIES

Initially, we adopt the assumptions of Amin and Jarrow (1991). Later in the chapter, we discuss how the three-factor HJM model from Chapter 9 can be used in a way consistent with the Amin and Jarrow approach.

Amin and Jarrow base their model, without loss of generality, on four independent standard Brownian motions, which are commonly used in the HJM framework and in the legacy term structure models reviewed in Chapter 13. Amin and Jarrow assume that interest rates in both countries 1 and 2 are driven by a two-factor interest rate model like the model of Chapter 8. They assume that interest rates are driven in country 1 by risk factors 1 and 2. They assume that interest rates in country 2 are driven by risk factors 2 and 3. The presence of risk factor 2 as a driver of interest rates in both countries means that we can model correlated interest rate movements in this framework, whether that correlation is positive or negative. In both countries 1 and 2, HJM no-arbitrage conditions are imposed and the resulting random evolution of interest rates in each country is similar to the modeling in Chapter 8. The authors note that “the analysis is easily generalized to a finite number of independent Brownian motions with any subset common across the two term structures” (Amin and Jarrow, 1991, 313). This is consistent with our findings in Chapter 3 that 5 to 10 risk factors drive interest rate movements.

The link between the two countries is the spot exchange rate. Amin and Jarrow assume that all three of the risk factors introduced so far are allowed to impact the spot exchange rate, allowing for correlation with the interest rate risk factors in both countries. To these three factors, they add a fourth independent risk factor, which drives the spot exchange rate but not interest rates in either country. Amin and Jarrow note that, since there are four sources of risks in this economy, in general it will require four unique assets to hedge total risk. With the addition of three assumptions (concerning complete markets, existence of a unique asset-specific martingale measure, and independence of the market prices of risk for the four risk factors), the authors derive no arbitrage conditions, showing that the result is consistent with the HJM conditions imposed on one country as in Chapters 6 through 9.

This framework allows the valuation of any European or American-style contingent claim in this two-country economy. The authors illustrate valuation techniques in the special case where the volatilities of the risk factors are nonstochastic, deterministic functions of time (like our interest rate assumptions of Chapter 6 for the one-factor HJM interest rate model). They go on to derive closed form solutions for options on the spot exchange rate, the pricing of forwards and futures, and the pricing of options on forwards and futures.

FOREIGN EXCHANGE FORWARDS

The pricing of foreign exchange forwards can be set by an arbitrage argument that doesn't rely on any of the assumptions above. Let's assume it's time t and the observable forward price of country 2 currency at time T is a function of the spot rate S and interest rates in the two countries is $H(t, T, S)$. Let's assume we want to invest, from the perspective of our position in country 1, in country 2 currency for T years and maximize our currency 2 holdings at time T . We define $P_1 = P_1(t, T)$ as the T maturity country 1 zero-coupon bond and $P_2 = P_2(t, T)$ as the country 2 T maturity zero-coupon bond. Assume that country 1 is the U.S. and country 2 is Japan. S and H both represent the dollar cost of one yen. There are two investment strategies:

Strategy 1

1. Borrow P_1 dollars, which will become \$1 at time T .
2. Sell \$1 forward at current time t for yen deliverable at time T , which will generate yen proceeds of $1/H$

Strategy 2

1. Borrow P_1 dollars.
2. At current time t , convert the dollars to P_1/S yen at the spot exchange rate S .
3. Invest P_1/S yen in a zero-coupon yen bond maturing at time T , resulting in time T yen proceeds of $P_1/(SP_2)$.

To avoid arbitrage, the results of these two strategies must be equal, so the forward rate H must equal:

$$H(t, T, S) = \frac{SP_2}{P_1}$$

This well-known expression makes it clear that the forward rate depends on all of the risk factors that determine domestic interest rates, foreign interest rates, and the spot foreign exchange rate. With this as background, we turn to numerical valuation of foreign currency derivatives in the HJM framework.

NUMERICAL METHODS FOR VALUATION OF FOREIGN CURRENCY DERIVATIVES

Jarrow and Turnbull (1996, Chapter 11) discuss the construction of a lattice to model movements of the spot foreign exchange rate in the special case where (1) the spot foreign exchange rate is lognormally distributed and (2) interest rates are constant. In the Amin and Jarrow example, four risk factors drive interest rates in both countries (two factors each, with one factor affecting both countries) and the spot exchange rate. In Chapter 9, we showed how to construct a bushy tree for interest rates in one country when there are three random factors. Using these insights, how should we value foreign currency derivatives? The elegant mathematics in Amin and Jarrow, even using a restricted number of interest rate factors, is complex. At the same time, the assumptions are too severe to be realistic. This leaves numerical methods as the only realistic alternative for valuation. As in Chapters 6 through 9, the two principal tools to use are Monte Carlo simulation and bushy trees.

To do Monte Carlo simulation with the Amin and Jarrow assumptions, we would simulate the movements of the four risk factors at M points in time and over N scenarios. At each point in time for each scenario, the Amin and Jarrow/HJM no-arbitrage restrictions give us the complete term structure of interest rates (as we derived in Chapter 8) from the two factors affecting each country plus a no-arbitrage simulated value of the spot rate. From the expression for H above, we would also have a complete term structure of forward foreign exchange rates. For most of the securities we have studied in Chapters 19 through 30, we can determine the cash flows in the base currency at each point in time. We discount at the correct HJM

discount factor (using the value of the money market fund equal to the compounded investment in the spot rate from time 0 onward) to get valuation.

What about the three-factor HJM example of Chapter 9? How would we use that in a foreign exchange context? Let's assume that there are three factors driving interest rates in country 1, three other (nonoverlapping) factors driving interest rates in country 2, and one unique factor driving the spot rate. In this special case, it is easy to describe the augmented bushy tree that would result. As in Chapter 9, at time step 1, we would have four nodes on the tree for country 1 interest rates with a complete term structure of interest rates at each node. We would have four independent nodes at time step 1 for country 2 as well. For the value of the spot rate, the random variation of the spot rate is driven by only one factor and the drift term is constrained by the no-arbitrage conditions in Amin and Jarrow. We can model the spot rate movements with an up shift and a down shift, properly constrained by no-arbitrage restrictions. Therefore, at time step 1 we will have two values for the spot rate, four country 1 term structures, and four country 2 term structures. From the expression for forward rates H given previously, this means that there are $2 \times 4 \times 4 = 32$ term structures of forward foreign exchange rates at time step 1.

What about time step 2? Let's assume that the volatility assumptions for the spot rate are such that a bushy tree results (i.e., the lattice does not recombine). At time step 2, then, we have four values for the spot rate, 16 nodes or term structures for country 1, and 16 term structures for country 2. This means that at time step 2 there will be $4 \times 16 \times 16 = 1,024$ term structures for forward foreign exchange rates. At time step 3, there will be $8 \times 64 \times 64 = 32,768$ term structures for forward rates. While this means that a lot of numbers will get thrown around, this is what computers are for: they do what they are told. In the world of 64-bit operating systems, it is a routine task.

What happens when one or more of the three factors in country 1 and 3 factors in country 2 are shared? And what if some of these factors affect the spot rate? This will affect the construction of the bushy tree (just as it impacts the results of a Monte Carlo simulation), but the result is still completely consistent with the no-arbitrage restrictions of Amin and Jarrow/HJM.

Before closing the chapter, we finish with a brief review of the results from a one-factor term structure model.

LEGACY APPROACHES TO FOREIGN EXCHANGE OPTIONS VALUATION

Hilliard, Madura, and Tucker (HMT) (1991) assume that bond prices in country 1 and country 2 follow the one-factor Vasicek stochastic process and that the spot rate follows a lognormal process with the form given in the previous section.

Model Risk Alert

We include the HMT results here for expository purposes, but the one-factor Vasicek model is not accurate enough to meet best practice standards.

Prior to Amin and Jarrow (1991), the most popular model of foreign exchange options had been the Garman and Kohlhagen (1983) model, a variation

on the Black-Scholes model that assumes that the interest rates in country 1 and country 2 are constant at levels r_1 and r_2 . Garman and Kohlhagen show that the value of a European call option maturing at time T from the perspective of time t to purchase one unit of country 2's currency at a strike price of K is

$$V(t, T, S, r_1, r_2) = Se^{-r_2\tau}N(h) - Ke^{-r_1\tau}N(h - \sigma_s\sqrt{\tau})$$

where

$$\tau = T - t$$

and

$$h = \frac{\ln\left(\frac{S}{K}\right) + \tau\left[r_1 - r_2 + \frac{\sigma_s^2}{2}\right]}{\sigma_s\sqrt{\tau}}$$

When interest rates change, the volatility used in the option calculation must change. Interest rates in both countries and their instantaneous correlation with the spot rate have an impact on the variance in the price of forward foreign exchange for delivery at time T . HMT show that the value of the call option under the stochastic interest rate case is

$$\begin{aligned} V(t, T, S, r_1, r_2, K) &= P_1(r_1, t, T)[H(t, T, S, r_1, r_2)N(h^*) - KN(h^* - \nu_b)] \\ &= SP_2(r_2, t, T)N(h^*) - P_1(r_1, t, T)KN(h^* - \nu_b) \end{aligned}$$

where

$$h^* = \frac{\ln\left(\frac{H(t, T, S, r_1, r_2)}{K}\right) + \frac{1}{2}\nu_b^2}{\nu_b}$$

and ν_b is the variance of the price of the forward foreign exchange rate H at time T from the perspective of time t .

Using our standard notation for each country for the volatility of interest rates in each country as of time T from the perspective of time t :

$$\nu_i^2(t, T) = \frac{\sigma_i^2}{2\alpha_i}[1 - e^{-2\alpha_i\tau}]$$

and

$$F_i = \frac{1}{\alpha_i}[1 - e^{-\alpha_i\tau}]$$

allows us to write ν_b as in HMT as

$$\nu_b^2 = \sigma_s^2\tau + M_2 + M_5 + 2(\rho_{s1}\sigma_s M_1 - \rho_{s2}\sigma_s M_4 - \rho_{12}M_3)$$

where HMT use the notation

$$M_1 = \frac{\sigma_1}{\alpha_1} [\tau - F_1]$$

$$M_2 = \frac{\sigma_1^2}{\alpha_1^2} [\tau - 2F_1] + \frac{\nu_1^2}{\alpha_1^2}$$

$$M_3 = \frac{\sigma_1 \sigma_2}{\alpha_1 \alpha_2} \left[\tau - F_1 - F_2 + \frac{1}{\alpha_1 + \alpha_2} \left(1 - e^{-(\alpha_1 + \alpha_2)\tau} \right) \right]$$

$$M_4 = \frac{\sigma_2}{\alpha_2} [\tau - F_2]$$

$$M_5 = \frac{\sigma_2^2}{\alpha_2^2} [\tau - 2F_2] + \frac{\nu_2^2}{\alpha_2^2}$$

IMPLICATIONS OF A TERM STRUCTURE MODEL-BASED FX OPTIONS FORMULA

In an environment where the correlation between risk exposures, as popularized by the value-at-risk (VaR) approach we discuss in later chapters, is a very high priority, the HJM term structure model approach to FX option valuation provides an immense improvement over the constant interest rate-based Garman-Kohlhagen approach:

- Correlation between country 1 interest rates, country 2 interest rates, and the spot rate are explicitly recognized.
- The option valuation can be stress tested for changes in any risk factor and any parameter in the Amin and Jarrow/HJM framework.
- The Amin and Jarrow formulation explicitly recognizes that the volatility of the spot rate at time T , the option exercise date, is directly affected by the volatility of all risk factors driving interest rates in the two countries. The more traditional approach ignores this source of volatility.
- Interest rate-related hedges stemming from FX options and forward positions can be done using the general HJM approach. The number of hedging instruments has to equal the number of risk factors for a “perfect hedge” (which comes about only when the model is true). We discuss this in more detail in later chapters, in a way fully consistent with all other fixed income and interest rate derivative hedging. Hedges using maturities other than the maturity of the FX option can be calculated and put into place. Like its relative the Black-Scholes model, the Garman-Kohlhagen model’s fixed interest rate assumption makes this impossible without some additional (and contradictory) assumptions.

All of these benefits are powerful ones, and they come with an ease of use far superior to complex but inaccurate closed form solutions.

THE IMPACT OF CREDIT RISK ON FOREIGN EXCHANGE RISK FORMULAS

The HMT FX option formula in this chapter, like the Garman-Kohlhagen model, implicitly assumes that our options counterparty is riskless. As JPMorgan Chase discovered in 1998 in a \$500 million FX dispute with SK Securities of Korea, it is much more likely that the default probability of the counterparty rises sharply as the amount of money owed on the FX position increases.² For this reason it is very important for FX rates to be explicitly modeled as drivers of default probabilities so that the default probabilities of all international counterparties will vary properly as exchange rates move.

A failure to incorporate this critical link with credit risk will result in a substantial understatement of total risk and of the correlation between the default probabilities in broad classes of borrowers.

We return to this point in Chapters 36 through 41 frequently as we value the Jarrow-Merton put option as a comprehensive measure of total risk—interest rate risk, market risk, liquidity risk, credit risk, and foreign exchange risk.

NOTES

1. For an excellent reference on this topic, see Chapter 11 in Jarrow and Turnbull (1996).
2. See van Deventer and Imai (2003) for a discussion of this incident.

Impact of Collateral on Valuation Models

The Example of Home Prices in the Credit Crisis

The 2006–2011 credit crisis has reemphasized what careful bankers have known for decades—collateral can have a significant impact on mitigating the credit risk of a particular loan structure. Lack of such collateral can have the opposite effect, which became obvious when slight declines in home prices triggered almost instantaneous defaults on mortgage loans made at 100 percent loan-to-value ratio levels. That is, those loans for which lenders were willing to lend 100 percent of the appraised value of the house on the assumption that increases in the home price would soon restore a prudent level of collateral.

In this chapter, we add collateralized transactions to our list of assets and liabilities that can be valued in a total risk management framework. Again, our objective in this effort is to be able to value the Jarrow-Merton put option on the value of the assets (and liabilities) of the firm as a measure of total integrated risk, including interest rate risk, market risk, liquidity risk, and credit risk.

THE IMPACT OF CHANGING HOME PRICES ON COLLATERAL VALUES IN THE CREDIT CRISIS

Up until 2006 in the United States and until the peak of the Japanese bubble in December 1989, bankers were willing to lend as much as 100 percent of the value of the home in the form of a first mortgage loan. This compares to a much more prudent, long-term “normal” level of an 80 percent loan-to-value ratio. What happened in the United States was a crushing collapse in home prices and an immediate response by mortgage loan borrowers. Exhibit 32.1 shows the change in the Case-Shiller home price index for Los Angeles (upper line) and for a 10-city composite (lower line) from 1987 to 2011.

As home prices fell, lenders began scrambling to check collateral values in anticipation of defaults and the default and prepayment behavior of home owners immediately changed to reflect the “new reality” of falling home prices. In Chapters 27 and 28, we analyzed prepayment and default modeling where behavior by consumers was potentially “irrational.” In this chapter, we return to the amortizing loan structure with rational prepayment from Chapter 28 and revisit our conclusions in the light of changing collateral values, namely the price of the home. While we discuss mortgages and home prices exclusively in this chapter, the analysis is nearly identical




EXHIBIT 32.1 Case-Schiller Home Price Index for Los Angeles and 10-City Composite, January 1987 to November 2011, Not Seasonally Adjusted

Sources: Kamakura Corporation; Standard & Poor's.

for any form of lending where collateral is part of the overall lending arrangement. The key question we need to address is this: How does the value of the collateral vary with the changes in prepayment and default experience on the loan? We address that issue in the next section.

MODELING VARIATIONS IN COLLATERAL VALUES

In Chapter 31, we used Amin and Jarrow (1991) to study FX derivatives modeling in a Heath, Jarrow, and Morton (HJM) context. In this chapter, we turn to another critical paper by Amin and Jarrow (1992), “Pricing Options on Risky Assets in a Stochastic Interest Rate Economy.” This paper presents a general framework for modeling the value of any asset under no-arbitrage assumptions in an HJM interest rate environment such as that illustrated in Chapters 6 through 9. The authors show that the expected return on the asset will be a premium to the spot risk-free rate of interest that depends on a number of factors:

- The risk factors that affect the value of the risk asset
- The volatility of the risky asset's returns as a function of these risk factors
- The market price of risk for each of these risk factors
- The level of the spot rate of the risk-free term structure

What does this mean for our mortgage loan example? It means that we need to tie the price of the home to each of these factors, most notably to the level of the spot rate of the risk-free rate of interest. Amin and Jarrow (1992, 220, equation (3.7)), specify the expected return on the risky asset, in this case the underlying home that is collateral on a mortgage loan, in this way:

$$\mu(t, x) - r(t) = - \sum_{i=1}^d \delta_i(t, x) n_i(t)$$

The μ term refers to the expected return on the risky asset, our home used as collateral. The $r(t)$ term refers to the risk-free rate. The δ terms are the volatilities with respect to the d risk factors impacting the value of the home via home price returns. The first n of these d terms will be related to the random interest rate economy. The $n(t)$ terms refer to the market prices of risk for each factor. We captured these terms in our HJM “bushy tree” in Chapters 6 through 9. We now use these insights of Amin and Jarrow on the rationally prepaid mortgage loan of Chapter 28.

THE IMPACT OF COLLATERAL VALUES ON A RATIONALLY PREPAID MORTGAGE

In Chapter 28, we analyzed rational prepayment behavior by the borrower on an amortizing mortgage with a principal of \$100, annual amortization, a four-period maturity, and a coupon rate of 2.00 percent. The value of this mortgage and the resulting cash flows are reproduced in Exhibit 32.2. We derived a mark-to-market value using the three-factor Heath, Jarrow, and Morton framework of Chapter 9. That value was 101.4234.

In that chapter, we ignored the value of the home and talked about the various ways in which borrower default and prepayment behavior can be jointly modeled. Best practice in that regard is multinomial logit. In this chapter, we deal with a special case that assumes a special type of rationality on the part of the borrower: (1) the borrower will always prepay if (starting at time $T = 1$) the risk-neutral value is greater than the principal value on the loan and (2) the borrower will default (a so-called “strategic default”) if the value of the home is too low relative to the terms on the loan. We discuss the possible meanings of “too low” in the context of our example.

Step 1 in this process is to use the Amin and Jarrow (1992) insights to model the price of the home. For ease of exposition, we assume that the market is completely indifferent to variations in home prices. This means that the market price of risk is zero and that the expected return on the home should be equal to the risk-free rate on each of the nodes of the bushy tree used in Chapters 9 and 28. Recall that the spot rates on each of the nodes on the bushy tree take on the values in Exhibit 32.3.

What volatility should we use for home price movements? For ease of exposition, we make a false assumption that was very commonly made in the credit crisis and hereby invoke a Model Risk Alert: we assume that the returns on home prices are independent from period to period (they are not). If we make this assumption and run a regression to predict the annual returns of home prices in Los Angeles as a function

EXHIBIT 32.2 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	26.2624	26.2624	26.2624	26.2624
2		102.0000	26.2624	26.2624	26.2624
3		102.0000	26.2624	26.2624	26.2624
4		26.2624	26.2624	26.2624	26.2624
5			0.0000	26.2624	26.2624
6			0.0000	26.2624	26.2624
7			0.0000	26.2624	26.2624
8			0.0000	26.2624	26.2624
9			0.0000	26.2624	26.2624
10			0.0000	52.0098	0.0000
11			0.0000	26.2624	26.2624
12			0.0000	26.2624	26.2624
13			26.2624	26.2624	26.2624
14			26.2624	26.2624	26.2624
15			77.2524	26.2624	26.2624
16			26.2624	26.2624	26.2624
17				0.0000	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	0.0000
41				0.0000	0.0000

(Continued)

EXHIBIT 32.2 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				0.0000	0.0000
46				0.0000	0.0000
47				0.0000	0.0000
48				0.0000	0.0000
49				26.2624	26.2624
50				52.0098	0.0000
51				52.0098	0.0000
52				26.2624	26.2624
53				26.2624	26.2624
54				26.2624	26.2624
55				52.0098	0.0000
56				26.2624	26.2624
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				26.2624	26.2624
62				26.2624	26.2624
63				26.2624	26.2624
64				26.2624	26.2624
Risk-Neutral Value =					101.4234

of the one-year risk free rate, we find an r -squared of zero and a standard error of the regression equal to 0.13271528. We use this value for the sigma of home price returns. For simplicity, we want to model home prices with an upshift and downshift like we did for interest rates in Chapter 6. We need the mean return on the price of the house to equal the one-period risk-free rate of interest r and we need the variance of home price returns such that $\sigma = 0.13271528$. The upshifts and downshifts that satisfy these constraints, given that the length of the period is $\Delta = 1$, are the following:

$$\begin{aligned}\text{Upshift} &= r + \sigma \\ \text{Downshift} &= r - \sigma\end{aligned}$$

How do we calculate these shifts in the context of our bushy interest rate tree in Chapter 9? Recall that we have four nodes on the bushy tree at time $T = 1$, 16 at $T = 2$, and 64 at $T = 3$. Our first shift is from Time 0 to Time $T = 1$. There is only one spot rate at time 0, so we have two values of home prices that can prevail at time $T = 1$, the upshift value and the downshift value. The upshift value is the spot rate

EXHIBIT 32.3 Spot Rate Process

State	Row Number	0	1	2	3
S-1, S-1, S-1	1	0.3003%	2.1653%	4.6388%	6.1106%
S-1, S-1, S-2	2		0.6911%	1.4142%	5.5667%
S-1, S-1, S-3	3		1.0972%	2.6089%	5.1710%
S-1, S-1, S-4	4		1.3611%	4.9179%	8.0375%
S-1, S-2, S-1	5			2.2496%	3.1041%
S-1, S-2, S-2	6			0.5984%	2.5913%
S-1, S-2, S-3	7			0.8121%	2.0291%
S-1, S-2, S-4	8			1.2459%	4.0628%
S-1, S-3, S-1	9			1.9217%	5.2210%
S-1, S-3, S-2	10			1.4149%	1.9784%
S-1, S-3, S-3	11			0.8591%	3.1798%
S-1, S-3, S-4	12			2.8695%	5.5017%
S-1, S-4, S-1	13			2.2088%	5.4252%
S-1, S-4, S-2	14			1.7005%	4.8848%
S-1, S-4, S-3	15			1.1432%	4.4917%
S-1, S-4, S-4	16			3.1593%	7.3396%
S-2, S-1, S-1	17				4.1517%
S-2, S-1, S-2	18				0.9421%
S-2, S-1, S-3	19				2.1313%
S-2, S-1, S-4	20				4.4295%
S-2, S-2, S-1	21				2.2036%
S-2, S-2, S-2	22				0.5531%
S-2, S-2, S-3	23				0.7667%
S-2, S-2, S-4	24				1.2004%
S-2, S-3, S-1	25				2.6231%
S-2, S-3, S-2	26				−0.3734%
S-2, S-3, S-3	27				0.2593%
S-2, S-3, S-4	28				2.1679%
S-2, S-4, S-1	29				1.7764%
S-2, S-4, S-2	30				1.2703%
S-2, S-4, S-3	31				0.7153%
S-2, S-4, S-4	32				2.7228%
S-3, S-1, S-1	33				2.2804%
S-3, S-1, S-2	34				1.7718%
S-3, S-1, S-3	35				1.2141%
S-3, S-1, S-4	36				3.2315%
S-3, S-2, S-1	37				2.2804%
S-3, S-2, S-2	38				1.7718%
S-3, S-2, S-3	39				1.2141%
S-3, S-2, S-4	40				3.2315%
S-3, S-3, S-1	41				3.6354%
S-3, S-3, S-2	42				0.6093%
S-3, S-3, S-3	43				1.2483%

(Continued)

EXHIBIT 32.3 (Continued)

State	Row Number	0	1	2	3
S-3, S-3, S-4	44				3.1757%
S-3, S-4, S-1	45				4.5070%
S-3, S-4, S-2	46				1.2864%
S-3, S-4, S-3	47				2.4797%
S-3, S-4, S-4	48				4.7858%
S-4, S-1, S-1	49				3.5013%
S-4, S-1, S-2	50				0.3117%
S-4, S-1, S-3	51				1.4934%
S-4, S-1, S-4	52				3.7773%
S-4, S-2, S-1	53				2.5831%
S-4, S-2, S-2	54				2.0730%
S-4, S-2, S-3	55				1.5136%
S-4, S-2, S-4	56				3.5370%
S-4, S-3, S-1	57				2.5831%
S-4, S-3, S-2	58				2.0730%
S-4, S-3, S-3	59				1.5136%
S-4, S-3, S-4	60				3.5370%
S-4, S-4, S-1	61				4.5061%
S-4, S-4, S-2	62				2.2686%
S-4, S-4, S-3	63				2.7252%
S-4, S-4, S-4	64				5.0230%

EXHIBIT 32.4 Percentage Shift in Home Prices

Period Time $T =$	
1	
Upshift	Downshift
0.135718489	-0.129712077

$0.00300321 + \sigma$ and the downshift value is the spot rate $0.00300321 - \sigma$. We show the resulting percentage shift in home prices in Exhibit 32.4.

What do we do to model home price movements from time $T = 1$ to time $T = 2$? We have two potential home price values as of time 1, $\text{Homereturn}(1, \text{up})$ and $\text{Homereturn}(1, \text{down})$. We have four different interest rate shifts from time 0 to time 1, though, so there are four potential mean shift values and, therefore, four upshifts in home prices and four downshifts in home prices from time $T = 1$ to time $T = 2$. We label them $\text{Homereturn}(2, \text{shift } 1, \text{up})$, $\text{Homereturn}(2, \text{shift } 1, \text{down})$, $\text{Homereturn}(2, \text{shift } 2, \text{up})$, $\text{Homereturn}(2, \text{shift } 2, \text{down})$, $\text{Homereturn}(2, \text{shift } 3, \text{up})$, $\text{Homereturn}(2, \text{shift } 3, \text{down})$, $\text{Homereturn}(2, \text{shift } 4, \text{up})$, $\text{Homereturn}(2, \text{shift } 4, \text{down})$. Using the four values for the one-period spot rate at time $T = 1$ and the same home price σ gives these eight potential changes in home prices from time $T = 1$ to time $T = 2$ in Exhibit 32.5.

EXHIBIT 32.5 Returns on Collateral by State and Collateral Value Shift

Period Time $T =$			
1		2	
Upshift	Downshift	Upshift	Downshift
0.135718489	-0.129712077	0.154368006	-0.111062561
		0.139625806	-0.12580476
		0.143687587	-0.121742979
		0.146325938	-0.119104629

EXHIBIT 32.6 Returns on Collateral by State and Collateral Value Shift

Period Time $T =$						
1			2			
Row Number	Upshift	Downshift	Shift Scenario	Upshift	Shift Scenario	Downshift
1	141.9648111	108.7859903	Upshift	163.8796359	Upshift	126.1978357
			1-Up		1-Down	
2			Upshift		Upshift	
			2-Up		2-Down	
3	125.5790667	96.70393968	Upshift	123.9753219	Upshift	95.10019488
			3-Up		3-Down	
4			Upshift		Upshift	
			4-Up		4-Down	
5	124.4171868	95.54205978	Downshift	124.7042024	Downshift	95.82907535
			1-Up		1-Down	
6			Downshift		Downshift	
			2-Up		2-Down	
7	124.7042024	95.82907535	Downshift		Downshift	95.54205978
			3-Up		3-Down	
8			Downshift		Downshift	
			4-Up		4-Down	

Let's assume that the home price was \$125 when the loan of \$100 was initially made, a loan to value ratio of 80 percent. What possible home values prevail at time $T = 2$? We apply the upshift and downshift to get two home price values at time $T = 1$. When we apply the eight shift percentages in Exhibit 32.5 to both of the potential $T = 1$ home prices, we have 16 possible values for home prices at time $T = 2$. Those values are shown in Exhibit 32.6.

The lowest home price value prevailing at time $T = 1$ is \$108.79 and at time $T = 2$ is \$95.10. Neither of these home price values is low enough to block rational prepayment or to cause strategic default. That is because the lender was prudent and would only lend \$100 against a \$125 house. You can see the scheduled loan principal balances in Exhibit 32.7, which is reproduced from Chapter 28.

What if the initial home price had been \$100 instead? The simulated home prices would have been 20 percent lower, at the levels shown in Exhibit 32.8.

The home prices simulation drop to a low of \$87.03 at time $T = 1$ and \$76.08 at time $T = 2$. At this point, the analysis of what the borrower will do gets complex, as shown in the shaded cells of Exhibit 32.9.

At time $T = 1$, would the borrower prepay on the loan, paying \$102, to live in a house worth \$87.03? Surprisingly, we don't know enough to answer the question. If the borrower is Warren Buffett and the home is his childhood home, the borrower would repay no matter what (because he is cash rich and because he wants to stay in that specific house). If the borrower is living paycheck to paycheck, he may only have enough money to make the scheduled payment of \$26.2624. The remaining funds that would be needed to prepay the loan (\$75.7376, the principal still outstanding after making the scheduled payment) would have to be borrowed. If the house is now worth only \$87.03, he would need to get a loan equal to 87 percent of home value. If lenders will now only lend 80 percent of home value, the homeowner does not have enough money to prepay.

What about time $T = 2$? The home, at its low point, is only worth \$76.08. If the borrower makes the scheduled payment of \$26.2624 at time $T = 2$, the remaining principal outstanding is \$50.99. If the borrower has the money to pay \$26.2624, it's rational to keep the house and make the payment. What if the borrower only has \$20? Before making the payment, the principal outstanding is \$75.7376. The house is worth only \$76.08. The borrower loses almost nothing by walking away from the loan and moving in with his brother. In lots of circumstances, this is what may well happen and that is why the multinomial logit probabilities of default and prepayment are such a critical overlay over rational prepayment and rational strategic default. They help predict very accurately rational acts by borrowers that appear "irrational"

EXHIBIT 32.7 Scheduled Amortization

Coupon Rate:	2.0000%
Time 0 Book Value	100
Periods to Maturity	4
Periodic Payment Amount	26.2624

Period Number	Initial Principal	Amount Paid	Interest	Principal	Ending Principal
0	100.0000				100.0000
1	100.0000	26.2624	2.0000	24.2624	75.7376
2	75.7376	26.2624	1.5148	24.7476	50.9900
3	50.9900	26.2624	1.0198	25.2426	25.7474
4	25.7474	26.2624	0.5149	25.7474	0.0000

EXHIBIT 32.8 Returns on Collateral by State and Collateral Value Shift

Row Number	Period Time $T =$					
	1			2		
	Upshift	Downshift	Shift Scenario	Upshift	Shift Scenario	Downshift
1	113.5718489	87.02879226	Upshift	131.1037087	Upshift	100.9582685
			1-Up		1-Down	
2			Upshift		Upshift	
			2-Up		2-Down	
3			Upshift	129.4294098	Upshift	99.28396966
			3-Up		3-Down	
4			Upshift		Upshift	
			4-Up		4-Down	
5			Downshift	130.1903562	Downshift	100.044916
			1-Up		1-Down	
6			Downshift		Downshift	
			2-Up		2-Down	
7			Downshift	99.18025752	Downshift	76.08015591
			3-Up		3-Down	
8			Downshift		Downshift	
			4-Up		4-Down	

EXHIBIT 32.9 Maturity of Cash Flow Received

Row Number	Current Time				
	0	1	2	3	4
1	0.0000	26.2624	26.2624	26.2624	26.2624
2		102.0000	26.2624	26.2624	26.2624
3		102.0000	26.2624	26.2624	26.2624
4		26.2624	26.2624	26.2624	26.2624
5			0.0000	26.2624	26.2624
6			0.0000	26.2624	26.2624
7			0.0000	26.2624	26.2624
8			0.0000	26.2624	26.2624
9			0.0000	26.2624	26.2624
10			0.0000	52.0098	0.0000
11			0.0000	26.2624	26.2624

(Continued)

EXHIBIT 32.9 *(Continued)*

Row Number	Current Time				
	0	1	2	3	4
12			0.0000	26.2624	26.2624
13			26.2624	26.2624	26.2624
14			26.2624	26.2624	26.2624
15			77.2524	26.2624	26.2624
16			26.2624	26.2624	26.2624
17				0.0000	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				0.0000	0.0000
46				0.0000	0.0000
47				0.0000	0.0000
48				0.0000	0.0000
49				26.2624	26.2624
50				52.0098	0.0000
51				52.0098	0.0000
52				26.2624	26.2624
53				26.2624	26.2624

EXHIBIT 32.9 (Continued)

Row Number	Current Time				
	0	1	2	3	4
54				26.2624	26.2624
55				52.0098	0.0000
56				26.2624	26.2624
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				26.2624	26.2624
62				26.2624	26.2624
63				26.2624	26.2624
64				26.2624	26.2624
Risk-Neutral Value =					101.4234

simply because we don't have complete information about the borrower (he only has \$20 and he can move in with his brother if he needs to).

How do we carry the analysis out to time $T = 3$? From the perspective of time $T = 2$, we know that there are 16 different possible home values. There are 16 different spot rates that could prevail, so there are $16 \times 16 = 256$ mean values of the home that could be simulated. We have an upshift and a downshift in home prices for each of these mean values, for a total number of potential home prices equal to $16 \times 16 \times 2 = 512$. From a computer science perspective, this is a small number and this is exactly what an enterprise-wide risk management system is designed to do.

CONCLUSIONS ABOUT THE IMPACT OF COLLATERAL VALUES

In the mortgage market during the 2006–2011 credit crisis, lenders and investors overlooked the huge impact on value that a collapse in collateral values can bring about. Others, typically agents originating but not keeping mortgages, knew about the risks but had no reason to care, since they did not keep the mortgages. Home price risk was identified as a key driver of the failure of U.S. banking companies by Jarrow and colleagues (2003) in their FDIC loss distribution model, but this basic fact of banking fundamentals was forgotten in one of the largest acts of mass amnesia in the history of finance.

For all types of loans and all types of collateral, simulating the interaction of collateral values and borrower behavior is critical to getting an accurate assessment of the Jarrow-Merton put option. We turn next to a related issue, the valuation of lines of credit.

Pricing and Valuing Revolving Credit and Other Facilities

The events of the 2006–2011 credit crisis and pervasive regulatory references to “exposure at default” highlight the need to thoroughly understand what drives the balances outstanding on revolving credit facilities. Once that understanding is gained, we can proceed to valuation and simulation in order to add revolving credit facilities to the list of instruments that can be valued in our ongoing quest to price the Jarrow-Merton put option of Chapter 1 as our primary measure of comprehensive risk.

We turn to that task in this chapter, but we start with a bit of history that we hinted at in the introduction to this book. Perhaps the most dramatic drawdown of a commercial paper backup facility was the August 16, 2007 announcement by Countrywide Financial that it was drawing down 100 percent of its backup facilities. One of the authors was literally in the head office of Countrywide on that date. The announcement was more than an admission that Countrywide was having trouble rolling over its commercial paper; it was an admission that Countrywide was out of options and this was a borrowing of last resort. As chronicled in the Introduction, Countrywide was acquired shortly thereafter by Bank of America, which in turn received massive injections of capital from the U.S. government in October 2008.

U.S. institutions were not the only firms that had to draw on back-up lines of credit during the credit crisis. As part of its support for troubled firms, the Federal Reserve established the Commercial Paper Funding Facility. Royal Bank of Scotland was a major borrower, taking down more than \$21 billion under this facility, as shown in the graph in Exhibit 33.1 provided by Kamakura Corporation.

Both Countrywide and Royal Bank of Scotland were forced to draw on these facilities because their credit risk had risen so high that the supply of funds to them from traditional sources was largely cut off. Exhibit 33.2, provided by Kamakura Corporation, shows the one-year default probabilities for Royal Bank of Scotland and Countrywide from January 1, 2007, through October 31, 2009. The Countrywide graph ends with its January 2008 acquisition by Bank of America. The one-year default probability for RBS exceeded 60 percent at its peak.

With this history as background, we plunge into our evaluation of loan commitments and revolving credits. As we will see, financial distress is just one of the many reasons why such a facility may be drawn down.




EXHIBIT 33.1 Borrowings from the Commercial Paper Funding Facility by Entities Sponsored by the Royal Bank of Scotland Group, October 1, 2008, to October 31, 2009
Sources: Kamakura Corporation; Federal Reserve.




EXHIBIT 33.2 Royal Bank of Scotland Group

ANALYZING REVOLVING CREDIT AND OTHER FACILITIES

In Chapter 30, we dealt with the risk measurement and valuation of nonmaturity deposits in detail. Those instruments (and their asset side counterparty, retail charge card balances) have balances that fluctuate up and down randomly at the option of the depositor. The accounts have no explicit maturity, although the Jarrow-van

Deventer methodology (1998) implicitly assumes that the deposit franchise has a finite ending date. On both the balance side and the interest rate side, the non-maturity deposit formulation of Jarrow and van Deventer (1998) explicitly incorporates open-market interest rates as a major determinant of deposit balances and deposit interest rates.

Can we use this kind of analysis to value revolving credits? On the surface, there are many parallels. Deposit balances and revolving credit balances can be zero for an extended period of time and then suddenly spring into use. Similarly, other deposits will show a randomly fluctuating balance over a long period of time just like an “evergreen” line of credit.

It is tempting to use the Jarrow-van Deventer methodology for nonmaturity deposits if it were not for one major assumption they make that doesn’t hold in this case—they assume that the borrower of the nonmaturity deposits (the bank) is riskless because of government deposit insurance. In a parallel way, they assume that the credit risk of charge card borrowers is nonzero, but known with certainty. Unfortunately, nothing could be further from the truth in the case of revolving credits. In what follows, we draw on the insights of Chava and Jarrow (2003, 2008).

FLUCTUATING CREDIT RISK AND REVOLVING CREDIT DRAWDOWNS

The changing credit quality of a borrower, along with randomly fluctuating cash needs that are linked in part to credit quality, is at the heart of the analysis of revolving credits. Cash needs are fluctuating for various reasons, above and beyond the credit distress we discussed above:

- The fundamental business of the borrower may have highly seasonal revenues (as in retailing, agriculture, and tourism) but expenses that are nonseasonal.
- The fundamental business of the borrower may have very steady revenues but highly seasonal cash needs (like farming).
- Cash flows of the business are not perfectly predictable due to normal uncertainties of client payments on the asset side and payments to suppliers on the liability side.
- Scheduled payments may not occur due to operational issues (a wire transfer failure) or credit quality problems on the part of either the company or its clients.
- The deteriorating credit quality of the borrower may shut off some sources of funding that may need to be replaced with a bank line of credit.
- Conversely, improving credit quality on the part of the borrower may open up new sources of funding that decrease the need for usage on the line of credit.

All of these factors contribute greatly to the need for a special understanding of revolving lines of credit, but it is the last three factors that are most important. If the only factors driving line of credit usage were the first three factors, we could use the Jarrow-van Deventer approach to nonmaturity deposit valuation in this context without hesitation. Alternatively, we could simply project the drawdown dates and implied forward interest rates and use the Heath, Jarrow, and Morton (HJM) discounting of Chapter 9 and 19 to calculate valuation directly.

INCORPORATING LINKS BETWEEN CREDIT QUALITY AND LINE USAGE

In van Deventer and Imai (2003) there is an extensive discussion of the JPMorgan Chase–SK Securities incident of 1998 where the amount of money owed to JPMorgan Chase by SK Securities and the default probability of SK’s special purpose vehicle borrowing the money were strongly correlated. The very nature of the revolving line of credit has this issue at its heart—one of the primary uses of proceeds and one of the primary motivations of line of credit usage is for commercial paper backup lines. These lines are intended to provide an alternative source of funding in the event of a disruption in the commercial paper market in general or in case the borrower is unable to issue commercial paper, presumably because of credit quality problems. The Countrywide and Royal Bank of Scotland examples are just two of very many such incidents in the 2006–2011 credit crisis.

Providers of revolving lines of credit walk a fine line—they know that this change in credit quality is the primary motivator of the line and the lender guesses that the potential deterioration in credit quality will not be so severe that the lender is not fully compensated for the risk they are taking. It goes without saying that the default probabilities of all borrowers with lines of credit and the usage under the lines of credit are driven by a set of macroeconomic factors common to many of the borrowers. We discussed this issue extensively in Chapters 19 and 20. In both the Countrywide and Royal Bank of Scotland cases, it was home price declines that triggered a severe cash need by both firms.

The cash needs under the line of credit are the sum of a number of influences:

$$C(t) = C_1(t) + C_2(\lambda) + C_3(\lambda_1, \lambda_2, \dots, \lambda_n) + Z(t)$$

Cash needs at any point in time consist of a time-dependent expected amount, which can vary seasonally, plus an amount that depends on the default intensity of the borrower λ , plus a third amount that depends on the default intensities of the major counterparties of the business, which we write $\lambda_1, \lambda_2, \dots, \lambda_n$. The final term is a random term. We don’t mean to suggest any particularly functional form for these cash flow needs. In Chapters 36 and 47, we suggest precisely how these cash flow needs can be modeled and forecast.

Unfortunately, the multiperiod nature of the usage on a line of credit means that the Merton model we reviewed in Chapter 18 is completely unsuited as a framework for line of credit valuation and estimating exposure at default. What we need instead is the kind of continuous time variation that is described by a function like that above or in the Jarrow-van Deventer nonmaturity deposit model.

Once we specify how the cash needs of a company vary with its own default risk and the relevant macroeconomic factors, we can proceed directly to valuation of the line of credit, simulation of cash needs, expected exposure at default, and the full probability distribution over time and in dollar terms of exposure of default. This analysis is completely analogous to the “own firm” risk assessment that we explore in Chapters 36 and 37.

IS A LINE OF CREDIT A PUT OPTION ON THE DEBT OF THE ISSUER?

Lines of credit are often thought of as a put option on the floating rate debt of the borrower. This view is almost, but not exactly, correct. Using the analogy with

the nonmaturity deposit modeling of Chapter 30, part of the cash needs listed in the previous section has nothing to do with credit risk. They typically reflect the seasonality of the business. Given random interest rates and a constant default intensity of the borrower λ , this noncredit risk-related cash need can be valued directly in a closed form solution using the HJM examples in Chapters 9 and 19. The second component of cash need comes from potential changes in the credit quality of the borrower's major business counterparties. In a sense, the portfolio of accounts receivable that are held by the borrower are a collection of debt instruments like those we modeled in Chapter 19 on credit risky bonds. These too have a known solution in the reduced form modeling framework; see the worked examples in Chapter 19.

In some cases, the credit quality of the business counterparties of the borrower affects cash flow in a more complex way than a default/no default kind of model. A rise in default probability of ABC Company, for example, may lengthen the terms of payment, for example, without triggering default. These kinds of influences of the default probability of ABC Company can be modeled like small portfolios of digital default swaps that are triggered, not so much by the event of default, but by the level of the default probability reaching a certain range. The graph of the default probabilities of Countrywide and Royal Bank of Scotland, in fact, could be used to (with hindsight) do exactly this.

The final default intensity-related component of cash needs of the company is the term $C_2(\lambda)$ in the previous cash need function. This function is undoubtedly a complex nonlinear function of the default probability that reflects sources of funding like the commercial paper market. As we discuss in Chapter 37, issuers of commercial paper rated A3/P3 have dramatically less access to commercial paper than more highly rated issuers. As the Countrywide and RBS examples show, market participants don't need to wait for a ratings change to withdraw their interest in the commercial paper of a given issuer.

Similarly, a much larger universe of institutions can buy investment grade bonds (we prefer the term "low-default probability bonds") than can buy noninvestment grade bonds (or "high-default probability bonds"). All of these factors affect the function $C_2(\lambda)$. Another factor comes into play as well. Even if the company holds relatively liquid assets, it may become obvious to other market participants that the company is in trouble if the default probability exceeds a certain level. The infamous Long-Term Capital Management incident in 1998 illustrates the result—in a market that is not perfectly liquid, other market participants will lower their bid prices for the assets held by the company because they know credit problems will force the company to take any price. (We return to this theme in detail in Chapter 37.) In summary, the cash flows generated by the $C_2(\lambda)$ term have a nature much like that of a range floater in credit derivatives terms. For a given range of default intensity, cash flows will take on a particular value. In the Jarrow model, the straightforward nature of the default intensity function gives us a good chance of a closed form valuation for this component of draws on the line of credit. Chava and Jarrow (2003, 2008) use nested logistic regressions to model the probability of a credit crisis for ABC Company and, given that the liquidity crisis has occurred, the probability of a default in their valuation of lines of credit. (The analysis parallels the analysis of irrational exercise of fixed income options in Chapter 28.) Instead of modeling partial prepayments, however, we are modeling partial drawdowns of the line of credit.

Multinomial logistic regression is the best practice approach to modeling events like these where there are n possible outcomes, not just two outcomes (i.e., no drawdown or 100 percent drawdown).

The final component of cash flow needs is the random term $Z(t)$. If we have properly included all of the macro factors driving cash flows in the other terms of the cash flow function, then we can make an argument that $Z(t)$ has mean zero and that diversification arguments with respect to the $Z(t)$ term's impact on valuation will prevail. This is an insight of Jarrow, Lando, and Yu (2005).


The upshot of this microfocused discussion of the determinants of cash needs on a line of credit is that a finely detailed analysis of cash needs will allow us to do a high quality valuation and simulation of lines of credit and exposure at default. Most of what we need to do has already been discussed in other chapters in this book. What little we need to add to do a complete job on lines of credit risk analysis will be discussed in detail in Chapters 36 and 37. We next turn to equity-linked instruments in our quest to value the Jarrow-Merton put option.

Modeling Common Stock and Convertible Bonds on a Default-Adjusted Basis

A very large proportion of the world's financial institutions are substantial direct or indirect owners of common stock and convertible bonds. Financial institutions, via their own pension funds, often have a large exposure to equities that is overlooked in risk analysis. In addition, the direct investment in equity-related securities is large and linked in a complex way with other assets and liabilities on the balance sheet.

Many sophisticated investors have major concerns with the inconsistency in their management of fixed income and equity portfolios. Credit risk is typically given lots of attention on the fixed income side and ignored on the equity side. Risk analysis of a position in IBM bonds and IBM common stock is typically done in a way that recognizes the losses of default on the bonds and ignores the simultaneous impact of default on the common stock of IBM. This inconsistency should be very troubling to management, boards of directors, and regulators of the financial institutions involved. The inconsistency stems from the largely independent development of financial theory for fixed income securities and for equities. In the latter case, a fundamental premise of the capital asset pricing model and arbitrage pricing theory is that the returns on common stocks are normally distributed. (Model Risk Alert: This assumption is false.) Exhibit 34.1, provided by Kamakura Corporation, shows that there is a 0.32 percent probability that the stock price of Citigroup is zero at the end of one year. This is the best estimate, on an annualized basis, that Citigroup will be bankrupt over that time horizon, using the reduced form models of Chapter 16. The other maturities in Exhibit 34.1 show the bankruptcy probabilities, all on an annualized basis, at different time horizons.

The existence of bankruptcy is inconsistent with the assumption of normally distributed returns. While bankruptcy risk was explicitly noted in the early models of Merton (1973), it has been largely ignored in practice. This has massive implications for measured risk. For example, if one calculated the mean monthly returns and volatility of monthly returns for Lehman Brothers and Bear Stearns from 1990 to 2007, they would have implied that the probability of a -100 percent return in one month (due to bankruptcy) was 0.000000 percent. The failure of both firms in 2008 shows that the explicit recognition of bankruptcy in the modeling of common stock and convertible bonds is essential for realism.

**EXHIBIT 34.1** Citigroup Inc.

Credit risk on the equity side has another more subtle form as well, a topic we introduced in Chapter 2. Many fund managers are judged on their performance versus an equity benchmark, such as the S&P 500 index in the United States. If a company is a member of the list of 500 companies making up the index, it has to be of fairly good credit quality. If the company's default probability rises to a certain level, the odds that it is removed from the S&P 500 index become very high. If the company is to be dropped, the announcement is normally made after the close of business. The stock will drop dramatically at the opening the next day, leaving the fund manager with a large loss versus the index because the weakened company is no longer in the index, but the fund manager still has the position from the previous day. This credit risk-related loss can be substantial and has also been almost completely overlooked by many financial institutions. Campbell, Hilscher, and Szilagyi (2011), in a paper given the Markowitz award for the best paper in the *Journal of Investment Management*, show that high default risk equity portfolios persistently underperform other equity portfolios after making all of the standard industry adjustments for risk-adjusted return calculations.

With this as background, we move on to a discussion of how to incorporate common stock and convertible bonds on the list of instruments that we can value as we progress toward our objective—valuing the Jarrow-Merton put option on company assets and liabilities as the comprehensive, integrated measure of interest rate risk, market risk, liquidity risk, foreign exchange risk, and credit risk.

MODELING EQUITIES: THE TRADITIONAL FUND MANAGEMENT APPROACH

For more than 40 years, it has been the fund management industry standard to model the return on a common stock as a linear combination of the return on various other factors. This approach is justified by assuming that stock prices are lognormally

distributed and, therefore, their returns are normally distributed. A linear combination of normally distributed inputs produces the desired normal distribution. This approach dates to the original formulation of the capital asset–pricing model, where the original single factor was the return on the market. In this model, the return on the common stock of company J is the sum of the risk-free rate and its beta times the return on the market less the risk-free rate plus a random error term. For decades, there has been an enormous amount of research on correctly measuring the correlation between the returns on the common stock of company J and the return on the market. The beta is the measure of this correlation of returns.

Fund managers quickly realized that there were multiple factors driving common stock returns, some of which were macroeconomic factors (like those in the reduced form credit models of Chapter 16) and some of which were company specific. The company-specific factors included things like company size, its price earnings ratio, its industry sector, and so on. Many of these factors, too, are incorporated in the reduced form default probability models of Chapter 16. It is fairly common for a fund manager to be using a multifactor model of common stock returns to manage tracking error, the amount of potential deviation between the return on the fund manager’s portfolio and the return on the index.

What is wrong with this approach?

It totally ignores both the direct and indirect impact of default. Clearly, when a company has a 1 percent probability of bankruptcy, a normal distribution can’t be an accurate portrayal of the potential distribution of returns on the common stock, because there is a 1 percent probability that the return will be minus 100 percent!¹ Moreover, in the indirect risk category, there is a large drop in the common stock price that would occur if the company is dropped from the S&P index for credit risk–related or other reasons.

The great strength of this approach, however, is that it already incorporates many of the macroeconomic factors and company-specific factors that are related to default risk—they are just not used properly for default-adjusted tracking error management. We turn now to general assumptions about common stock price movements in the derivatives world.

MODELING EQUITIES: THE DERIVATIVES APPROACH

In the derivatives world, key assumptions about the movement in equity prices also employ the lognormal assumption about stock price movements that results in normally distributed returns on common stock. As we noted in our term structure analysis of Chapters 3 through 10 and in Chapter 18 on the Merton model of risky debt, the analytical benefits of this assumption can be so powerful that it is worth the cost of glossing over some real-world realities.

Typically, in the derivatives world, the stock price itself is used as a random factor without modeling the N macroeconomic and company-specific factors driving movements in the stock price. This is due in large part because the primary focus of derivatives analysts, dating from the original options model of Black and Scholes (1973), is on synthetic replication of the derivative with an appropriate position in the common stock. Jarrow and Turnbull (1996) discuss the importance of this argument in detail.

It is argued that because movements in the price of the derivative can be hedged by appropriate positions in the common stock, we don't need to know the drivers of returns on the common stock.

For most of the past 40 years, equity derivatives valuation formulas have largely assumed away both the potential default of the counterparty on an option on common stock and the potential default of the issuer of the common stock underlying the derivative. Lately, the issue of these two potential defaults has been getting a lot more attention in the literature and the real world for two reasons. First, the 2006–2011 credit crisis has made it obvious that default risk is a massive risk management issue. Second, because of the insights triggered by the Jarrow and Turnbull (1995) reduced form credit model and research by many other talented researchers, it is now well understood how to deal with credit risk issues explicitly.

MODELING EQUITIES: A CREDIT RISK-ADJUSTED APPROACH

The Merton (1974) model of risky debt, which we discussed in Chapter 18, is not just a model of debt prices. It is implicitly a single-factor model of equity prices where the value of a company's common stock is a function of the value of company assets. Similarly, the Shimko, Tejima, and van Deventer (1993) model implies that the value of a company's common stock is a function of a one-factor Vasicek interest rate term structure and the value of company assets, which in turn can have any arbitrary correlation with interest rates. As we discussed in Chapters 16, 17, and 18, tests of the consistency of these models with relative movements in common stock prices and credit spreads on bonds have proved to be disappointing. Consequently, the Merton insights have not been used by equity market participants to the extent they have been used in the debt markets.

Jarrow and Turnbull (1995) present a framework for modeling equity securities on a full credit-adjusted basis. Using a simplified assumption of a constant default intensity, they assume that stock prices are lognormally distributed and driven by random interest rates and a company-specific risk factor, which can have any correlation with interest rates, provided the company does not go bankrupt. They assume that the stock price drops to zero in the event of bankruptcy. Jarrow and Turnbull show that, under these assumptions, the risk-neutral drift term in the common stock's return formula has to increase by the instantaneous probability of default. This has a number of important implications for modeling not only the common stock but also options on a credit-risky company's common stock and convertible bonds, which we discuss below.

Using the Jarrow-Turnbull framework for modeling common stock returns, we can directly address all of the concerns expressed in the first sections of this chapter while preserving (in modified form) the multifactor approach to explaining lognormally distributed returns.

Jarrow (1999) takes this approach still further and models common stock in a framework with a random default intensity driven by interest rates and macro factors as discussed in Chapter 16. A common stock pays dividends until bankruptcy, when the value of common stock goes to zero. Amin and Jarrow (1992) provide the general framework for the valuation of risky assets and derivatives on them in an economy where the risk-free term structure is driven by n random factors (like we analyzed in

Chapters 6 through 9) and d other random factors. Our analysis of foreign exchange risk (Chapter 31) and the value of collateral (Chapter 32) take advantage of these insights. Amin and Jarrow show that the excess return on a risky asset over the instantaneous spot rate of interest is a function of these d risk factors, the volatility δ of the risky asset's return with respect to each risk factor, and the market price of risk n associated with each risk factor:

$$\mu(t, x) - r(t) = - \sum_{i=1}^d \delta_i(t, x) n_i(t)$$

If this constraint were not imposed on the risky asset's expected return, one could achieve excess risk-adjusted returns by buying the risky asset and hedging each of the risk factors. If one estimates the excess return over the risk-free rate using one-year returns on the common stock of Citigroup, the joint importance of random interest rates and other macro factors is extremely obvious. For expositional purposes, we use 244 overlapping annual intervals from January 1991 to April 2011 and fit a linear relationship between the excess return variable and the change in one-year U.S. Treasury yields. The excess return is calculated by subtracting the one-year U.S. Treasury yield at the beginning of the one-year interval from the one-year return on Citigroup common stock. We ignore dividends and the many econometric problems with this specification to emphasize more important aspects of the results, which are summarized in Exhibit 34.2.

First, the change in the one-year interest rate is statistically significant with a coefficient that is large in magnitude. It is not a surprise to find that interest rates are an important driver of the common stock returns of Citigroup. Second, interest rates explain only 7.6 percent of the total variation in Citigroup common stock returns. Other macroeconomic factors make a large contribution to returns as does the non-zero default probability of the firm. This default probability, in turn, depends on an overlapping set of macro factors such as those we discussed in Chapters 19 and 20. The annual standard deviation of these missing variables is an extremely large 0.34.

Jarrow (1999) shows how to explicitly incorporate the possibility that the underlying risky asset could fall to zero value with a default intensity λ , consistent with the Jarrow and Turnbull (1995) reduced form model and our exposition in Chapter 16.

A full exploitation of the virtues of the Jarrow and the Amin and Jarrow (1992) approach has been embedded in state-of-the-art risk management systems that we describe in Chapters 36 through 41. In this chapter, we restrict our detailed comments on equity derivatives on the common stock of a credit risky company to the Jarrow and Turnbull (1995) and closely related approaches.

OPTIONS ON THE COMMON STOCK OF A COMPANY THAT CAN GO BANKRUPT

Jarrow and Turnbull (1995) use their model for the common stock of a company that can go bankrupt, to value put and call options on that credit-risky common stock.

EXHIBIT 34.2 Citigroup Inc. Regression Statistics

Regression Statistics									
Multiple R		0.282568644							
R Square		0.079845038							
Adjusted R Square		0.076042745							
Standard Error		0.340355258							
Observations		244							
ANOVA									
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>				
Regression	1	2.432580696	2.432580696	20.99917957	7.36384E-06				
Residual	242	28.03369182	0.115841702						
Total	243	30.46627252							
		Coefficients	Standard Error	<i>t</i> Stat	<i>P</i> -value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%
Intercept		-0.010609215	0.022581687	-0.469815013	0.638910254	-0.055090964	0.033872533	-0.055090964	0.033872533
Change in 1-Year U.S. Treasury Yields		6.831973477	1.490887962	4.582486177	7.36384E-06	3.895199806	9.768747147	3.895199806	9.768747147

They reach a very powerful and yet intuitive conclusion. The zero-coupon bond prices in the call options formula become the zero-coupon bond prices extracted from the yield curve of the risky company (see Chapters 5, 16, 17, and 19 for how to do this) instead of from the risk-free yield curve. These risky zero-coupon bond prices directly reflect the default intensity of the issuer, as we discussed in Chapters 16, 17, and 19.

Using the Jarrow and Turnbull (1995) formula for options on the common stock of a firm that can go bankrupt, we can explain much of the volatility smile issues with the constant interest rate/no default Black-Scholes options model. As many analysts have noted, equity options prices seem to be based on a probability distribution that reflects much greater downside risk than a simple lognormal distribution on stock price would imply. This is due to the risk of default (among other things), and the Jarrow-Turnbull model will improve the volatility smile problems because it deals with this increased downside risk directly. Moreover, the Jarrow-Turnbull model can be used to imply the probability of default from equity options on a company that can go bankrupt.

CONVERTIBLE BONDS OF A COMPANY THAT CAN GO BANKRUPT

The same approach proposed by Jarrow-Turnbull (1995) can be applied to convertible bonds, which contain a complex set of options controlled by both the issuer and by the holder of the bonds. After many years of trying to model convertible bonds as a combination of the straight bond and an option on the common stock, hedge funds and other analysts have employed both structural and reduced form modeling approaches. Three recent papers by Ayache, Forsyth, and Vetzal (2003); Andersen and Buffum (2003); and Bermudez and Webber (2003) show great promise in tackling one of the most complex securities on the balance sheet of financial institutions. Ultimately, one should use a bushy tree like the three-factor HJM example, coupled with a bushy tree of stock prices much like the collateral bushy tree of home prices that we employed in Chapter 32 on the impact of collateral.

Convertible bonds normally contain a call option granted to the issuer to call the bonds at a preset schedule of call prices. Separately, the holder of the convertible bonds has an option to convert the bonds normally to a preset number of common shares. This option is maintained for a prespecified number of days even if the bonds are called by the issuer to allow time for conversion. Both of these options are significantly affected by the probability of default, which is why initial approaches that ignored default led to large losses at hedge funds investing in convertible bonds.

Many of the recent papers on convertible bond valuation combine reduced form default intensities with the finite difference method of securities valuation that we discussed in earlier chapters. Andersen and Buffum (2003) note that “While it is in principle possible to build convertible bond models using the structural approach . . . the reduced-form approach is, by far, the most natural for trading applications and shall be the sole focus of this paper” (page 2). At this stage it is safe

to say that tremendous progress is being made on the valuation of convertible securities on a full credit-adjusted basis and they can be included in valuation of the Jarrow-Merton put option, our integrated measure of credit risk and interest rate risk.

NOTE

1. We note that, even in the event of bankruptcy, in many cases common stock prices do not drop to zero in large part because of the common violation of absolute priorities of debt holders over equity holders. This is another of the many challenges facing analysts using the Merton model of risky debt from Chapter 18.

Valuing Insurance Policies and Pension Obligations

Our final chapter in individual instrument valuation and simulation is focused on insurance policies and pension obligations, since insurance companies and pension funds make up a substantial proportion of the world's financial institutions. To reiterate, our objective in reviewing valuation and simulation methodologies is to incorporate the insurance liabilities and pension liabilities in the Jarrow-Merton put option on the financial institution's assets and liabilities. We want to use a consistent methodology to calculate this comprehensive measure of integrated interest rate risk, market risk, liquidity risk, foreign exchange risk, and credit risk.

This is particularly interesting to discuss for insurance and pension obligations because the links with what we have already covered are so strong and yet so few members of the financial community would acknowledge that that is the case. We hope that this chapter assists in narrowing this perception gap.

First, we start with life insurance policies and then move to pension obligations and property and casualty insurance.

LIFE INSURANCE: MORTALITY RATES VS. DEFAULT PROBABILITIES

The analogy between the mortality rate on a life insurance policy and the default probability on a bond of a particular issuer is very strong. Exhibit 35.1 shows annual mortality rates for a 50-year-old male (upper line) and 50-year-old female (lower line) taken from the U.S. Society of Actuaries RP-2000 mortality tables.

The mortality rate graph is surprisingly similar in shape and meaning to the evolution of the one-year default probability for Tokyo Electric Power prior to its nationalization by the Japanese government on May 8, 2012.

Like default or bankruptcy, mortality on a life insurance policy is a zero/one kind of event. The fundamental approach to measuring mortality risk was proposed by D. R. Cox (1972). The "Cox process" that he developed is in fact the same Cox process used in reduced form modeling in the Jarrow (1999, 2001) and Duffie and Singleton (1999) credit models that we discussed in Chapter 16. The mathematics is identical. The term structure of mortality rates is very directly related to the term structure of default probabilities that we can observe for Tokyo Electric Power Corporation on May 8, 2012, shown in Exhibit 35.2. Default probability changes over time are shown in Exhibit 35.3.




EXHIBIT 35.1 Society of Actuaries RP-2000 Mortality Rates for Male and Female “Combined Healthy” Annuitants, Ages 50 to 119




EXHIBIT 35.2 Tokyo Electric Power Co Inc./Term Structure

This parallel is so strong and so obvious to financial academics that Jarrow (2012) described the relationship between credit default swaps, which we discussed in Chapter 20, and life insurance in this way in his paper “Problems with Using CDS to Infer Default Probabilities”:

It is commonly believed that credit default swap (CDS) implied probabilities provide unbiased estimates of a corporation’s or sovereign’s actual default



EXHIBIT 35.3 Tokyo Electric Power Co Inc./Profile

probabilities or credit risk. The purpose of this paper is to show why this common belief is false using an intuitive analogy between CDS and life insurance premiums. For life insurance, the difficulties of using insurance premiums to imply the mortality probability are crystal clear. The difficulties are unraveling the impact of risk premium, counterparty risk, market frictions, and strategic trading. These difficulties are well known in the academic literature (references are provided below). Of course, all is not lost because there is no need to use these implied probabilities. Readily available actuarial based mortality tables are used to determine mortality probabilities. And insurance premium implied mortality probabilities are not used for risk management nor in regulation.

The other strong parallel between life insurance and credit risk, as Jarrow implies, is the progressive use of logistic regression to determine the probability of the event (bankruptcy as in Chapter 16 or mortality in this chapter) with more precision than we would get from a simple table (a transition matrix with default probability by ratings grade or a mortality table like the graph in Exhibit 35.1). A recent paper by Guizhou Hu (n.d.) shows how mortality rate prediction can be improved by incorporating explanatory variables that are much more detailed than those in mortality tables (age, male or female, smoker or nonsmoker); among the variables he mentions are the following:

- Systolic blood pressure
- Total cholesterol
- Body mass index
- Smoker or nonsmoker
- Albumin
- Physical activity—i.e., light, moderate, or heavy
- Income

All of these input variables are specific to the insured, just like the accounting and stock price inputs to the Jarrow-Chava model of Chapter 16 are specific to the firm being modeled.

Cyclicalities in Default Probabilities and Mortality Rates

Logistic regression can also confirm the importance of cyclicalities in mortality just as it does (as required by the Basel Capital Accords) for default. Any Japanese banker can tell you that the mortality rate of his co-workers is much different when the Nikkei stock index is below ¥10,000 than it was when it was nearly ¥39,000 in December 1989. Similarly, just as there is seasonality in the default of small businesses, there is seasonality in mortality that is logical and predictable (due to factors like weather, stresses of the school calendar year, low stress of summer vacations, and so on).

Valuing Life Insurance Policies

Life insurance policies are closely related to digital default swaps that we discussed in Chapter 20. The protection buyer on the digital default swap called a “life insurance policy” is (usually) the insured. The insured is also the reference name on the digital life insurance policy. The counterparty on the digital life insurance policy is the life insurance company. The amount paid on the digital life insurance policy in the event of default is the policy amount. The only significant difference is the maturity—term life insurance policies, which have an explicit maturity, are almost literally identical to digital default swaps. A whole life policy runs until mortality or until canceled. Ignoring the cancellation provisions, it’s like an infinite maturity digital default swap except that we know the term structure of mortality/default probabilities approaches 100 percent in time. In fact, the graph for age 120 in Exhibit 35.1 would have been 100 percent, the amount specified in the Society of Actuaries RP-2000 tables for that age.

What about cancellation privileges? This is a real option of the insured that can be exercised rationally or irrationally as discussed in the prepayment context in Chapters 27 to 29. It is rational to cancel when (1) the mortality risk of the insured becomes lower than that implied in the current policy, allowing the insured to reinsure at a lower rate; or (2) when the insured becomes so wealthy that they no longer need insurance (they save money by not paying the insurance company costs embedded in the policy and effectively self-insure). Irrational cancellation (like irrational prepayment) is usually for a good reason unknown to the insurance company. The financial condition of the insured, for example, may have deteriorated so much that they can’t afford the premiums, even though they need the insurance.

What about investment options on a whole life policy? We know from finance theory that an option to switch from one instrument at its market price to another instrument at its market price has zero value, since there is no benefit to the options holder (if we ignore transactions costs). Many of the investment options in a life insurance policy have this nature. To the extent there is a real benefit to the holder of the insurance policy from the investment options offered, we can value them using

the techniques in other chapters in this book using the default-adjusted approaches from Chapter 19 (corporate bonds) or Chapter 20 (default swaps or CDOs).

One of the key benefits of life insurance to the buyer is the ability to defer taxes on the investment income of a whole life insurance policy. This tax deferral is of course a real benefit to the buyer of the insurance policy, relative to making the identical investments in the name of the insured. The value of this benefit to the buyer of the insurance policy depends on the state of the tax code now and in the future, the source and variability of income now and in the future, and the default probability of the insurance company. Clearly, valuing these tax benefits with precision is a challenge, and one that we're happy to defer to another time.

PENSION OBLIGATIONS

Defined benefit pension obligations have a very interesting mix of terms that combine part of the life insurance problem with more traditional financial instruments. The U.S. Social Security System is exactly such an obligation.¹ One advantage of a pension fund in an economic sense is the ongoing stream of current payments for future pension benefits from both employees and their employers. The present value of this cash inflow from any one employee, say Mr. Jones, depends on three factors: (1) the mortality probability of Mr. Jones (he could be unlucky enough to pass away before leaving employment and receiving pension benefits); (2) the probability that Mr. Jones will change jobs and end up contributing to a different pension fund; and (3) the probability of default of Mr. Jones' employer. The liabilities of a pension fund, the future payments to the beneficiaries of the fund, depend on the mortality rate of the pension beneficiary, just like a life insurance policy. The difference is that the payments are made in the opposite circumstance of a life insurance policy—they are paid while the beneficiary is still living, and a “default event” terminates this stream of cash flows just like a mortgage default terminates payments on a mortgage held by a bank. In this sense, the future streams of pension payments are each a digital default swap of sorts again, except the key probability used in valuation is the probability of no mortality prior to the payment, not the probability of bankruptcy/mortality relevant to valuing a life insurance policy.

For defined benefit pension plans, which is our focus in this section, there are normally a complex set of rules or circumstances that determine when and how pension benefits are increased over time. Similarly, even with government pension insurance, there is the probability that the pension fund defaults on its obligations to the pension beneficiaries. Both of these complexities can be dealt with using a combination of the previous life insurance analysis and the techniques of other chapters. In particular, the increased use of the logistic regression technique and reduced form modeling approach allows for comprehensive credit-adjusted risk management of both the assets and liabilities of a pension fund.

One of the greatest fallacies in the world of legacy pension analytics is the assumption of a portfolio return on the assets of the pension fund. This assumed return has huge political implications, in the case of a government-sponsored pension fund, if an analysis of the fund uses this assumed return in the calculation of whether or not the pension fund is “underfunded.” A more rational analysis marks to market

both the assets and the liabilities of the pension fund using the techniques of this book. Discounting the mortality adjusted cash flows of the liabilities at the risk-free rate is the liability risk-neutral value on the assumption that the liabilities must and will be paid. Discounting the liabilities at any higher discount rate implies it is economically and politically acceptable for the pension fund to default on its obligations to pensioners. Similarly, any form of valuation on the asset side that allows assets to be marked at a higher value than their market value (like the previous assumed return) is a thinly disguised ploy to overstate the financial health of the pension fund.²

PROPERTY AND CASUALTY INSURANCE

Property and casualty insurance differs from life insurance in two dimensions: the term of the insurance contract is typically shorter and the events being insured are quite diverse, ranging from auto insurance to catastrophic weather-related events. Aside from events that are acts of nature, like weather and earthquakes, logistic regression again plays a powerful role in both pricing and risk management of liabilities on a property and casualty insurance contract. The price of an auto insurance policy depends on factors such as the following:

- The age of the driver
- The sex of the driver (in some jurisdictions)
- The health of the driver (if the data is available)
- The income of the driver (again, if the data is available)
- The driving history of the driver
- The amount of driving to be done (and the possibility that the true amount is not disclosed)
- The type of car being insured (i.e., a bus or a sports car)
- The value of the car being insured
- The place in which the car will be stored (a deserted urban parking lot in a high-crime area or a locked garage in a gated community)
- The place in which the car will be driven (i.e., Manhattan or rural France)

In addition, it is an empirical question whether macroeconomic factors have an impact on both the probability of an insured event and the payout on the occurrence of the event. Clearly, we know from the statistics on new auto sales that new sales of cars decline in bad times so the average age of all cars being driven increases during recessions. This in turn would be expected to affect the probability of a claim and the payout on the claim. For this reason, depending on the nature of the insured event, it is often the case that macro factors drive the probability of an insurance payout just like they affect the probability of default in the reduced form credit models of Chapter 16. When this is the case, it is essential that these macroeconomic factors are incorporated in policy pricing and total balance sheet risk management or the risk of providing the policy can be dramatically underestimated. If these macro factors are relevant, they will increase the correlation in the occurrence of insured events and the fat tails of the loss distribution.

THE JARROW-MERTON PUT OPTION

In many insurance companies and pension funds, there is a sharp division between the investment or asset side of the organization and the insurance or liability side of the organization. The investment function is dominated by finance experts and the liability side is dominated by insurance experts and actuaries. This contrasts sharply with risk management at banks and securities firms where asset and liability management on a joint basis has been standard for 40 years.

Looking at the risk on only one side of the balance sheet at a time is dangerous because common macro factors drive the risk on both sides of the balance sheet. If you are looking at only half the picture, it is impossible for management, the board of directors, and regulators to have an accurate view of total risk. How can we value the Jarrow-Merton put option on the assets and liabilities of the firm as a comprehensive integrated measure of interest rate and credit risk if half of the information is missing or based on a completely different and totally inconsistent set of assumptions (e.g., assuming that the assets of the pension fund will return a compounded rate of 12 percent annually for the next 30 years)?

Perhaps the most tragic recent example of the damage caused by these “Great Walls of China” in the insurance industry is the joint impact of the 2011 Sendai tsunami and Fukushima nuclear disaster in Japan. Japanese life insurance companies were very heavily invested in the bonds of Tokyo Electric Power Company at the same time mortality rates and property and casualty damages rates skyrocketed in the aftermath of the tsunami. If the Japanese government had not rescued Tokyo Electric Power, the default probabilities for TEPCO shown above make it clear that a “double whammy” of wrong way risk on the asset and liability sides would have occurred.

Fortunately, as we have seen in this chapter, the mathematics of credit risk models is taken directly from insurance expert Cox and there is no conceptual barrier to integrated risk management in the insurance industry. There is literally no difference in the mathematical approach, so there are only emotional and political reasons for the long-standing divide between the actuaries and the investment side of the organization when it comes to risk management. The best-practice insurance companies are solving this problem by moving actuaries to the investment side of the organization and investment managers to the insurance side of the organization. Integrated risk management teams that have both actuaries and finance experts are being established and a common view of integrated risk is widely recognized as necessary and desirable. Insurance firms are increasingly bringing in chief risk officers from the banking business who are determined to break down the legacy Chinese walls between the actuaries and the rest of the organization. We expect that the traditional “walls” down the middle of insurance companies and pension funds will break down rapidly in the years ahead. That, in the end, is good for all concerned.

We now step back and do exactly that, ourselves. We look at the entire integrated risk of the entire organization in the next six chapters, now that we have completed our transaction-level survey of best-practice valuation techniques.

NOTES

1. We encourage readers to explore the excellent paper on this topic by Samir Soneji of Dartmouth and Gary King of Harvard, "Statistical Security for Social Security," Harvard University Working Paper, June 30, 2011.
2. A recent study, Stanford Institute for Economic Policy Research, "Going for Broke: Reforming California's Public Employee Pension Systems," April 2010, made this particularly clear with respect to the pension funds overseen by the State of California.

PART

Five

Portfolio Strategy and Risk Management

Value-at-Risk and Risk Management Objectives Revisited at the Portfolio and Company Level

In Chapter 1, we defined risk management in a practical way:

Risk management is the discipline that clearly shows management the risks and returns of every major strategic decision at both the institutional level and the transaction level. Moreover, the risk management discipline shows how to change strategy in order to bring the risk return trade-off into line with the best long- and short-term interests of the institution.

In Chapter 1, we noted that this definition of risk management includes within it the overlapping and inseparable subdisciplines, such as:

- Credit risk
- Market risk
- Asset and liability management
- Liquidity risk
- Capital allocation
- Regulatory capital calculations
- Operational risk
- Performance measurement
- Transfer pricing

Our primary focus in this book is to show how to execute the practice of risk management in a way that is fully integrated and makes no distinction between these subdisciplines. We do that in this chapter.

THE JARROW-MERTON PUT OPTION AS A MEASURE OF TOTAL RISK: AN EXAMPLE

Value-at-risk (VaR) as a risk measure has dominated the past 20 years of risk analytics, especially for traded instruments. The authors grant that this focus was

well-intended. Alas, as the large JPMorgan Chase credit trading losses of 2012 and the listing of large institution losses in the introduction to this book make clear, the VaR concept has not worked well in preventing losses and failures at the institutions that have relied on it. In the rest of this chapter, we draw a clear distinction between VaR and best practice, state-of-the-art risk management.

Unlike VaR, the Jarrow-Merton put option as a measure of total risk is an observable, tradable number. Exhibit 36.1 shows put option prices on the common stock of Citigroup prevailing on July 25, 2012.

This data is available, often in real time, on websites accessible even to retail investors, let alone CEOs of major financial institutions. The chart makes a number of things much more clear than the kind of value-at-risk calculations outlined by Jorion's (2006) classic overview of the subject:

- All assets and liabilities contribute to total risk, not just traded assets.
- Total risk is what matters, not just a narrow definition of market risk.
- Risk has a term structure and the price of protection against total risk changes as the time horizon expands beyond the myopic 10- to 30-day focus of most VaR analysis.

The Jarrow-Merton put option allows a risk analyst to definitively answer practice questions on capital adequacy. Exhibit 36.2, the data behind the graph in Exhibit 36.1, gives us the ability to deal with a number of questions very specifically.




EXHIBIT 36.1 Citigroup Put Option Prices as Percent of Stock Price (\$25.7582) and Days to Maturity, July 25, 2012

EXHIBIT 36.2 Citigroup Put Option (Last) Price as Percent of Stock Price, July 25, 2012

Strike Price	Days to Maturity of Option Contract									
	In-the-Money Amount									
	0	1	23	58	86	114	140	177	233	541
13	0.000%		0.078%	0.194%		0.427%	0.893%	0.971%		
14	0.000%		0.078%	0.194%	0.660%	0.776%	0.893%	2.096%	1.708%	
15	0.000%		0.039%	0.621%	0.505%	0.699%	0.854%	1.475%	1.980%	5.358%
16	0.000%		0.078%	0.272%	0.776%	0.971%	1.359%	1.863%	1.631%	
17	0.000%		0.078%	0.388%	0.893%	1.281%	1.320%	2.291%	3.183%	
18	0.000%		0.078%	0.466%	0.621%	1.398%	1.902%	2.485%	3.572%	8.619%
19	0.000%		0.116%	0.660%	1.320%	1.592%	2.485%	3.339%	4.309%	
20	0.000%		0.155%	0.971%	2.058%	2.679%	3.222%	3.999%	5.047%	10.948%
21	0.000%	0.039%	0.233%	1.126%	2.213%	3.416%	3.999%	4.814%	6.250%	
22	0.000%	0.039%	0.427%	1.747%	2.912%	3.999%	5.590%	5.862%	7.532%	
23	0.000%	0.039%	0.776%	2.446%	3.999%	4.620%	6.289%	7.105%	8.153%	15.723%
24	0.000%	0.078%	1.398%	3.339%	5.047%	6.212%	8.386%	9.434%	11.103%	
25	0.000%	0.349%	2.407%	4.620%	6.095%	7.454%	8.813%	9.977%	12.811%	19.023%
26	0.939%	1.747%	4.038%	5.823%	9.473%	10.870%	12.423%	11.530%	12.617%	
27	4.821%		6.483%	8.696%	10.094%	11.491%	12.617%	13.976%	16.111%	24.458%
28	8.703%	8.852%	9.589%	11.336%	14.753%	15.141%	17.082%	16.111%	15.723%	
29	12.586%		13.006%	13.006%	14.170%	17.470%	19.800%	17.664%	18.053%	
30	16.468%	17.858%		17.470%	17.470%	18.829%	22.129%	23.488%	23.876%	31.835%
31	20.350%			22.517%		21.935%	21.935%	22.517%	27.331%	
32	24.232%		26.982%	25.041%		25.235%	26.205%		30.631%	37.076%
33	28.115%		28.340%	31.058%	30.476%		32.417%	32.223%	30.670%	
34	31.997%		34.746%		34.940%	35.329%	35.717%			

(Continued)

EXHIBIT 36.2 (Continued)

Days to Maturity of Option Contract										
Strike Price	In-the-Money Amount									
	0	1	23	58	86	114	140	177	233	541
35	35.879%						38.628%	36.843%		44.258%
36	39.761%		40.181%	42.899%				43.870%		
37	43.644%			45.811%				46.393%	47.596%	50.314%
38	47.526%			50.314%						
39	51.408%			52.061%						
40	55.290%			55.710%				58.428%		61.806%
41	59.173%			59.787%		59.593%				
42	63.055%									
43	66.937%									
44	70.819%									
45	74.702%							77.140%		79.858%

Sources: Yahoo! Finance; Kamakura Corporation.

If dynamic new Citigroup Chairman of the Board Michael O'Neill asked this question: "How much additional capital does Citigroup need to ensure that our stock price is at least \$25 in January 2013?" Exhibit 36.2 contains the answer. On July 25, 2012, Citigroup's stock was trading at \$25.7582. A put option with a strike price of \$25 and a maturity of January 18, 2013 (177 days from July 25, 2012), costs \$2.57 per share or 9.977 percent of the current Citigroup stock price. What trade would Citigroup need to do to buy this insurance? We note that this trade cannot be done in practice for legal reasons but it does give us the market price of total risk protection. The hypothetical trade is this, assuming no change in Citigroup stock price if a new equity issue is done:

1. Issue N new shares at \$25.7582 per share, bringing the total number of shares from the current number of X to $X + N$
2. Buy put options with a total cost of \$2.57 per share (9.977 percent of share price) on $X + N$ shares for a total cost of \$2.57 ($X + N$)
3. On January 18, 2013, if Citigroup's share price is below \$25, exercise the puts

With the current number of shares outstanding at Citigroup at about 2.93 billion shares, puts on those shares would be \$2.57 (2.93 billion) = \$7.5 billion. We then solve for the number of new shares N such that the amount of the new issue is \$7.5 billion plus \$2.57 N (since we need puts on the new shares too).

$$25.7582N = 2.57(2.93\text{billion}) + 2.57N$$

The number of shares that need to be issued to protect a stock price level of \$25 per share on January 18, 2013, is $N = 325$ million. If Citigroup's total risk were lower, N would be a lot lower number and the put price of \$2.57 would be much lower.

Value-at-risk provides next to no information on this practical question.

A FOUR-QUESTION PASS-FAIL TEST FOR FINANCIAL INSTITUTIONS' CEOs AND BOARDS OF DIRECTORS

Above and beyond the practical question posed above, the four most important risk management questions that a financial institution's CEO and Board should be able to answer are easy to summarize in a pass-fail test.

Question 1: What happens to the market capitalization and net income of the firm if any of these risk factors change: home prices, foreign exchange rates, commercial real estate prices, stock index levels, interest rates, commodity prices?

If the answer is "I don't know," the institution fails the test.

Question 2: Using an insider's knowledge of the assets and liabilities of the firm, both "on balance sheet" and "off balance sheet," what is the best estimate, monthly for the next 10 years, of the probability that the firm will fail in each of these 120 monthly periods?

If the answer is “I don’t know,” the institution fails the test.

Question 3: Using only information available to an outsider, what is the best estimate of the probability of the failure of the firm in both the short run and the long run?

If the answer is “I don’t know,” the institution fails the test.

Question 4: If the firm is able to answer Questions 1, 2, and 3, what hedging position is necessary to insure that the macro factor sensitivity of the firm and default probability of the firm reach the target levels set by the Board of Directors?

If the answer is “I don’t know,” the institution fails the test.

Why Do These Four Questions Matter?

When one looks at the ex post analysis of why institutions like Merrill Lynch, UBS, and Citigroup got into such trouble during the current crisis, the feedback from the Board and CEO levels is amazingly consistent. We summarized key quotations in this regard in Chapter 1. In short, management admitted that they could not answer these questions and that is why the institutions failed. The purpose of this book is to provide a framework that allows an automated fully integrated highly accurate answer to these questions that is consistent with the observable Jarrow-Merton put prices outlined previously.

But aren’t there more than four questions one should be able to answer? Here are some other questions that financial executives using best practice risk management can answer and should answer.

An Alphabet of 26 Extra-Credit Questions

A modern approach to risk management can provide answers to a longer list of additional questions than these four key questions posed above. To illustrate the power of a multiperiod credit adjusted approach to risk management, we pose additional “extra credit” questions that a financial institution can answer with a high degree of accuracy using a best practice approach.

- A. What is the probability that the firm will have to increase its reserve for loan losses by more than X per quarter, quarterly for the next 10 years?
- B. Given that macroeconomic factors drive actual losses, what pattern for the allowance for loan losses gives the most stable net income for the institution?
- C. To what extent does lending in commercial real estate provide diversification versus the bank’s mortgage lending businesses?
- D. What interest rate risk posture provides the most effective offset of credit losses given the links between interest rates and defaults?
- E. Given cyclical credit losses, what timing of common stock offerings (1) meets capital needs through the cycle and (2) results in the highest average common stock price on new issues?
- F. How do regulatory capital ratios like the Basel II risk measures vary over the business cycle and to what extent are they correlated with the firm’s actual risk of default?

- G. What 10 macro factors represent the highest threat to the safety and soundness of the institution, from highest to lowest?
- H. What risk limits should the firm have in place with respect to these 10 macro factors?
- I. What hedging instruments are available to mitigate the risk of these macro factors if their impact on the institution's default risk grows too high?
- J. How should these macro factor risk exposures be displayed and reported in the institution's annual report and regulatory filings like the 10-k report of the U.S. Securities and Exchange Commission?
- K. How should compensation systems for senior management be modified in light of the risk that the institution has from its macro factor exposure?
- L. How should the institution replace external and internal ratings, which obscure macro factor links to default risk, with a more transparent "credit risk CAT scan" of risk?
- M. Looking backward in time, which macro factors had the largest historical impact on the firm's market capitalization and default risk?
- N. How do macro factor movements affect the supply of liabilities (alleged "core" deposits, certificates of deposit, commercial paper, bonds, derivatives) to the firm?
- O. What percentage decline can be expected in the supply of liabilities to the firm if there is an adverse movement of 10 percent, one by one, in the 10 most important macro factors affecting the institution?
- P. What strategies can be put in place to mitigate this macro factor-driven credit risk and liquidity risk?
- Q. Which class of business counterparties has the greatest sensitivity to each of these macro factors, and how should pricing of credits to these counterparties reflect the risk?
- R. Is the institution as diversified as it could be with respect to this macro factor risk? If not, how could diversification be improved?
- S. How is operational risk linked to movements in the same top 10 macro factors?
- T. To what extent do legacy silo risk systems and ratings systems obscure the answers to the questions above?
- U. Can the system's infrastructure be improved so that an integrated risk assessment of macro factor sensitivity is both more accurate and more efficiently produced?
- V. Can the management team and board be strengthened so there is an enhanced understanding of macro factor-driven risk across the organization?
- W. Has the firm worked closely with auditors to ensure that the auditors have a clear understanding of macro factor-driven risk?
- X. Is the firm's self-assessment of its own risk level consistent with the pricing on put options on its common stock and credit default swaps?
- Y. If the answer is no, what is the reason for the inconsistencies?
- Z. If the market prices of puts and CDS are overly pessimistic, what investor education campaign is necessary to more effectively communicate the true risk of the institution?

A firm whose management team can answer this alphabet of questions truly represents best practice in risk management and shareholder value creation.

IS YOUR VALUE-AT-RISK FROM VALUE-AT-RISK?

How does a firm's ability to answer these critical 4 + 26 risk management questions differ if value-at-risk is the primary risk measure? We summarize the well-known problems of value-at-risk in this section.

VaR has as many incarnations as there are risk managers. In the mid-1990s, JPMorgan made the decision to publicize the technique and to spin off an affiliate that would resell the VaR technology and related information. Hundreds of financial institutions adopted the VaR technology as a supplement for or replacement of traditional trading floor-based risk limits like PVBP (present value of a basis point move in interest rates) or delta-type measures with respect to other risk factors, like stock prices and foreign exchange rates. Philippe Jorion (2006) has done a masterful job in his book at documenting various ways to implement the VaR concept. The purpose of this section is much simpler—to answer the question “Does it work?” Could a sophisticated user of the VaR concept have anticipated the 2006–2011 credit crisis and taken action early on to mitigate the impact of the crisis? The answer is . . . maybe. That's not a sufficiently strong endorsement to justify the considerable sums that many institutions have invested in legacy VaR technology. Why is the answer so lukewarm? We turn now to the reasons.

First of all, there are three principal variations of VaR technology, and they have different strengths and weaknesses:

1. *Historical VaR*, which estimates the n th percentile in potential losses based on historical returns on the relevant traded assets.
2. *Variance-covariance VaR*, which takes the historical volatility and correlation of returns on traded assets to calculate the n th percentile in potential losses
3. *Monte Carlo VaR*, which simulates the environment forward, rather than using historical data like A and B.

Most of the initial publicity around the VaR concept focused on a 10-day time horizon for the calculation, since it was advertised as an index of short-term trading risk for assets traded in an active market. Most implementations assume that the portfolio of assets and liabilities that one is analyzing stays constant over the one-period horizon (with length 10 days).

How well did these techniques serve their users in the credit crisis? On January 28, 2008, a story on Bloomberg.com, “Death of VaR Evoked as Risk-Taking Vim Meets Taleb's Black Swan,” reported that Merrill Lynch's historical VaR measure was \$92 million at a time that CDO losses had already reached \$18 billion at Merrill. With all due respect to Nassim Taleb, this was no “black swan.” The error in the VaR measure at Merrill was due to simplifying assumptions in the calculation that were simply not true, leading to incorrect risk assessment. With respect to historical VaR, here are the major problems that would lead one to underestimate risk by a factor of 200 like Merrill did:

- Problem 1: Future returns almost never equal historical returns. Home prices dropped, mortgage default rates skyrocketed, and the historical returns (which are typically measured over a period of less than a year) did not adequately reflect the possibility of these events.

Problem 2: Much of the losses in the 2006-2011 crisis have stemmed from assets that either always have been or are currently “nontraded:” subprime mortgage loans, Alt-A loans, and structured products using them as reference collateral. They were simply left out of the calculation

Problem 3: The balance sheet you end with is not the balance sheet you start with, even though that’s the assumption made by most VaR users. A veteran risk manager at one of the largest life insurance firms in North America once asked, “This VaR concept doesn’t make any sense. If we extend the time horizon to a year, a huge percentage of our assets would have matured by then. How could anyone assume a constant balance sheet for VaR?”

Most of these problems apply to variance-covariance VaR as well, with an additional monkey wrench thrown into the gears:

Problem 4: Returns on most securities are *not* normally distributed. If one assumed that the returns on Bear Stearns common stock were normally distributed at their historical mean and variance, one would have implied a probability of failure (a monthly return of –100 percent) of 0.000000. For CDO tranches (see Chapter 20) with narrow bands of losses, the outcome is almost exclusively “all or nothing,” since it becomes highly unlikely for the tranche to suffer partial losses when the size of the tranche is a small percentage of total notional principal of the collateral underlying the CDO. The net impact of the normality assumption is a dramatic underestimation of risk.

When one turns to Monte Carlo VaR, still another problem can be important:

Problem 5: Returns in period N are not independent of the returns in period $N - 1$, even though this assumption of independence is very, very common in financial theory and practice. Take the Case-Shiller home price index from www.sandp.com and test the hypothesis that home price returns are independent of returns on prior periods to see for yourself—the returns are very, very highly cyclical and dependent on past returns.

Can VaR be saved? The answer is yes if the analyst is very careful and has a powerful enterprise-wide software package available. Here are the steps that one needs to take to get maximum accuracy from the VaR way of thinking:

Fix Number 1: *Do the VaR calculation on a multiperiod basis* that allows for a dynamic balance sheet, with cash flow reinvested, new assets and liabilities generated, and embedded options exercised.

Fix Number 2: *Allow for defaults to occur*, since they are the biggest source of “non-normality” in the real world. Kamakura Risk Manager uses default probabilities from the Kamakura Risk Information Services default service (or any other source) for this purpose.

Fix Number 3: *Recognize that the drivers of risk, typically macroeconomic factors*, can be cyclical with probability distributions in period N that depend on prior periods.

Fix Number 4: *Look at the entire distribution of value*, not just one percentile level, at multiple points in time.

Fix Number 5: *Stress test VaR and mark-to-market value with respect to all relevant macro factors* (like home prices) to avoid being lulled into complacency. This requires explicit links between the macro factors, default probabilities, and credit spreads.

With these fixes, a multiperiod default-adjusted VaR calculation can add some insight that creates incremental risk-adjusted shareholder value added. We now turn to some additional important differences between VaR and the Jarrow-Merton put option concept.

VaR VS. THE PUT OPTION FOR CAPITAL ALLOCATION

What are the real differences between using the VaR approach to capital allocation and using the put option approach? This section answers that question.

A worked example shows the differences very clearly. Let's assume that there are only 100 scenarios that can impact our firm ABC Company. In best practice net present value terms, in 93 scenarios there is no value destruction and capital is unchanged. In seven scenarios, numbered 94 to 100, we have these losses:

Scenario	Net Present Value of Losses
94	1
95	3
96	4
97	6
98	12
99	17
100	27

We can then easily determine that the average loss over all 100 scenarios is 0.70. The ninety-ninth percentile VaR is 17, and the ninety-fifth percentile VaR is 3. We need only \$3 of capital to survive 95 percent of the time, but we need \$17 to survive 99 percent of the time. This is simple enough. How does the put option approach compare?

Let's make it very clear that, when we are talking about a put option, we could equivalently call it an insurance policy that reimburses us for any value destruction in our current portfolio. Puts come at various strike prices and insurance policies come with deductibles and caps. Exhibit 36.3 shows the payments that an insurance company would make to us in each of the loss scenarios for four policies: three policies where the losses are capped (at 5, 10, and 20) and one policy with no caps.

How much would each of these put options or insurance policies cost? Since our losses are stated in net present value terms, in a competitive market with perfect competition the put/policy should be priced at its expected value over the 100 scenarios. For the policy capped at a maximum loss of 5, the cost is 0.28 (the

EXHIBIT 36.3 Loss Scenarios for Four Policies

Scenario	Net Present Value of Losses	Losses Capped at 5	Losses Capped at 10	Losses Capped at 20	No Cap on Losses
94	1	1	1	1	1
95	3	3	3	3	3
96	4	4	4	4	4
97	6	5	6	6	6
98	12	5	10	12	12
99	17	5	10	17	17
100	27	5	10	20	27

expected losses of $1 + 3 + 4 + 5 + 5 + 5 + 5 = 28/100$ trails). For the \$10 cap, the cost is 0.44. For the 20 loss cap, the cost is 0.63. For a policy that reimburses us for all losses, the cost of the put is exactly the expected loss on the portfolio: 0.70. How much capital would be needed if we were comfortable with insurance only up to 20 in losses? We issue additional capital of 0.63, and we use the 0.63 to buy the insurance policy.

WHY ARE THE VaR AND PUT APPROACHES SO DIFFERENT: SELF-INSURANCE VS. THIRD-PARTY INSURANCE

How do we explain the fact that the ninety-ninth percentile VaR capital needed is 17 and the put option with a cap at 20, where the company also loses money in only 1 of the 100 scenarios, costs only 0.63, not 17? The VaR capital calculation assumes that you are “self-insuring” the company against losses. Even though 17 or more in losses happens in only 2 of the 100 scenarios, you issue 17 in capital, which turns out to be unnecessary 98 times out of 100.

The Jarrow-Merton put premium or insurance cost of protecting your portfolio of losses up to \$20 reflects not only the amount of the losses, but also their probability of occurrence. The insurance provider can diversify, and a competitive marketplace will ensure that the benefits of diversification are passed along to the clients of the insurance provider. In a nutshell, the VaR approach assumes you are self-insured and can’t diversify. The put option approach assumes you buy protection from a third party who is diversified. It’s all very simple.

As we noted previously, when it comes to capital allocation using the put option concept, we don’t even have to make the calculation for bank holding companies that are listed on stock exchanges in the United States. Put options on their common stock will be publically observable, and, for any “deductible” (strike price). There will be a different amount of capital required for that level of protection as we saw in Exhibits 36.1 and 36.2 for Citigroup.

Another approach to protect one’s capital has tremendous moral hazard for your counterparty—you can buy credit default swaps on yourself. Rumor has it that AIG was “sophisticated” enough to be willing to do this. But what if Nassim Taleb is right and the CDS market is the equivalent of buying marine insurance from another

passenger on the *Titanic*? The put/insurance policy can still be valued and purchased. We do a credit risk “CAT scan” that looks through all the assets and liabilities of the firm and all of the counterparties to see how they are affected by key macro factors like interest rates, home prices, foreign exchange rates, and stock indices. This is exactly what we describe below. We then buy the appropriate insurance protection on these macro factors. Almost all of them have exchange traded futures and (in some cases) options, so we don’t have a *Titanic* problem.

In short, the difference between a VaR approach to capital needs and a put option or insurance approach to capital needs is the difference between self-insurance or insurance in an efficient market that recognizes diversification. In most cases, the latter assumption will be a much more accurate indicator of how much capital protection is necessary.

As we explained previously, a put option or insurance policy that insures the firm against the same loss profile is the probability-weighted present value of all scenarios. Given identical loss projections, this normally means that the “true” VaR will be much higher than the put price (unless the losses are the same in all scenarios) because the put is probability weighted and the VaR is not. Then what about this quote from Bloomberg.com on January 28, 2008?

Merrill’s highest one-day value-at-risk in the third quarter was \$92 million, indicating that the firm’s maximum expected cost during the 63-trading day period would be \$5.8 billion. In fact, the firm wrote down \$8.4 billion from the value of collateralized debt obligations, subprime mortgages and leveraged finance commitments, 45 percent more than the worst-case scenario. (www.bloomberg.com/apps/news?pid=newsarchive&sid=axo1oswvqx4s.)

How could that happen?

What was reported in the Bloomberg story was “false” VaR, a gross underestimate of risk that has a host of problems that we listed in the previous section “Is Your Value-at-Risk from Value-at-Risk?” The Bloomberg story goes on to state that most of the large U.S. securities firms were using “false” VaR based on one to four years of historical data. Given the history of U.S. home prices, a four-year historical false VaR would have predicted minimal losses (as it did) because over the four-year period ending December 31, 2007, home prices were strongly up. Even though home prices peaked in Los Angeles in September 2006, they did not begin to have a powerful impact on mortgage defaults until the second half of 2007. False VaR is based on an average of history. True VaR, which is what we assumed in our previous example, is a complete and accurate listing of all possible future outcomes and their probabilities, Nobel laureate Kenneth Arrow’s famous “state space.” If we base true VaR on that complete listing of outcomes and probabilities, the put price or cost of insuring risks via a third party that can diversify will almost always be less than the high percentile (say ninety-ninth) true VaR number. That true VaR number is the amount of self-insurance you need to survive with 99 percent probability.

Wasn’t the miss with respect to VaR just a black swan as Nassim Taleb argued? No, it was just bad math. It was an assumption that the world was flat. As we explained previously, historical VaR contains an implicit assumption that neither Lehman nor Bear Stearns could fail, because (based on historical equity returns) the

implied probability of -100 percent stock returns within a month for each of them was 0.000000 percent. Similarly, if the monthly changes in home prices have been strongly positive over the measurement period, the historical VaR approach will implicitly assume away the probability of a decline. The “true” VaR calculation and a rationally priced put will consider all possible future scenarios, not just those that have actually come about in the past 12 or 48 months.

CALCULATING THE JARROW-MERTON PUT OPTION VALUE AND ANSWERING THE KEY 4 + 26 QUESTIONS

“What is the hedge?” In other chapters in this book, we posed this question as the best single-sentence test of risk management technology. For that reason, the Jarrow-Merton put option is a very powerful concept because the put option they propose, if properly structured, *is* the hedge. What is the equivalent of the Merton and Jarrow put option in the interest rate risk context? It is the value of an option to buy the entire portfolio of the financial institution’s assets and liabilities at a fixed price at a specific point in time.

Here are some examples of the Jarrow-Merton put option as a practical risk management concept:

- *Instead of the 10-day value-at-risk* of a trading portfolio, what is the value of a 10-day put option on my current portfolio with an exercise price equal to the portfolio’s current market value? The price of the put option will increase sharply with the risk of my portfolio, and the put option’s price will reflect all possible losses and their probability, not just the ninety-ninth percentile loss as is traditional in value-at-risk analysis
- *Instead of stress testing the 12-month net income* of the financial institution to see if net income will go below \$100 million for the year, what is the price of a put option in month 12, which will produce a gain in net income exactly equal to the shortfall of net income versus the \$100 million target? The more interest rate risk in the balance sheet of the financial institution, the more expensive this put option will be. The put option will reflect all levels of net income shortfall and their probability, not just the shortfalls detected by specific stress tests
- *Instead of the Basel II, Basel III, or Solvency II risk-weighted capital ratio* for the institution, what is the price of the put option that insures solvency of the institution in one year’s time? This put option measures all potential losses embedded in the financial institution’s balance sheet and their probability of occurrence, including both interest rate risk and credit risk, as we discuss at the end of this chapter.
- *Instead of expected losses on a collateralized debt obligation tranche’s B tranche*, what is the price of a put option on the value of the tranche at par at maturity? This put option reflects all losses on the tranche, not just the average loss, along with their probability of occurrence.
- *Instead of expected losses on the Bank Insurance Fund* in the United States, the Federal Deposit Insurance Corporation (FDIC) has valued the put option of retail bank deposits at their par value as discussed in the FDIC’s loss distribution model announced on December 10, 2003.

The Jarrow-Merton put option concept helps to reconcile what many regard as conflicting objectives to be managed from a risk management perspective:

- *Net interest income (or net income)*, which is a multiperiod financial accounting-based exhibit that includes the influence of both instruments that the financial institution owns today and those that it will own in the future.
- *Market value of portfolio equity*, which is the market value of the assets a financial institution owns today less the market value of its liabilities.
- *Market-based equity ratio*, which is the ratio of the mark-to-market value of the equity of the portfolio (“market value of portfolio equity” in bank jargon) divided by the market value of assets. This is most closely related to the capital ratio formulas of the primary capital era, Basel I, Basel II, Basel III, and Solvency II.
- *Default probability of the institution*, which is another strong candidate as a single measure of risk.

We next discuss how the valuation of these options can be done using the technology in the first 38 chapters of this book.

VALUING AND SIMULATING THE JARROW-MERTON PUT OPTION

How do we value the Jarrow-Merton put option? In some of the cases mentioned above, we already have known valuation approaches from Chapters 20 through 35 using an N -factor Heath, Jarrow, and Morton (HJM) interest rate model. In those chapters, we used the three-factor HJM model from Chapter 9 for exposition purposes. For the full on- and off-balance sheet portfolio of a financial institution, we simply repeat the process with finer time granularity and more transactions. This is what computers are for. We need to calculate the risk-neutral expected values of the cash flows that would be paid on the Jarrow-Merton put option that we are analyzing.

We outlined the steps in this simulation in Chapter 19 in the context of a portfolio of risky bonds:

For the reduced form model, the steps are as follows:

1. Simulate the risk-free term structure using an N -factor HJM approach.
2. Choose a formula and the risk drivers for the liquidity component of credit spread for all of the relevant asset classes as discussed in Chapter 17.
3. Choose a formula and the risk drivers for the default-intensity process in the Jarrow model as discussed in Chapter 16 for all asset classes.
4. Simulate the random values of the drivers of liquidity risk and the default intensity for M time periods over N scenarios, consistent with the risk-free term structure.
5. Calculate the default intensity and the liquidity component at each of the M time steps and N scenarios for each counterparty, from retail to small business to corporate to sovereign. Note that these will be changing randomly over time because interest rates and other key macro factors (like home prices) are changing as well. This is essential to capturing cyclicalities in traded asset prices and defaults.

6. Apply the Jarrow model as we have done in Chapters 16 and 17 to get the zero-coupon bond prices for each maturity for each transaction for each counterparty.
7. Calculate the dates and cash flow amounts for the Jarrow-Merton put option structure that is most relevant.
8. Calculate the put option's value by discounting by the appropriate risk-free zero-coupon bond price for each of the M time steps and N scenarios.

This is a general valuation procedure that we can apply for all instruments discussed in Chapters 20 through 26.

WHAT'S THE HEDGE?

Let's say management has become comfortable with the Jarrow-Merton put option as a comprehensive measure of risk. How do we address the key test of risk management technology: "What's the hedge?"

If we have done a comprehensive job of fitting our default probability models as in Chapters 16–17 and a thorough job of fitting our interest rate models as discussed in Chapters 5 through 10, we can stress test the put option value that we derived in the previous section. We pick a macroeconomic risk factor to stress test. If the financial institution is a lender in the United States, the S&P 500 is a commonly used macro factor that can be proven to be a statistically significant driver of default probabilities for a large range of counterparties. We can calculate what position in S&P 500 futures is needed to control changes in the Jarrow-Merton put option as a measure of risk in exactly the same way van Deventer and Imai (2003) discuss macro hedging of a portfolio of risky credits:

1. Select the base case value X for the S&P 500 and all other macro factors.
2. Value the appropriate Jarrow-Merton put option using the approach of the previous section.
3. Select the stress test value Y for the S&P 500 and use all of the same macro factor values as in step 1. The same Monte Carlo "seed value" as step 1 also has to be used.
4. Get the stress-tested value of the Jarrow-Merton put option.
5. The change in the Jarrow-Merton put option is the Value in Case 1 – Value in Case 2.
6. The change in the S&P 500 is $X - Y$.
7. The proper hedge depends on what type of S&P 500 hedging instruments are being used (futures contracts or not, maturity of futures contracts, etc.). The number of contracts to be used for the hedge is the number such that the change in the value of the futures contracts if the S&P 500 moves from X to Y will exactly offset the change in the value of the Jarrow-Merton put option from its Case 1 value to its Case 2 value.

Using this approach, we always know the proper integrated hedge of interest rate risk and credit risk. Our objectives have been achieved.

**LIQUIDITY, PERFORMANCE, CAPITAL ALLOCATION, AND
OWN DEFAULT RISK**

Using the Jarrow-Merton put option approach that we have built in this chapter on the foundation of Chapters 1 through 35, we can now definitively address key risk management issues like those in our previous 4 + 26 questions:

What is the liquidity risk of my organization?

What amount of capital should I have in the institution as a whole?

What amount of capital should I have in each business unit?

What is the performance of each business unit?

We turn to that task in the next three chapters.

NOTE

1. www.bloomberg.com/apps/news?pid=newsarchive&sid=axo1oswvqx4s.

Liquidity Analysis and Management

Examples from the Credit Crisis

In Chapter 36, we applied the Jarrow-Merton put option concept as a comprehensive measure of integrated interest rate, market risk, foreign exchange risk, and credit risk. We listed four key questions and 26 supplementary questions that can be answered as a result of the risk management process that are outlined in the first 36 chapters of this book. We reviewed the observable market data on the cost of the Jarrow-Merton put option for Citigroup at various time horizons. Finally, we showed the multiperiod simulation process that allows us to value the Jarrow-Merton put option, giving us an alternative measure of the dollar amount of money that would be necessary to eliminate the risk we face. We also noted in that chapter that the put option concept can be applied to risk management defined as shortfalls in net income, capital ratios, or provisions for loan losses.

In this chapter, we apply the same concept to shortfalls in cash—liquidity risk—which is tightly linked with the credit risk of the institution. We take great care to avoid the single greatest mistake in liquidity risk analysis—basing the analysis solely on the cash flow history of a firm that has never had a “near death” experience from liquidity risk. Indeed, our focus in this chapter is primarily on institutions that *have* had such near death experiences. We start with a review of the five biggest liquidity risk problems in North America during the 2006–2011 credit crisis.

LIQUIDITY RISK CASE STUDIES FROM THE CREDIT CRISIS

In the introduction to this book, we discussed the six biggest fallacies in risk management. As an introduction to five case studies of liquidity shortfalls, we survey each of the risk fallacies that contributed to the 2006–2011 credit crisis, which, in turn, caused the liquidity shortfalls that are outlined in this section:

- *If it hasn't happened to me yet, it won't happen to me, even if it has happened to someone else.* Many market participants believed that home prices in the United States would never go down, but they did.
- *Silo risk management allows my firm to choose the “best of breed” risk model for our silo.* Risk management was so fragmented at the largest U.S. and European institutions that the pervasive impact of home price falls on a wide range of counterparty risks and security values was not modeled realistically.

- *I don't care what's wrong with the model. Everyone else is using it.* The copula approach to CDO valuation was hopelessly flawed and confined to a largely desktop review of traded products. A macroeconomic factor-driven total balance sheet analysis of credit risk was rarely performed.
- *I don't care what's wrong with the assumptions. Everyone else is using them.* The assumption that home prices and other macro factors had returns that were independent from period to period was dramatically wrong. There is a strong persistence in trends in returns, causing macro factor swings to be much larger than what was modeled.
- *Mathematical models are superior to computer simulations.* Few banks had the systems capability to realistically model the impact of a large number of random macro factors. What analysis was done was dominated by closed form solutions that could only be obtained by making unrealistic assumptions (like the independence assumption above) about the nature of macro factor movements.
- *Big North American and European banks are more sophisticated than other banks around the world and we want to manage risk like they do.* The evidence in this section is proof positive that, sadly, their example is not one that should be followed.
- *Goldman says they do it this way and that must be right.* Michael Lewis' (2011) instant classic, *The Big Short: Inside the Doomsday Machine*, thoroughly documents how a few securities dealers were able to whipsaw smaller market participants with contrarian (and more accurate) views for an extended period of time with arbitrary mark-to-market positions imposed by those dealers. This aggravated the crisis, delayed recognition of its magnitude, and triggered some of the biggest flows of bailout funds.¹

We now turn to five case studies in liquidity risk during the 2006–2011 credit crisis.

CASE STUDIES IN LIQUIDITY RISK

Largest Funding Shortfalls

In this section, we review the data used in the Kamakura analysis of 21 cases studied in liquidity risk and then review each of the top five funding shortfalls.² Under the Dodd-Frank Act of 2010, the Board of Governors of the Federal Reserve was required to disclose the identities and relevant amounts for borrowers under various credit facilities during the 2006–2011 financial crisis. These credit facilities provide, perhaps, the best source of data about liquidity risk and funding shortfalls of the last century. We use this data to determine to what extent there was a funding shortfall at the largest institutions active in U.S. financial markets during the credit crisis.

The data used in the study consist of every transaction reported by the Federal Reserve as constituting a “primary, secondary, or other extension of credit.” Included in this definition are normal borrowings from the Fed, the primary dealer credit facility, and the asset-backed commercial paper program. Capital injections under the Troubled Asset Relief Program (TARP) and purchases of commercial paper under the Commercial Paper Funding Facility are not included in this definition put forth by the Federal Reserve. We analyze borrowings under the Commercial Paper Funding Facility separately.³

Kamakura took the following steps to consolidate the primary, secondary, and other extensions of credit:

- From www.twitter.com/zerohedge, Kamakura downloaded the daily reports, in PDF format, from the Federal Reserve on primary, secondary, and other extensions of credit from February 8, 2008, until March 16, 2009, approximately 250 reports in total.
- Kamakura converted each report to spreadsheet form.
- These spreadsheets were aggregated into a single database giving the origination date of the borrowing, the name of the borrower, the Federal Reserve District of the borrower, the nature of the borrowing (ABCP, PDCF, or normal), the maturity date of the borrowing, and (in the case of Primary Dealer Credit Facility) the name of the institution holding the collateral.
- Consistency in naming conventions was imposed; that is, while the Fed listed two firms as “Morgan Stanley” and “MS Co,” Kamakura recognized to the maximum extent possible that they are the same institution and used a consistent name.
- To the maximum extent possible, the name of the ultimate parent was used in order to best understand the consolidated extension of credit by the Fed to that firm.⁴

Kamakura calculated the daily borrowings of a large number of prominent international financial institutions during the period covered by the Federal Reserve data. For each of these selected institutions, Kamakura calculated these statistics:

- Maximum Outstanding Borrowing
- Average Outstanding Borrowing
- First Borrowing Date
- Average Borrowing on Nonzero Borrowing Days
- Number of Borrowing Days

Exhibit 37.1 ranks these selected institutions by maximum borrowings during the period covered. AIG leads the list with a maximum borrowing of \$208 billion during this period.

Readers should be under no illusion that this was just a temporary facility and that these institutions could have survived without government aid. Nothing could be further from the truth, as nicely summarized by the report by the Office of the Special Inspector General of the Troubled Asset Relief Program entitled “Emergency Capital Injections Provided to Support the Viability of Bank of America, Other Major Banks, and the U.S. Financial System,” October 5, 2009. For each of the top five borrowers in Exhibit 37.1, we discuss their borrowings in detail below. Exhibit 37.2 ranks the same institutions by the average amount of borrowing on those days during the period when borrowings were outstanding.

We now turn to a name-by-name summary of the top five liquidity crises in North America during the period covered by the Fed data.

American International Group (AIG)

Excellence in risk management and good corporate governance require that financial institutions analyze their own probability of default. The proposed Basel III

EXHIBIT 37.1 Kamakura Corporation, Analysis of Primary, Secondary, and Other Borrowings from the Federal Reserve Ranked by Maximum Borrowing

Rank	Institution	Maximum Outstanding	Average Outstanding	First Borrowing Date	Average Of Non-Zero Borrowing Days	Number of Borrowing Days
1	AIG	208,616,483,142	72,793,773,121	20080916	161,186,211,912	182
2	Consolidated JPMorgan, Bear Stearns, and WaMu	101,125,000,000	17,728,196,729	20080306	23,666,720,690	290
3	State Street	77,802,450,046	9,000,920,496	20080220	19,821,699,234	183
4	Morgan Stanley	61,292,078,000	7,223,467,296	20080317	16,172,540,667	180
5	Dexia New York Branch	50,000,000,000	12,875,000,000	20080417	15,534,805,389	334
6	Consolidated BAC, Countrywide, and Merrill Lynch	48,141,416,451	10,091,065,527	20080310	14,120,484,053	288
7	Barclays	47,942,000,000	2,452,411,911	20080317	6,177,012,500	160
8	Depfa Bank PLC New York Branch	47,800,000,000	9,716,203,474	20080930	24,022,269,939	163
9	BNY Mellon	41,588,201,426	2,495,094,570	20080922	5,880,252,114	171
10	Merrill Lynch	39,956,500,000	7,620,802,284	20080310	16,164,122,739	190
11	Wachovia Bank NA	36,000,000,000	5,962,593,052	20080416	23,791,336,634	101
12	Lehman Brothers	28,000,000,000	195,555,571	20080318	5,253,926,333	15
13	Citigroup	24,200,000,000	6,712,770,766	20080318	12,524,289,901	211
14	Goldman Sachs	24,200,000,000	2,155,973,945	20080318	11,584,766,667	75
15	Royal Bank of Scotland Group, including CP*	20,460,500,000	5,531,579,799	20080229	9,604,281,409	364
16	Bank of America	13,000,000,000	1,895,449,347	20080619	4,106,806,919	186
17	Bank of Scotland PLC New York Branch	12,000,000,000	2,363,275,434	20080917	8,140,170,940	117
18	Royal Bank of Scotland PLC New York & RBS Citizens	8,400,000,000	49,284,697	20080229	2,837,390,429	7
19	Societe Generale New York Branch	8,000,000,000	2,584,880,893	20080325	3,960,863,118	263
20	Countrywide Financial	6,295,000,000	574,813,896	20080317	2,271,078,431	102
21	HSH Nordbank AG New York Branch	5,250,000,000	1,662,622,829	20080327	3,004,650,224	223

*Average calculated through October 31, 2009.

EXHIBIT 37.2 Kamakura Corporation, Analysis of Primary, Secondary, and Other Borrowings from the Federal Reserve Ranked by Average of Non-Zero Borrowing Days

Rank	Institution	Maximum Outstanding	Average Outstanding	First Borrowing Date	Average Of Non-Zero Borrowing Days	Number of Borrowing Days
1	AIG	208,616,483,142	72,793,773,121	20080916	161,186,211,912	182
2	Depfa Bank PLC New York Branch	47,800,000,000	9,716,203,474	20080930	24,022,269,939	163
3	Wachovia Bank NA	36,000,000,000	5,962,593,052	20080416	23,791,336,634	101
4	Consolidated JPMorgan, Bear Stearns, and WaMu	101,125,000,000	17,728,196,729	20080306	23,666,720,690	290
5	State Street	77,802,450,046	9,000,920,496	20080220	19,821,699,234	183
6	Morgan Stanley	61,292,078,000	7,223,467,296	20080317	16,172,540,667	180
7	Merrill Lynch	39,956,500,000	7,620,802,284	20080310	16,164,122,739	190
8	Dexia New York Branch	50,000,000,000	12,875,000,000	20080417	15,534,805,389	334
9	Consolidated BAC, Countrywide, and Merrill Lynch	48,141,416,451	10,091,065,527	20080310	14,120,484,053	288
10	Citigroup	24,200,000,000	6,712,770,766	20080318	12,524,289,901	211
11	Goldman Sachs	24,200,000,000	2,155,973,945	20080318	11,584,766,667	75
12	Royal Bank of Scotland Group, including CP*	20,460,500,000	5,531,579,799	20080229	9,604,281,409	364
13	Bank of Scotland PLC New York Branch	12,000,000,000	2,363,275,434	20080917	8,140,170,940	117
14	Barclays	47,942,000,000	2,452,411,911	20080317	6,177,012,500	160
15	BNY Mellon	41,588,201,426	2,495,094,570	20080922	5,880,252,114	171
16	Lehman Brothers	28,000,000,000	195,555,571	20080318	5,253,926,333	15
17	Bank of America	13,000,000,000	1,895,449,347	20080619	4,106,806,919	186
18	Societe Generale New York Branch	8,000,000,000	2,584,880,893	20080325	3,960,863,118	263
19	HSH Nordbank AG New York Branch	5,250,000,000	1,662,622,829	20080327	3,004,650,224	223
20	Royal Bank of Scotland PLC New York & RBS Citizens	8,400,000,000	49,284,697	20080229	2,837,390,429	7
21	Countrywide Financial	6,295,000,000	574,813,896	20080317	2,271,078,431	102

*Average calculated through October 31, 2009.

liquidity risk ratios are concrete symbols of the regulators' focus on default risk. Liquidity risk is the symptom that a firm has some other severe disease, be it credit risk, market risk, interest rate risk, fraud, or something else. Default becomes inevitable when it becomes apparent that an institution cannot liquidate its assets with sufficient speed or volume to meet cash needs from liabilities that have been withdrawn (in the case of deposits) or that will not be rolled over (commercial paper, bank lines, bonds, canceled insurance policies, and so on). For an institution that has not failed, data from institutions that have failed or have had "near death experiences" usually provide more insights on liquidity risk than the institution's own history of liability amounts and costs. This section features the funding shortfalls at AIG during the credit crisis for exactly that reason.

The primary, secondary, or other extensions of credit by the Federal Reserve to AIG during the period February 8, 2008 to March 16, 2009 can be summarized as follows:

First borrowing date:	Continuous from September 16, 2008
Average from 2/8/2008 to 3/16/2009	\$72.8 billion
Average when drawn	\$161.2 billion
Maximum drawn	\$208.6 billion

Kamakura Risk Information Services (KRIS) version 5.0 Jarrow-Chava default risk models showed a cumulative five-year default probability for AIG of 65.34 percent on the date of its first Fed borrowing, September 16, 2008, as shown in Exhibit 37.3.

The pattern of total outstanding borrowings from the Federal Reserve shows a steady increase from the first borrowing date on September 16, 2008, as shown in Exhibit 37.4.

In addition to these borrowings ("primary, secondary, or other extensions of credit") from the Federal Reserve, AIG-related entities were substantial beneficiaries of support from the Commercial Paper Funding Facility run by the Federal Reserve, as shown in Exhibit 37.5.




EXHIBIT 37.3 American International Group




EXHIBIT 37.4 AIG Primary, Secondary, and Other Extensions of Credit by the Federal Reserve, 2008–2009

Sources: Kamakura Corporation; Federal Reserve.




EXHIBIT 37.5 AIG Total Borrowings Outstanding under the Federal Reserve Commercial Paper Funding Facility, October 27, 2008, to January 25, 2010

Sources: Kamakura Corporation; Federal Reserve.

EXHIBIT 37.6 Chronology: JPMorgan Chase, Bear Stearns, and Washington Mutual

September 2, 2004	Washington Mutual chief risk officer Jim Vanasek sends internal Washington Mutual memo stating “At this point in the mortgage cycle with prices having increased far beyond the rate of increase in personal incomes, there clearly comes a time when prices must slow down or perhaps even decline.” (Levin Report, p. 66)
August 2006	Internal Washington Mutual presentation on Option ARM credit risk reported that from 1999 to 2006 Option ARM borrowers selected the minimum payment more than 95 percent of the time. (Levin Report, p. 59)
December 22, 2006	FDIC dedicated examiner at Washington Mutual Steve Funaro raises questions of senior management about unexpected increases in “early payment defaults” and demands from Wall Street firms doing securitizations to repurchase the loans. (Levin Report, p. 82)
December 31, 2006	Washington Mutual’s high risk loans begin incurring record rates of delinquency and default. (Levin Report, p. 48)
January 31, 2007	Washington Mutual chief risk officer Ron Cathcart identifies five top priority risk issues “in light of the slowdown and decline in home prices in some areas.” (Levin Report, p. 82)
January 31, 2007	The OTS found that, as of January 31, 2007, Washington Mutual had \$62 billion in outstanding Option ARMS in its investment portfolio, of which 80 percent were negatively amortizing. (Levin Report, p. 220)
February 18, 2007	Washington Mutual chief risk officer for mortgage lending notes in a memo “There is a meltdown in the subprime market” and that “many submarkets within California actually have declining home prices. . . .” (Levin Report, p. 128)
February 28, 2007	David Beck, Washington Mutual’s head of Wall Street securitizations, writes that securitizing second lien loans is “not a viable exit strategy” and that a Washington Mutual May 2006 securitization had already experienced 7 percent foreclosures. (Levin Report, p. 83)
June 7, 2007	Bear Stearns suspends redemption rights for hedge fund heavily invested in subprime debt market after losing 19 percent of value in April alone. (BusinessWeek.com)
June 17, 2007	Two Bear Stearns subprime hedge funds collapse. (Levin Report, p. 47)
June 2007	Washington Mutual shuts down Long Beach Mortgage Corporation as a separate entity. (Levin Report, p. 55)
August 1, 2007	Bear Stearns’ two troubled funds file for bankruptcy protection and the company freezes assets in a third fund. (Investopedia.com)
September 30, 2007	Washington Mutual halts subprime lending completely. (Levin Report, p. 55)

EXHIBIT 37.6 (Continued)

December 20, 2007	Bear Stearns reports its first quarterly loss in 84 years, \$854 million, after write downs of \$1.9 billion on mortgage holdings. (<i>Financial Times</i>)
December 31, 2007	At the end of 2007, 84 percent of the total value of Washington Mutual Option ARMS were negatively amortizing, meaning that the borrowers were going deeper into debt rather than paying off their balances. (Levin Report, p. 59)
January 16, 2008	JPMorgan Chase says it has cut the value of its investments in the U.S. subprime market by \$1.3 billion. (BBC)
March 13, 2008	Bear Stearns reports a \$15 billion loss in cash and cash equivalents in two days. Liquid assets also dropped \$2 billion mainly as a result of the loss of investor confidence. (<i>Wall Street Journal</i>)
March 14, 2008	Federal Reserve and JPMorgan Chase agree to provide emergency funding for Bear Stearns. Under the agreement, JPMorgan Chase would borrow from the Federal Reserve discount window and funnel the borrowings to Bear Stearns. (<i>Forbes</i> and DataCenterKnowledge.com)
March 16, 2008	JPMorgan Chase agrees to pay \$2 per share for Bear Stearns on Sunday, March 16, a 93 percent discount to the closing price on Friday March 14. JPMorgan Chase agreed to guarantee the trading liabilities of Bear Stearns, effective immediately. (<i>New York Times</i>)
March 16, 2008	Federal Reserve agrees to provide up to \$30 billion of financing to support Bear Stearns' "less liquid" assets. (<i>New York Times</i>)
March 24, 2008	JPMorgan Chase raises bid for Bear Stearns from \$2 per share to \$10 per share. JPMorgan Chase also agreed to bear the first \$1 billion of losses on Bear Stearns assets, with the Federal Reserve bearing the next \$29 billion of losses. (<i>The Times of London</i>)
March 24, 2008	Federal Reserve Bank of New York forms Maiden Lane I to help JPMorgan Chase acquire Bear Stearns. (Levin Report, p. 47)
April 8, 2008	Washington Mutual, the largest savings and loan association, receives a \$7 billion capital injection. (Reuters)
May 29, 2008	Bear Stearns shareholders approve sale, and the acquisition by JPMorgan Chase is completed. (Levin Report, p. 47)
July 26, 2008	Washington Mutual borrows \$10 billion from the Federal Reserve's discount window, the Federal Home Loan Banking system, and open market operations to bolster on-hand liquidity. (<i>Financial Times</i>)
July 31, 2008	Depositors withdraw \$10 billion from Washington Mutual in wake of IndyMac failure. (Levin Report, p. 57)
August 12, 2008	JPMorgan Chase reports \$1.5 billion in write downs in July, and bankers state July was the worst month for mortgage-backed securities since the start of the crisis. (<i>Financial Times</i>)

EXHIBIT 37.6 *(Continued)*

September 8, 2008	Washington Mutual ousts CEO Killinger. As a result of the credit crisis, the company's share price fell from \$40 in the summer of 2007 to \$3 in fall 2008. (<i>Financial Times</i>)
September 15–23, 2008	Depositors withdraw \$17 billion from Washington Mutual. (Levin Report, p. 57)
September 25, 2008	Washington Mutual fails, subsidiary bank is seized by the FDIC and sold to JPMorgan Chase for \$1.9 billion. JPMorgan Chase immediately wrote off \$31 billion in losses on the Washington Mutual assets. (<i>The Guardian</i> and Levin Report, p. 47).
September 26, 2008	JPMorgan Chase issues \$8 billion in common stock in conjunction with Washington Mutual takeover. (<i>The Guardian</i>)
September 28, 2008	Washington Mutual Inc., the holding company left after banking operations were seized, files for Chapter 11 bankruptcy. (<i>Seattle Times</i>)
October 28, 2008	United States uses TARP to buy \$125 billion in preferred stock at nine banks. The nine banks held over \$11 trillion in banking assets or roughly 75 percent of all assets owned by U.S. banks. (Levin Report, p. 47; SIGTARP report, p. 1).

The AIG borrowings in the form of primary, secondary, and other extensions of credit were extraordinary in both amount and in duration. The borrowings under the Commercial Paper Funding Facility, while less dramatic, are further confirmation that the liquidity risk and funding shortfall at AIG were very large. Every careful risk manager whose institution participates in the credit markets and securitized asset markets like AIG should have no illusions about the liquidity risk that can arise when large losses are incurred on the asset side.

Consolidated JPMorgan Chase, Bear Stearns, and Washington Mutual

This section focuses on the U.S. dollar funding shortfall that took place at JPMorgan Chase, Bear Stearns, and Washington Mutual during the period from February 8, 2008, to March 16, 2009. What happened surprised many: in combination with Bear Stearns and Washington Mutual, JPMorgan Chase was the most significant borrower from the Federal Reserve after AIG. JPMorgan Chase's peak borrowings, on a consolidated basis, were \$101.1 billion, nearly four times the \$28 billion that the Federal Reserve was willing to lend to Lehman Brothers only after Lehman Brothers declared bankruptcy on September 14, 2008.

The key dates in the chronology relevant to JPMorgan Chase and the two firms it acquired are summarized in Exhibit 37.6. We call your attention to the references to deposit runoff at Washington Mutual, which total \$27 billion. Washington Mutual's




EXHIBIT 37.7 JPMorgan Chase and Co.

borrowings from the Federal Reserve were only \$2 billion for eight days, so deposit runoffs were met largely via Federal Home Loan Bank borrowings and asset sales.

We now turn to primary, secondary, or other extensions of credit by the Federal Reserve to JPMorgan Chase, Bear Stearns, and Washington Mutual during the period February 8, 2008, to March 16, 2009. The three firms' consolidated borrowings from the Federal Reserve can be summarized as follows:

First borrowing date:	\$175 million on March 3, 2008
Average from 2/8/2008 to 3/16/2009	\$17.7 billion
Average when drawn	\$23.7 billion
Maximum drawn	\$101.1 billion on April 3, 2008
Number of days with outstanding borrowings	290 days

Exhibit 37.7 shows the one-month and one-year default probabilities for JPMorgan Chase from KRIS version 5.0 Jarrow-Chava reduced form credit model. Default probabilities began rising in an erratic pattern two months after the bankruptcy of Lehman Brothers on September 15, 2008, with the rise continuing through March 16, 2009, the last data point provided by the Federal Reserve.

JPMorgan Chase's first borrowing from the Federal Reserve was \$175 million on March 3, 2008, probably a "practice run" in anticipation of the Sunday March 16 announcement that JPMorgan Chase would absorb Bear Stearns in a Fed-supported rescue. The following graph shows the dual peaks in the consolidated borrowings in the names of JPMorgan Chase, Bear Stearns, and Washington Mutual. The first peak borrowing for the combined firms from the Fed was \$101.1 billion on April 3, 2008, two weeks after the rescue of Bear Stearns. The second peak was \$67.5 billion on October 15, 2008, one month after the failure of Lehman Brothers and three weeks after Washington Mutual was absorbed. Part of the means of the pay down following this second peak was a \$25 billion capital injection in late October 2008 under




EXHIBIT 37.8 JPMorgan Chase, Bear Stearns, and Washington Mutual Primary, Secondary and Other Extensions of Credit by the Federal Reserve, 2008–2009

Sources: Kamakura Corporation; Federal Reserve.

TARP's Capital Purchase Program, according to the October 5, 2009, report of the Special Inspector General of the Troubled Asset Relief Program.

There were no borrowings from the commercial paper funding facility by JPMorgan Chase, Bear Stearns, or Washington Mutual (see Exhibit 37.8).

JPMorgan Chase borrowed nearly four times more money, \$101.1 billion on April 3, 2008, than the \$28 billion withheld from Lehman Brothers until after Lehman Brothers declared bankruptcy. Sixth months later, JPMorgan Chase borrowed a second peak of \$67.5 billion on October 15, 2008, more than double what was refused to Lehman Brothers. Both before and after Lehman Brothers, JPMorgan Chase suffered larger funding shortfalls than Lehman Brothers. Clearly, JPMorgan Chase was too big to fail and Lehman Brothers was not.

State Street

This section focuses on the U.S. dollar funding shortfall that took place at State Street during the period from February 8, 2008, to March 16, 2009. The analysis shows that much of State Street's funding needs came from an unexpected source: losses in the firm's investment management business, which normally would be borne by the investors in the funds managed by State Street. We reach a conclusion that we have reached often in this liquidity risk series: State Street borrowed nearly three times more money from the Federal Reserve than Lehman Brothers was refused until after its bankruptcy was announced September 14, 2008.

The key dates in the credit crisis chronology relevant to State Street are summarized in Exhibit 37.9.

EXHIBIT 37.9 Chronology: State Street

November 1, 2007	State Street facing three lawsuits over bond fund losses. (<i>USAToday</i>)
January 3, 2008	State Street ousts head of State Street Global Advisors after losses. (Bloomberg)
January 4, 2008	State Street announces settlement of \$618 million on lawsuit tied to losses in pension funds managed by State Street due to declines in prices on mortgage-backed securities. (<i>New York Times</i>)
April 10, 2008	State Street hires new chief risk officer after derivatives index losses. (SeekingAlpha.com)
October 15, 2008	State Street may set aside as much as \$450 million in losses from propping up fixed income funds. (Bloomberg)
December 4, 2008	State Street announces plans to lay off 1,800 staff. (Finextra.com)
January 20, 2009	State Street stock plunges 59 percent when firm announces fourth quarter earnings and unrealized fixed income losses of \$6.3 billion. State Street also revealed it had spent \$3 billion more to prop up stable value funds. (Bloomberg).
January 21, 2009	State Street downgraded by Standard & Poor's to A+ from AA-. (Standard & Poor's).

State Street's borrowings from the Federal Reserve can be summarized as follows:

Borrowing dates:	First borrowing February 20, 2008, for \$605,000,000, with intermittent borrowings thereafter, followed by a \$3.3 billion borrowing on September 22, 2008. Borrowings were continually outstanding thereafter until the end of the Fed data series on March 16, 2009.
Average from 2/8/2008 to 3/16/2009	\$9.0 billion
Average when drawn	\$19.8 billion
Maximum drawn	\$77.8 billion on October 1, 2008
Number of days with outstanding borrowings	183 days

Exhibit 37.10 shows the one-month and one-year default probabilities for State Street from KRIS version 5.0 Jarrow-Chava reduced form credit model. Default probabilities began rising sharply after the January 20, 2009, earnings announcement that revealed \$9 billion of previously undisclosed losses.

Exhibit 37.11 shows the dramatic spike in borrowings that began on September 22, 2008, and peaked at \$77.8 billion on October 1, 2008. State Street received \$2 billion in capital under TARP's Capital Purchase Program on October 28, 2008, according to the SIGTARP Report. Most important, State Street's dramatic need for funds in the wake of the failures of Lehman Brothers, Washington Mutual,




EXHIBIT 37.10 State Street Corp.




EXHIBIT 37.11 State Street Primary, Secondary, and Other Extensions of Credit by the Federal Reserve, 2008–2009

Sources: Kamakura Corporation; Federal Reserve.

Wachovia, FNMA (Fannie Mae), and FHLMC (Freddie Mac) should be remembered by every serious analyst of liquidity risk at large financial institutions.

The Federal Reserve's disclosure of borrowings under the Commercial Paper Funding Facility lists 16 transactions in which State Street Bank & Trust is listed as "parent/sponsor at time of purchase."

We find once again, after reviewing State Street's borrowings from the Federal Reserve, that State Street borrowed \$50 billion more, a maximum of \$77.8 billion, than the \$28 billion that was refused to Lehman Brothers until that firm declared bankruptcy. This is further confirmation that the definition of "too big to fail" was dramatically revised immediately after Lehman Brothers' demise, since State Street's peak borrowing was on October 1, 2008, only 17 days after Lehman Brothers' September 14 bankruptcy announcement.

Morgan Stanley

This section focuses on the funding shortfall that took place at Morgan Stanley during the period from February 8, 2008, to March 16, 2009. While the bankruptcy of Lehman Brothers and the rescue of Merrill Lynch by Bank of America are emphasized by many analysts of the credit crisis, the analysis below confirms that Morgan Stanley, too, was in severe financial distress. The analysis confirms that Lehman Brothers borrowed less than half as much from the Federal Reserve after its bankruptcy than Morgan Stanley did when it was at the point of highest distress.

The key dates in the credit crisis chronology relevant to Morgan Stanley are summarized in Exhibit 37.12.

EXHIBIT 37.12 Chronology: Morgan Stanley

October 3, 2007	Morgan Stanley cuts 600 jobs. (FoxNews.com)
October 17, 2007	Morgan Stanley cuts another 300 bankers in credit trading, structured products, and leveraged lending areas. (Marketwatch.com)
November 7, 2007	Morgan Stanley reports \$3.7 billion in subprime losses. (Bloomberg)
December 19, 2007	Morgan Stanley announces \$9.4 billion in write downs from subprime losses and a capital injection of \$5 billion from a Chinese sovereign wealth fund. (<i>Financial Times</i>)
September 21, 2008	Goldman Sachs and Morgan Stanley convert to bank holding companies. (Levin Report, p. 47)
October 13, 2008	Morgan Stanley issued to Mitsubishi UFJ Financial Group 7,839,209 shares of Series B Non-Cumulative Non-Voting Perpetual Convertible Preferred Stock for an aggregate purchase price of \$7.8 billion plus \$1.2 billion in nonconvertible preferred stock. (October 16, 2008, and October 10, 2009, investor presentation by Morgan Stanley CFO Colm Kelleher)
October 28, 2008	U.S. uses TARP to buy \$125 billion in preferred stock at 9 banks. (Levin Report, p. 47)

Morgan Stanley's borrowings from the Federal Reserve can be summarized as follows:

Borrowing dates:	First borrowing \$2 billion on March 17, 2008, with borrowings continuously outstanding from September 15, 2008 to February 12, 2009
Average from 2/8/2008 to 3/16/2009	\$7.2 billion
Average when drawn	\$16.2 billion
Maximum drawn	\$61.3 billion on September 29, 2008
Number of days with outstanding borrowings	180

Exhibit 37.13 shows the Morgan Stanley one-month (upper line) and one-year default probabilities (lower line) for the KRIS Jarrow-Chava version 5.0 default probability models. Note that the default probabilities for Morgan Stanley continued to rise in November and December, after the September 14, 2008, bankruptcy filing by Lehman Brothers. Ultimately, the one-month default probability for Morgan Stanley peaked at over 14 percent.

Morgan Stanley's first borrowing from the Federal Reserve was \$2.0 billion on March 17, 2008, which was one day earlier than Lehman Brothers' first borrowing and \$400 million more than Lehman Brothers borrowed on March 18. Lehman Brothers' peak borrowing prior to filing for bankruptcy was \$2.73 billion on March 24, 2008, but Morgan Stanley's peak borrowing was \$61.3 billion on September 29, 2008. The Federal Reserve reported that Lehman Brothers had outstanding borrowings from the Federal Reserve on only 10 days prior to its bankruptcy filing on September 14, 2008, but Morgan Stanley had borrowings outstanding for 180 days.



EXHIBIT 37.13 Morgan Stanley

Sources: Kamakura Corporation; Federal Reserve.

Exhibit 37.14 shows the dramatic spike in Morgan Stanley's borrowings from the Federal Reserve.

After funding needs peaked at \$61.3 billion on September 29, 2008, Morgan Stanley received a capital injection of \$10 billion under the Capital Purchase Program on October 28, 2008. Morgan Stanley began borrowing under the Commercial Paper Funding Facility in late October, as shown Exhibit 37.15. Finally, as noted in the chronology in Exhibit 37.12, Morgan Stanley issued \$7.8 billion in convertible preferred stock and \$1.2 billion in nonconvertible preferred stock to Mitsubishi UFJ Financial Group on October 13, 2008.

Morgan Stanley began borrowing under the Commercial Paper Funding Facility on October 27, 2008, and did a total of 22 transactions, which are listed in Exhibit 37.15.

Morgan Stanley's peak funding shortfall of \$61.3 billion was more than double the \$28 billion funding shortfall at Lehman Brothers on September 15, 2008, the day after Lehman Brothers filed for bankruptcy. It is clear that the U.S. government dramatically revised its decision making on which firms were too big to fail in the aftermath of the Lehman Brothers bankruptcy. If the definition of "too big to fail" that was in place on September 14, 2008, had remained unchanged, it is very likely that Morgan Stanley would have failed as well.

Dexia Credit Local New York Branch

This section concentrates on a major European institution, specifically on the funding shortfall that took place at Dexia SA during the period from February 8, 2008, to




EXHIBIT 37.14 Morgan Stanley Primary, Secondary, and Other Extensions of Credit by the Federal Reserve, 2008–2009

Sources: Kamakura Corporation; Federal Reserve.

EXHIBIT 37.15 Extensions of Credit by the Federal Reserve under the Commercial Paper Funding Facility

Trade and Settlement Date	Maturity Date	Commercial Paper Type	Amount (USD Millions)	Discount Rate (Percent)	Credit Enhancement Surcharge (Percent)	Total Discount Rate (Percent)	Issue Name	Parent/Sponsor at Time of Purchase
10/27/2008	1/26/2009	CP	149.3	2.70	0.00	2.70	MORGAN STANLEY	Morgan Stanley
10/28/2008	1/26/2009	CP	139.3	2.70	0.00	2.70	MORGAN STANLEY	Morgan Stanley
11/4/2008	2/2/2009	CP	249.0	1.60	0.00	1.60	MORGAN STANLEY	Morgan Stanley
11/5/2008	2/3/2009	CP	199.2	1.55	0.00	1.55	MORGAN STANLEY	Morgan Stanley
11/6/2008	2/4/2009	CP	199.2	1.54	0.00	1.54	MORGAN STANLEY	Morgan Stanley
11/7/2008	2/5/2009	CP	199.2	1.54	0.00	1.54	MORGAN STANLEY	Morgan Stanley
11/10/2008	2/9/2009	CP	199.2	1.53	0.00	1.53	MORGAN STANLEY	Morgan Stanley
11/12/2008	2/10/2009	CP	199.3	1.47	0.00	1.47	MORGAN STANLEY	Morgan Stanley
11/13/2008	2/11/2009	CP	199.2	1.52	0.00	1.52	MORGAN STANLEY	Morgan Stanley
11/14/2008	2/12/2009	CP	219.2	1.54	0.00	1.54	MORGAN STANLEY	Morgan Stanley
11/17/2008	2/17/2009	CP	199.2	1.48	0.00	1.48	MORGAN STANLEY	Morgan Stanley
11/18/2008	2/17/2009	CP	199.3	1.48	0.00	1.48	MORGAN STANLEY	Morgan Stanley

11/19/2008	2/17/2009	CP	149.4	1.48	0.00	1.48	MORGAN STANLEY	Morgan Stanley
11/20/2008	2/18/2009	CP	149.5	1.42	0.00	1.42	MORGAN STANLEY	Morgan Stanley
11/21/2008	2/19/2009	CP	124.5	1.49	0.00	1.49	MORGAN STANLEY	Morgan Stanley
11/24/2008	2/23/2009	CP	249.1	1.49	0.00	1.49	MORGAN STANLEY	Morgan Stanley
11/25/2008	2/23/2009	CP	199.3	1.49	0.00	1.49	MORGAN STANLEY	Morgan Stanley
11/26/2008	2/24/2009	CP	199.3	1.42	0.00	1.42	MORGAN STANLEY	Morgan Stanley
11/28/2008	2/26/2009	CP	159.4	1.41	0.00	1.41	MORGAN STANLEY	Morgan Stanley
12/2/2008	3/2/2009	CP	199.3	1.39	0.00	1.39	MORGAN STANLEY	Morgan Stanley
12/3/2008	3/3/2009	CP	179.4	1.37	0.00	1.37	MORGAN STANLEY	Morgan Stanley
12/4/2008	3/4/2009	CP	368.8	1.33	0.00	1.33	MORGAN STANLEY	Morgan Stanley

EXHIBIT 37.16 Chronology: Dexia SA

October 31, 2007	Dexia's subsidiary FSA reports \$190 million in third quarter mark-to-market losses. (Reuters)
August 6, 2008	Dexia injects \$300 million in FSA. (Dexia.com)
August 7, 2008	Dexia's FSA Faux Pas. (<i>Forbes</i>)
September 18, 2008	Dexia announces Lehman Brothers-related losses estimated to be around €350 million. (Dexia press release)
September 30, 2008	European governments bail out Dexia in €6.4 billion rescue deal. (<i>London Telegraph</i>)
October 1, 2008	Dexia chairman and president resign. (Associated Press)
November 14, 2008	Dexia sells Financial Security Assurance after €1.5 billion Third Quarter Loss. (<i>Financial Times</i>)
March 9, 2009	Dexia posts €3.3 billion Euro losses for 2008. (Neurope.de)

March 16, 2009. The analysis confirms that Dexia, through the New York Branch of Dexia Credit Local, was one of the largest borrowers from the Federal Reserve during the credit crisis, a fact that hasn't received as much attention as it should. At its peak, Dexia was borrowing \$50 billion from the Federal Reserve, almost double the funding shortfall at Lehman Brothers on September 15, 2008, the day after its bankruptcy filing. This amount of aid is four times the Dexia rescue amount that was announced by the Netherlands, Belgium, and Luxembourg on September 30, 2008.

The key dates in the credit crisis chronology relevant to Dexia are summarized in Exhibit 37.16.

Dexia's losses have continued long after the September 30, 2008, bailout by European governments. On May 27, 2011, Bloomberg reported that Dexia would make a \$5.1 billion provision for losses on accelerated asset sales.

Dexia's borrowings from the Federal Reserve can be summarized as follows:

Borrowing dates:	First borrowing \$6 billion on April 17, 2008, with borrowings continuously outstanding through March 16, 2009, even after the September 30, 2008, bail-out by three European governments.
Average from 2/8/2008 to 3/16/2009	\$12.9 billion
Average when drawn	\$15.5 billion
Maximum drawn	\$50.0 billion on March 12, 2009
Number of days with outstanding borrowings	334

Exhibit 37.17 shows the Dexia one-month (lower line) and one-year default probabilities (upper line) for the KRIS Jarrow-Chava version 5.0 default probability models from February 8, 2008, through the effective nationalization of Dexia on September 30, 2008.




EXHIBIT 37.17 Dexia SA




EXHIBIT 37.18 Dexia Primary, Secondary, and Other Extensions of Credit by the Federal Reserve, 2008–2009
Sources: Kamakura Corporation; Federal Reserve.

Dexia’s first borrowing from the Federal Reserve was \$6 billion on April 17, 2008. Dexia Credit Local New York Branch borrowed from the Federal Reserve for 311 days. The peak borrowing was \$50 billion on March 12, 2009, five months and 12 days after Belgium, the Netherlands, and Luxembourg announced their €6.4 billion rescue of Dexia SA (Exhibit 37.18).

Dexia began borrowing under the Commercial Paper Funding Facility on October 27, 2008, and did a total of 42 transactions, which are listed in Exhibit 37.19.

EXHIBIT 37.19 Kamakura Corporation, Extension of Credit by the Federal Reserve under the Commercial Paper Funding Facility

Trade and Settlement Date	Maturity Date	Commercial Paper Type	Amount (USD Millions)	Discount Rate (Percent)	Credit Enhancement		Total Discount Rate (Percent)	Issue Name	Parent/Sponsor at Time of Purchase
					Surcharge (Percent)				
10/27/2008	1/26/2009	CP	2,978.2	1.87	1.00		2.87	DEXIA DELAWARE LLC	Dexia SA
10/29/2008	1/27/2009	CP	2,978.7	1.83	1.00		2.83	DEXIA DELAWARE LLC	Dexia SA
10/30/2008	1/28/2009	CP	2,979.5	1.73	1.00		2.73	DEXIA DELAWARE LLC	Dexia SA
10/31/2008	1/29/2009	CP	1,987.0	1.60	1.00		2.60	DEXIA DELAWARE LLC	Dexia SA
11/4/2008	2/2/2009	CP	993.5	1.60	1.00		2.60	DEXIA DELAWARE LLC	Dexia SA
11/5/2008	2/3/2009	CP	496.8	1.55	1.00		2.55	DEXIA DELAWARE LLC	Dexia SA
11/6/2008	2/4/2009	CP	298.1	1.54	1.00		2.54	DEXIA DELAWARE LLC	Dexia SA
11/7/2008	2/5/2009	CP	794.9	1.54	1.00		2.54	DEXIA DELAWARE LLC	Dexia SA
11/10/2008	2/9/2009	CP	496.8	1.53	1.00		2.53	DEXIA DELAWARE LLC	Dexia SA
11/12/2008	2/10/2009	CP	1,490.7	1.47	1.00		2.47	DEXIA DELAWARE LLC	Dexia SA
11/14/2008	2/12/2009	CP	1,192.4	1.54	1.00		2.54	DEXIA DELAWARE LLC	Dexia SA
11/17/2008	2/17/2009	CP	496.8	1.48	1.00		2.48	DEXIA DELAWARE LLC	Dexia SA
11/18/2008	2/17/2009	CP	745.3	1.48	1.00		2.48	DEXIA DELAWARE LLC	Dexia SA
11/19/2008	2/17/2009	CP	248.5	1.48	1.00		2.48	DEXIA DELAWARE LLC	Dexia SA
11/20/2008	2/18/2009	CP	497.0	1.42	1.00		2.42	DEXIA DELAWARE LLC	Dexia SA
11/24/2008	2/23/2009	CP	795.0	1.49	1.00		2.49	DEXIA DELAWARE LLC	Dexia SA
11/28/2008	2/26/2009	CP	994.0	1.41	1.00		2.41	DEXIA DELAWARE LLC	Dexia SA
12/1/2008	2/27/2009	CP	994.1	1.42	1.00		2.42	DEXIA DELAWARE LLC	Dexia SA
1/27/2009	4/27/2009	CP	2,983.2	1.24	1.00		2.24	DEXIA DELAWARE LLC	Dexia SA
1/28/2009	4/28/2009	CP	2,983.4	1.22	1.00		2.22	DEXIA DELAWARE LLC	Dexia SA
1/29/2009	4/29/2009	CP	1,988.9	1.22	1.00		2.22	DEXIA DELAWARE LLC	Dexia SA
2/2/2009	5/4/2009	CP	994.3	1.26	1.00		2.26	DEXIA DELAWARE LLC	Dexia SA
2/3/2009	5/4/2009	CP	497.2	1.26	1.00		2.26	DEXIA DELAWARE LLC	Dexia SA
2/5/2009	5/6/2009	CP	497.2	1.25	1.00		2.25	DEXIA DELAWARE LLC	Dexia SA

2/9/2009	5/11/2009	CP	497.1	1.27	1.00	2.27	DEXIA DELAWARE LLC	Dexia SA
2/10/2009	5/11/2009	CP	994.3	1.27	1.00	2.27	DEXIA DELAWARE LLC	Dexia SA
2/12/2009	5/13/2009	CP	696.0	1.27	1.00	2.27	DEXIA DELAWARE LLC	Dexia SA
2/17/2009	5/18/2009	CP	1,093.8	1.27	1.00	2.27	DEXIA DELAWARE LLC	Dexia SA
2/18/2009	5/19/2009	CP	497.2	1.27	1.00	2.27	DEXIA DELAWARE LLC	Dexia SA
2/23/2009	5/26/2009	CP	795.4	1.23	1.00	2.23	DEXIA DELAWARE LLC	Dexia SA
2/26/2009	5/27/2009	CP	994.4	1.25	1.00	2.25	DEXIA DELAWARE LLC	Dexia SA
2/27/2009	5/28/2009	CP	994.4	1.25	1.00	2.25	DEXIA DELAWARE LLC	Dexia SA
4/27/2009	7/27/2009	CP	2,981.3	1.20	1.00	2.20	DEXIA DELAWARE LLC	Dexia SA
4/29/2009	7/28/2009	CP	2,983.5	1.20	1.00	2.20	DEXIA DELAWARE LLC	Dexia SA
4/30/2009	7/29/2009	CP	1,989.0	1.20	1.00	2.20	DEXIA DELAWARE LLC	Dexia SA
5/4/2009	7/31/2009	CP	1,491.9	1.21	1.00	2.21	DEXIA DELAWARE LLC	Dexia SA
5/6/2009	8/4/2009	CP	497.2	1.21	1.00	2.21	DEXIA DELAWARE LLC	Dexia SA
5/27/2009	8/25/2009	CP	596.7	1.21	1.00	2.21	DEXIA DELAWARE LLC	Dexia SA
5/28/2009	8/26/2009	CP	994.5	1.21	1.00	2.21	DEXIA DELAWARE LLC	Dexia SA
7/29/2009	10/27/2009	CP	1,989.0	1.20	1.00	2.20	DEXIA DELAWARE LLC	Dexia SA
7/31/2009	10/29/2009	CP	1,491.8	1.20	1.00	2.20	DEXIA DELAWARE LLC	Dexia SA
8/4/2009	11/2/2009	CP	497.3	1.20	1.00	2.20	DEXIA DELAWARE LLC	Dexia SA

Dexia's peak funding need of \$50 billion on March 12, 2009, was almost double the \$28 billion funding shortfall at Lehman Brothers on September 15, 2008, the day after Lehman Brothers filed for bankruptcy. Again, it is clear that the U.S. government dramatically revised its decision making on which firms were too big to fail in the aftermath of the Lehman Brothers bankruptcy. What is surprising is the Fed's willingness to provide double the aid that Lehman Brothers required to a non-U.S. institution after a supposed rescue by three European governments. In reality, obviously, the United States was an eager and active participant in the September 30, 2008, rescue of Dexia SA. The Internet is silent on the "thank you" to the United States from the shareholders of Dexia SA (who would otherwise have been completely wiped out) and the three European nations whose September 30, 2008, bailout was only roughly one-fourth of the aid provided by the Federal Reserve at its peak, excluding the other amounts extended through the Commercial Paper Funding Facility.

IMPLICATIONS OF THE CREDIT CRISIS HISTORY FOR LIQUIDITY RISK MANAGEMENT AND ANALYSIS

In the remainder of chapter, we emphasize the fact that the Jarrow-Merton put option concept can be used to measure liquidity risk in dollar terms as well serving as an index of total risk. In effect, we are going to use our knowledge of all of the financial institution's assets and liabilities to price the cost of an irrevocable line of credit to our institution that would supply any random cash needs from now until the maturity of the line of credit. In doing so, we will take advantage of our previous discussions from Chapter 33 on pricing and valuing revolving credits and exposure at default.

We discuss liquidity risk in a step-by-step process, covering these topics:

- Types of liquidity events
- Liquidity risk and credit risk linkages
- Measuring liquidity risk as a "line of credit" in the Jarrow-Merton put option sense
- Integrating managerial behavior and market funds supply in liquidity risk measurement

Types of Liquidity Events

In spite of the previous examples, liquidity risk analysis often seems to have a lot in common with the story about the group of blind men, each touching a different part of the elephant, each of them believing the elephant to be a different type of animal from the others' perceptions. In order to clearly reveal the true nature of the elephant, we want to clarify the types of liquidity conditions that can come about.

Liquidity risk crises tend to revolve around a combination of borrower credit quality (risky or safe) and market conditions (troubled or normal). Besides the 2006–2011 credit crisis events described in the previous sections, there are many examples of liquidity crises. The Asia crisis in 1997 resulted in liquidity crises from Malaysia to Korea as a combination of very difficult market conditions in the foreign exchange and money markets with a sharp rise in the credit risk of financial institutions. This increase in credit risk was in large part directly attributable to the sharp changes in

macro factors that drive default risk: foreign exchange rates and interest rates. This kind of incident fits precisely in the analytical framework that we have been using throughout this book because we have successfully linked credit risk to macroeconomic conditions. This is the central theme of the reduced form modeling techniques outlined in Chapter 16. Another related incident was the 1998 Long-Term Capital Management incident, where one troubled institution's holdings were so large that its need to sell off its portfolio rapidly depressed a wide array of security prices and raised the credit risk of a broad range of corporations and financial institutions.

The second type of crisis is a combination of normal market conditions and credit quality problems of a single or small number of institutions. One example is the 1992 forced acquisition of Security Pacific Corporation by Bank of America. In this case, sharp increases in credit losses in real estate at Security Pacific National Bank led regulators to cut off dividend payments to the parent company, Security Pacific Corporation. Investors were no longer willing to buy commercial paper of Security Pacific Corporation because its default risk had risen too high.

Another example of troubled institutions, in what were otherwise fairly normal markets, was the savings and loan crisis of the 1980s and early 1990s in the United States. A rise in interest rates caused a sharp decline in the market value of fixed rate home mortgages on the books of the savings and loan associations. The market value of their assets fell well below the principal amount of consumer deposits on their books. In anticipation of their failure, consumers began withdrawing their deposits in spite of the U.S. government guarantee of consumer deposits. Even if the savings and loan associations had been able to sell their mortgage holdings, they would not have had sufficient assets to meet deposit obligations. The institutions were taken over by Federal regulators and liquidated.

A final example is that of Barings in 1995, where a rogue trader in Singapore generated more than \$1 billion in losses from unauthorized futures trading. The end result, due to the impaired credit quality of Barings and the liquidity problems that it generated, was the forced sale of Barings to the ING Group.

A third category of liquidity problems is a systemic market crisis that isn't directly related to a troubled financial institution. Besides the 2006–2011 credit crisis, perhaps the best illustration was the impact of the September 11, 2001, attack on the United States, which caused a serious disruption in operational cash flows in the U.S. financial system. While insurance-related losses from September 11, 2001, did have a concrete impact on the credit quality of many financial institutions, most institutions did not suffer an impact on their credit quality. Nonetheless, they did have a cash shortfall. Other examples are disruptions caused by power failures (for example, New York and Toronto power failure of 2003), fires (First Interstate Bancorp headquarters, in 1988), and other natural disasters.

The latter kind of crisis produces a very acute need for cash in the short term, but the need comes from an institution that has fundamentally sound credit quality in the long term. This is a very important distinction that we explore in the next section.

Liquidity Risk and Credit Risk Linkages

As we saw in the previous section and in the five examples from the 2006–2011 crisis, most liquidity crises occur when the financial institution itself is troubled. An important related issue is in how much company the financial institution has in its

troubles—it can be alone in its difficulties (Security Pacific, 1992), or it can have lots of company (Korea, 1997–1998, 2006–2011 credit crisis). For that reason, we need to be more precise about the impact of a decline in credit quality of a financial institution on the willingness of other firms to lend it money. Is the decline in credit quality due to a movement in macro factors that affect many firms, or is it a Barings-type idiosyncratic risk that has gone wrong?

In either case, the market reaction to one or many firms in trouble is simple: the market reduces its extension of credit to troubled firms, and a lender of last resort (like the Federal Reserve in the preceding five examples) either steps in or stands aside to let the firms fail (as in the savings and loan crisis in the United States in the late 1980s and early 1990s).

Measuring Liquidity Risk as a Line of Credit in the Jarrow-Merton Put Option Sense

How do institutions buy time when they have a credit quality problem that they hope will be rectified in time? That is the essence of liquidity risk management discipline. The root cause of the problem is a credit quality problem that affects both the financial institution itself (in the case of a company-specific problem in normal market conditions like Security Pacific in 1992) and perhaps many other peer group institutions as well (such as the 2006–2011 credit crisis, Mexico in 1995, or Korea in 1997). If the credit problems are driven by the credit cycle, as we discussed in Chapter 16, it is logical to think that if management can assure itself three to five years to fix the problem, they are likely to be successful. If the liquidity posture of the bank is so focused on short-term financing, the financial institution may face a cash flow need that they can't meet and fail before the credit problem can be solved. This would have happened en masse in the 2006–2011 credit crisis if the Federal Reserve had not lent to troubled firms so aggressively.

The best way to assure one has years, not days, to deal with a company-specific or marketwide credit crisis is to secure protection against unforeseen funding needs with (at least in concept) an irrevocable line of credit. What is the coherent measure of the cost of liquidity risk protection for, say, 10 years?

Again, the answer is the Jarrow-Merton put option in the context of Chapter 19 (risky corporate credit risk assessment), Chapter 20 (credit default swaps), and Chapter 33 (lines of credit) where we priced a revolving line of credit to a corporate borrower. In this case, we want the price of a 10-year irrevocable line of credit for an appropriate amount of borrowings with a 10-year life and a funds provider who is not in the same market as the borrower. We remind readers of the cash needs from retail deposits who fled Washington Mutual and Wachovia, triggering their failure, in the 2006–2011 credit crisis. We specifically included cash needs from this source in our example of nonmaturity deposit analysis in Chapter 30. We know this about an irrevocable line of credit with a 10-year maturity:

- It solves the liquidity crisis for 10 years and buys up to 10 years of time to solve the credit risk, market risk, foreign exchange risk and/or interest rate risk problem permanently.
- If the general risk level of the financial institution is low and the funding profile of the financial institution is very long term, this line of credit will be very cheap

and it is confirmation that the liquidity risk of the institution is low. In an extreme case of a financial institution where all liabilities are longer than 10 years in maturity and there are no random cash needs, the price of this line of credit will be zero, because it will never be drawn down.

- If the general risk level of the financial institution is low and the funding profile of the financial institution is very short term, the line of credit will be moderately expensive.
- If the general risk level of the financial institution is high and the funding profile of the institution is very long (all liabilities longer than 10 years), then again, the line of credit will be fairly cheap because it will not be drawn if cash needs are not random.
- If the general risk level of the financial institution is high and the funding profile of the institution is very short, then the line of credit will be very expensive because the liquidity risk of the institution is very high.
- If the general risk level of the institution gets high enough, no lender will provide quotations for the line of credit because of the same phenomenon we studied previously—the supply of credit of all types falls as risk rises.

Integrating Managerial Behavior and Market Funds Supply in Liquidity Risk Measurement

How can we price the cost of a 10-year irrevocable line of credit for our own institution if we are not going to get actual market quotes to measure the degree of liquidity risk that we face? We can use exactly the same valuation and simulation techniques that we discussed in Chapter 36, but with a few important differences that revolve around building in a logical market and managerial responses to the changing interest rate risk and credit risk of the institution.

As the default probability of our financial institution rises, consumers will begin to withdraw consumer deposits. We need to specify the formula by which this will happen, exactly as we did in the nonmaturity deposit in Chapter 30.

As the default probability of the financial institution rises, our access to commercial paper and bond markets will decrease or even dry up completely. We need to specify how access decreases in response to the default probability. As the default probability of the institution rises, management will begin to sell assets to meet cash needs. As these sales increase in amount and urgency, the transactions cost of each sale will get larger and larger as the percent of value of the assets sold. If the urgency and volume of sales grows too high, there will be “no bid” in the market as was seen for much of 2007 and 2008 in the credit crisis in the market for CDOs and mortgage-backed securities.

We value the 10-year irrevocable line of credit as a measure of liquidity risk using a process exactly like that used in Chapter 36, but we factor in a rational market and managerial response to a change in our institution’s default probability:

1. Specify how our access to the commercial paper market varies with our default probability as discussed in this chapter.
2. Specify how our access to the bond market varies with our default probability as discussed in this chapter.
3. Specify how consumers will withdraw consumer deposits as our default probability rises.

4. Specify the order in which assets will be sold to meet cash needs and the transactions cost of selling those assets as a function of the asset class and the speed with which we have to sell it.
5. Simulate the risk-free term structure for N scenarios over M time steps.
6. Choose a formula and the risk-drivers for the liquidity component of credit spread for all of the relevant asset classes we own as discussed in Chapter 17.
7. Choose a formula and the risk drivers for the default intensity process in the Jarrow model as discussed in Chapter 17 for all asset classes that we own.
8. Simulate the random values of the drivers of liquidity premium and the default intensity for all of our counterparties for M time periods over N scenarios, consistent with the risk-free term structure.
9. Calculate the default intensity and the liquidity component at each of the M time steps and N scenarios for each counterparty, from retail to small business to corporate to sovereign. Note that these will be changing randomly over time because interest rates and other factors impact each of them. This is essential to capturing cyclicalities in bond prices and defaults.
10. Apply the Jarrow model as we have done in Chapters 16 and 17 to get the zero-coupon bond prices for each maturity for each transaction for each counterparty.
11. Calculate the market value of portfolio equity for our financial institution in scenario J and time period K .
12. Calculate the probability of default for our own financial institution in scenario J and time period K (we can shortcut this by using the logistic regression approach in Chapter 16 or we can be more fancy).
13. Calculate the amount of consumer deposits that we will have in scenario J and time period K based on this default probability. Do the same for commercial paper and our bond issuance capability.
14. Calculate the net needs for funds based on Steps 12 and 13.
15. Subtract from this need for funds liquid assets we can sell with no transactions costs, based on the amount of those liquid assets still remaining in time period K of scenario J .
16. The remaining funding need is X .
17. If the remaining availability on our irrevocable line of credit is more than X , we assume we draw down an additional X on the line of credit.
18. We calculate the interest rate owed on the line of credit in time period K and scenario J .
19. We simulate default/no default in period K of scenario J .
20. If we do not default, we proceed to time period $K + 1$ in Scenario J .
21. If we do default, the lender on the line of credit receives our adjusted market value of portfolio assets.
22. We do this for all periods and scenarios.
23. We take the present value of these risk-adjusted, risk-neutral cash flows by discounting at the risk-free yield appropriate for time period K and scenario J .
24. We calculate the value of the line of credit to the lender and the fee they need to charge us for the line of credit to be fairly priced.

This process is exactly the same process that we followed in Chapter 33 for a revolving line of credit to a corporate borrower with one major exception. Since we

are evaluating the cost of a line of credit to our own institution, we have full information on every asset and every liability on the balance sheet. We know how to value each one of them, as we have demonstrated in Chapters 19 to 35. Therefore, we can be much more precise in our calculations as the previous simulation process outlines.

DETERMINING THE OPTIMAL LIQUIDITY STRATEGY

What is the optimal maturity structure of our institution's liabilities from a liquidity risk perspective? This question is easy to answer since we now have an explicit process for measuring our current liquidity risk—the cost of a 10-year irrevocable line of credit using the Jarrow-Merton put option concept. The optimal liquidity strategy is that which maximizes the risk-neutral value of the firm. This will not necessarily be the liability maturity structure that minimizes the cost of eliminating liquidity risk; that is, the liability maturity structure that produces the lowest cost for the 10-year irrevocable line of credit.

SUMMING UP

In Chapter 36, we showed how to use the Jarrow-Merton put option as a comprehensive index of the level of integrated credit risk, market risk, foreign exchange risk, and interest rate risk. This measure of risk is more than that—it tells us the dollar cost of eliminating the risk.

We can do the same thing with liquidity risk by pricing an irrevocable line of credit for any maturity. We used a 10-year maturity since 10 years is likely to get us all the way through a full interest rate cycle and credit cycle, giving us time to recover from interest rate or credit risk-related losses that would have produced a liquidity crisis if we did not have (hypothetically) the irrevocable line of credit. In practice, the cost of liquidity protection for any user-defined maturity can be calculated using the same process.

Next, we take the discipline and insights of the Jarrow-Merton put option measure of risk to performance measurement in order to improve on the conventional wisdom of the early chapters in this book.

NOTES

1. For a less vivid but equally detailed account, see the Levin Report (2011), whose full title is U.S. Senate, Senator Carl Levin, Chairman, Majority and Minority Staff Report, Permanent Committee on Investigations, *Wall Street and the Financial Crisis: Anatomy of a Financial Collapse*, April 13, 2011, www.hsgac.senate.gov/imo/media/doc/Financial_Crisis/FinancialCrisisReport.pdf?attempt=2.
2. For a reader interested in all 21 institutions, www.kamakuraco.com contains daily liquidity needs of Bank of America, Countrywide Financial, Merrill Lynch, a consolidation of the latter three firms, Lehman Brothers, Morgan Stanley, Citigroup, Dexia SA, Depfa Bank plc, Barclays, Goldman Sachs, the combined JPMorgan Chase, Washington Mutual,

and Bear Stearns, Wachovia, State Street, BNY Mellon, HSH Nordbank AG, Societe Generale, and HBOS/Bank of Scotland.

3. A summary of the Federal Reserve programs that were put into place and summary statistics are available from the Federal Reserve at www.federalreserve.gov/newsevents/reform_transaction.htm.
4. For information regarding the Kamakura Credit Crisis Liquidity Risk data base, contact e-mail: info@kamakuraco.com.

Performance Measurement

Plus Alpha vs. Transfer Pricing

In Chapter 2, we introduced the common practice of performance measurement in large financial institutions. There are a number of key points about the way in which large financial institutions manage risk:

- Commercial banks use a mix of financial accounting and market-oriented performance measurements.
- Transfer pricing in commercial banks is a financial accounting–driven performance measure that has the virtue of moving interest rate risk into one central transfer pricing book. The system has a 40-year history, but it mistakenly implies that each new transaction should be judged versus the average marginal risk of the bank (via the bank’s marginal cost of funds) instead of the marginal risk of the asset being considered. The fact that a traditional transfer pricing system would reject the purchase of risk-free U.S. Treasuries shows this process is flawed at the transaction level, since all banks in the United States own at least some U.S. Treasuries.
- Insurance companies are moving to adopt the transfer pricing technology as a way of finally integrating risk assessment of insurance assets and liabilities.
- Fund managers, pension funds, and insurance companies (on the asset side of the balance sheet) almost all base performance measurement on a market price basis. A good performer is one who can generate plus-alpha performance over and above a predefined benchmark with equivalent risk. Banks (until recently) have almost never used this approach outside of their own fund management business, even in their own trading activities.
- Banks allocate capital to judge the performance of major business units, and almost no other financial institutions do. The capital allocation is done so that each business unit can be judged using the financial accounting measures that would be applied to the business unit if it were a stand-alone financial institution.

Our purpose in this chapter is to talk about the implications of the Jarrow-Merton put option as a potential guide to performance measurement. Matten (1996) presents an excellent summary of current common practice in commercial bank capital allocation that will be of interest to readers who would like to supplement the material in Chapter 2 before plunging into the rest of this chapter.

TRANSACTION-LEVEL PERFORMANCE MEASUREMENT VS. PORTFOLIO-LEVEL PERFORMANCE MEASUREMENT

There is a central tenet of performance management that we need to keep in mind for the rest of this chapter:

A performance measurement system which sends incorrect signals at the transaction level will be inaccurate at the portfolio level as well, since the portfolio is the sum of its individual transactions.

In Chapters 2 and in van Deventer and Imai (2003), buy low/sell high was outlined as the ultimate transaction-level guide to asset and liability selection. If an asset worth 102 can be purchased at a cost of 100, we should do it. If another asset worth 101 can be purchased at a cost of 100, we should buy that, too. What could be simpler than that? This basic performance measurement system has been at the heart of capitalism for thousands of years, and our objective in this chapter is to strip back some of the institutional baggage of financial institutions' performance measurement systems to reveal performance as clearly as these examples.

In Chapters 19 through 35, we showed that every single type of asset and liability issued by major financial institutions can be valued precisely both before the asset or liability is added to the balance sheet and every instant thereafter. We now know enough to do the simple assessment in the previous paragraph. Are there any real or imagined complexities that may complicate the buy low/sell high transaction-level performance assessment? There are only a few and we mention them briefly in this section:

- *Costs of origination* may have both a large fixed cost component and a marginal cost component, which impact our assessment of buy low/sell high. A classic example is the consumer deposit gathering franchise of a commercial bank. On the transaction level, a 90-day consumer certificate of deposit may cost 25 basis points less than a wholesale certificate of deposit (CD) with the identical maturity. Based on this information alone, the first-pass buy low/sell high decision is buy—originate the retail CD. The costs of origination and servicing need to be included in the analysis for a complete answer. What if there was a \$1 per month fixed cost of servicing the consumer CD that would not impact the wholesale CD? We just need to add this to the effective cost of the retail CD to have a more accurate answer. More realistically, the head of the retail banking division should be constantly re-evaluating this calculation: I can sell the building housing my River City branch for \$500,000 today. Or I can keep it, pay staff expenses of \$450,000 per year, and issue a large number of consumer CDs. Which is better? This is still a buy low/sell high calculation that requires no capital allocation to do correctly. It is a question that should be continually asked. A related question is whether the 25 basis point margin on the consumer CD is the margin that produces maximum present value (that marginal revenue equals marginal cost from basic economic theory is relevant here). There are other calculations that come into play that we mention next.
- *Future business may depend on doing current business*, with an actuarially likely arrival of good buy low/sell high business going forward. The example above is a good example to illustrate another complexity in assessing performance of any

business. The River City branch's existence may give us the real option to do lots of retail CD business going forward. This retail business depends both on our pricing and our own default probability as we saw in Chapter 37 on liquidity risk. The probability of a client rolling over one profitable retail CD has characteristics that are very similar to that of the reduced form credit models we discussed in Chapter 16, the credit default swaps we discussed in Chapter 20, and the insurance policies we valued in Chapter 35. Doing a really good buy low/sell high decision analysis depends on taking these facts into account in a careful way. It has nothing to do with capital allocation—it's a straightforward net present value calculation using the analytics presented in Chapters 19 to 35.

- *Some costs may be shared with other products.* What if the River City branch also makes profitable small business loans (valued correctly using the technology of Chapter 19 and 20), and that there is a fixed cost of monitoring them that is in part embedded in the \$450,000 staff costs of the River City branch? Again, assessing whether to close the branch or keep it and optimize pricing is a risk-adjusted present value calculation like those in Chapters 19 through 35—shared fixed costs simply broaden the boundaries of the calculation that has to be done and the transactions that have to be included. Capital allocation is again irrelevant. What is relevant is the available market return that has the same risk as keeping the River City branch open.
- *Cash might be tight.* Van Deventer and Imai (2003) discuss how the buy low/sell high transaction-level performance measurement strategy has to be tempered by the financial institution's ability to buy low. Chapter 37 discusses how, as a financial institution's risk rises, the supply of funds to the financial institution can fall or reach zero. At such a high level of risk, the financial institution has lost its real option to buy low/sell high. This is one of the reasons why financial institutions typically find that their optimal credit risk from a shareholder value-added perspective is consistent with default probabilities in a low range, a lesson they learned the hard way in the 2006–2011 credit crisis. We deal with this issue in more detail in the next chapter. This “own default risk” issue is why the buy low/sell high performance measurement test outlined previously is necessary, but not sufficient, to result in the origination of an asset or liability. The financial institution's own default risk impacts its ability to exercise its option to buy low/sell high.

The preceding examples show that, even at the transaction level, portfolio considerations can be important. The decision to issue retail CDs depends on the total volume that can be done (the portfolio of CDs), not just the pricing on one CD. The decision on whether to keep or sell the branch depends on transaction level benefits of the retail CD and small business loan portfolio transactions at the branch. Finally, the decision to buy low/sell high can depend on the entire existing portfolio of assets and liabilities, since they determine the current credit risk of the institution.

PLUS ALPHA BENCHMARK PERFORMANCE VS. TRANSFER PRICING

In Chapter 2, we discussed both the financial accounting–oriented transfer pricing system used by commercial banks and the market value–based benchmark performance system used by almost every other type of financial institution. The authors

believe that commercial banks would benefit from adopting the benchmark-based performance measurement approach, while being careful to avoid the pitfalls of such an approach as outlined in Chapter 2.

WHY DEFAULT RISK IS CRITICAL IN PERFORMANCE MEASUREMENT OF EQUITY PORTFOLIOS

The risk management techniques in common use for equity portfolio management include Nobel prize-winning concepts such as the capital asset-pricing model (CAPM), arbitrage pricing theory (APT), and the efficient frontier concept. The list of researchers behind these concepts reads like a who's who of finance: William Sharpe, Harry Markowitz, John V. Lintner, Jr., Jan Mossin, Steven Ross, and many other key contributors. Like any theory, however, the CAPM, APT, and efficient portfolio concept are based on the same flaw that has doomed simple implementations of value-at-risk (VaR) and the copula approach for CDO valuation: a normal distribution of equity returns is at the heart of CAPM, APT, and the efficient frontier. Default risk is ignored. The current credit crisis makes a more sophisticated and realistic assessment of risk in equity portfolios a mission critical function. We discuss how to do that in this section.

In our discussion of value at risk in Chapter 36, we observed that the assumption that equity returns are normally distributed with their historical mean and standard deviation is so highly stylized that it makes default disappear like magic. In Bear Stearns' case, for example, the historical mean and variance of monthly returns from January 1990 onward implied that the probability of a -100 percent return in a given month was zero to six decimal places. The fact that this -100 percent return did occur is not a "black Swan" à la Nassim Taleb; it is a very predictable result that is a measurable probability for any public firm. The decision to ignore default in measuring risk and return in equities is a simplification that has proven extraordinarily costly to many in the 2006–2011 credit crisis. Still we hear things like this all the time from equity portfolio managers:

Why do I need default probabilities? I manage an equity portfolio, not a fixed income portfolio.

On hearing this, John Y. Campbell of Harvard University laughingly responded, "Don't they know that equity is the most highly subordinated liability on the balance sheet of a public firm?" He's exactly right. A more subtle argument for ignoring default in equity portfolio management is embedded in this comment:

I don't need default probabilities to manage my equity portfolio. My benchmark is the S&P 500 and no firm in the S&P 500 has ever defaulted.

It turns out, thanks to the Standard & Poor's website, that this comment is dramatically untrue. Remember Dana Corporation? Calpine? Delphi? Winn-Dixie? All of these firms failed while they were a component of the S&P 500. There is a more important phenomenon, however, that is very much credit driven that affects equity managers. Consider firms like Washington Mutual, Wachovia, Merrill Lynch,

General Growth Properties, and Lehman Brothers. All of them were thrown out of the S&P 500 as their credit quality deteriorated. These changes in the composition in the index happen “after hours.” The stock price of the firm dropped from the index will be down 10 to 20 percent at the next day’s open, and the stock price of the firm replacing them in the index will be up 10 to 15 percent on the open. Even if the equity portfolio manager is running a perfect replica of the S&P 500, the manager will lose the equivalent of 20 percent (best case) to 35 percent on one of the 500 elements of the index because of this phenomenon. This is a serious “negative alpha” phenomenon. Changes in index composition are similar to what the always insightful Michael Lewis (2004) pointed out in his baseball book *Money Ball*: walking is a skill, not an accident. Similarly, changes in index composition are not accidents either. They are highly predictable, and the default probabilities of the companies in the index have a very high degree of correlation with their probability of being dropped from the index.

Just as important to equity managers is the disconnect between the credit default swap market and equity returns. The CDS market is focused on senior unsecured debt. What if Citigroup is bailed out continuously but not nationalized? Most observers expect that equity holders will lose everything (by being completely diluted) and debt holders will lose nothing. Only default probabilities that predict the probability of the company failing (i.e., the shareholders being wiped out) capture this impact correctly. The CDS quotes for this example, Citi, would be much lower than the correct default probabilities because they reflect anticipation of 100 percent recovery like that which has accompanied the creeping nationalization of AIG.

“PLUS ALPHA” PERFORMANCE MEASUREMENT IN INSURANCE AND BANKING

With the proper adjustments discussed in the prior section, the comparison of actual mark-to-market performance relative to a naïve index of the same risk is a highly useful and very simple concept. The only trick in implementing this approach in the commercial banking and insurance arena is the need to create benchmarks. For example, as van Deventer and Imai (2003) suggest, the head of retail banking who owns the three-year auto loan portfolio should be judged by the market value of a U.S. Treasury index that replicates the exact cash flow timing and amount of each auto loan on a full option-adjusted basis. We show in Chapter 40 that this is a modest computer science effort that is already commercially available and it is merely a question of institutional will to move in this direction.

The manager of the three-year auto loan portfolio is a good performer if over time the mark-to-market value of the auto loans, after expenses and after actual credit losses, is a larger number than the mark-to-market value of the artificial Treasury auto loan index.

The manager of the auto loan portfolio will use transaction-level analytics to decide which loans to originate and which loans to pass on. An important part of this relates directly to the reduced form modeling technology we discussed in Chapters 16 and 17:

- The *default probability* of each auto loan borrower has to be correctly assessed using state-of-the-art logistic regression such as that discussed in the public firm default probability context in Chapter 16.

- The proper *credit spread* for the auto loan has to be assessed in light of expenses of origination and monitoring and in conjunction with the macro factors driving default. In addition, the probability of prepayment has to be taken into account as discussed in Chapters 27 to 29.
- Once breakeven pricing is set for each borrower, the bank can *accept or reject the loan*. Skillful loan selection creates plus alpha performance at the portfolio level by creating value one loan at a time.

This is true buy low/sell high implementation, consistent with state-of-the-art state of the art credit risk and interest rate risk management. A benchmark performance that is plus alpha will show management exactly how much value has been created by this business above and beyond what would have been achieved by a matched maturity investment in U.S. Treasury securities with identical payment dates, payment amounts, and options characteristics. Of course the more traditional transfer pricing methodology used in commercial banking can be used to show this as well, but the transfer pricing yield curve will be the U.S. Treasury yield curve, not the marginal cost of funds yield curve of the bank. In addition, the transfer pricing rate will include a premium that reflects the prepayment option held by the auto loan borrower.

DECOMPOSING THE REASONS FOR PLUS OR MINUS ALPHA IN A FIXED INCOME PORTFOLIO

One of the most interesting set of cultural differences we have come across in finance is the difference between investment managers and bankers. One could write 100 books and 2 million jokes about those differences, but this section has a more modest ambition: to show how to improve fixed income performance attribution in both investment management and in banking by combining best practice from both types of institutions. This section is an introduction to that topic.

We can briefly summarize the key causes of the cultural differences between fixed income investment managers and bankers in a few bullet points:

- *Nature of assets managed*: The overwhelming majority of the assets held by fixed income fund managers is publicly traded. The overwhelming majority of assets held by banks, at least by number of transactions, is not publicly traded.
- *Nature of performance measurement*: In fixed income fund management, performance measurement is based on market valuation and both absolute returns and returns versus a benchmark portfolio. In spite of more than 40 years of debate in banking, that industry remains largely driven by financial accounting–based performance measurement, not market value–based performance measurement.
- *Granularity of performance measurement*: The CFA Institute, formerly known as the Association for Investment Management and Research, has long recommended calculating performance using daily time periods. Outside of the trading floor, banks are largely driven by monthly or quarterly time frames.
- *Volume of transactions*: Most investment managers would have no more than 5,000 or 10,000 positions, and many would have far less. One of the world's five largest banks has more than 700 million distinct assets and liabilities. This creates a much bigger information technology footprint in banking by necessity.

- *Staff count:* Staffing in investment management is lean and mean, because very little retail business is done, at least relative to banking. The largest banks in North America, by contrast, have a staff count in the hundreds of thousands.
- *Analytical rigor:* There are brilliant people in both banking and fixed income investment management who are well beyond best practice in their analysis. That being said, it is much more tempting to manage risk with a spreadsheet or risk analysis that is not much more powerful than a spreadsheet when the portfolio has a relatively small transaction count, something very common in the hedge fund industry in particular.

We recently came across a European paper on fixed income performance attribution that relied heavily on standard yield-to-maturity calculations and duration as the basis for improved fixed income performance attribution. The purpose of this section, in an introductory way, is to explain why it is particularly important in fixed income portfolio attribution to apply much more accurate and modern techniques than the duration concept, because the errors embedded in the yield-to-maturity concept and duration (as we discussed in Chapter 4) are large and can result in wildly inaccurate assessments of fixed income performance attribution:

- Yield-to-maturity discounts cash flows at all payment dates at the same interest rate
- If the same issuer has two bonds outstanding with different yields to maturity, interest payments on the same dates will be discounted at different rates for each bond
- Common spread calculations are usually the simple difference between the yields to maturity on two bonds with similar but not identical maturity dates

Now we can see the benefits of cross-pollination of risk management between bankers and fixed income fund managers:

- We take the *mark-to-market orientation* of fund managers.
- We take the exact day count *matched maturity transfer pricing* concept from banking.
- We take the *yield curve smoothing* concepts that are essential to transfer pricing in banking.
- We take the macroeconomic factor-driven *stress testing approach* from banking and bank regulators that has come out of the 2006–2011 credit crisis.
- We then use this technology to analyze the reasons for mark-to-market *performance changes* in a way already familiar to fixed income fund managers but with greater accuracy.

We know that changes in a large number of macroeconomic factors affect performance of fixed income portfolios. A small sample of them is listed here:

- Level and shape of the risk-free yield curve
- Home prices
- Commercial real estate prices
- Oil prices

- Other commodity prices
- Foreign exchange rates
- Market volatility
- Unemployment rates
- Changes in real gross domestic product
- Government surpluses or deficits

Best practice risk managers are increasingly measuring the impact of each of these factors on mark-to-market performance by tracing their impact on credit spreads, prepayment, and default risk of borrowers from retail to small business to major corporations to sovereigns. In the example of this section, we keep it simple.

- We analyze the performance of one bond
- We assume that we only need to analyze three drivers of performance: (1) changes in the risk-free yield curve and (2) the credit spread for the bond issuer, along with (3) the passage of time.
- We look from the perspective of July 15, 2010, and ask the question, “How much did these risk factor movements contribute to the gain or loss in the value of our bond since April 15, 2010?” We could just as easily do the analysis for a portfolio versus a benchmark. We keep it very simple here for expository purposes.

A more complete analysis has a much longer list of drivers of performance that includes all of the macro factors above plus many more.

A WORKED EXAMPLE OF MODERN FIXED INCOME PERFORMANCE ATTRIBUTION

We now go through a worked example of a more modern approach to fixed income performance attribution with one bond issued by ABC Company and three drivers of performance: the risk-free yield curve, the spread, and the passage of time.

We assume the ABC bond has a par value of 1,000, a 10 percent coupon, semiannual coupon payment dates of June 30 and December 31, and a maturity on June 30, 2019. We observe in the marketplace these prices:

April 15, 2010: \$1532.60 (net present value = price plus accrued interest)

July 15, 2010: \$1580.18

We ask this question: What factors have caused the net present value of the bond to change from \$1532.60 to \$1580.18? How much was due to the passage of time, how much to changes in the risk-free rate, and how much to spread changes?

We answer this question by using the maximum smoothness forward rate approach to yield curve smoothing outlined in Chapter 5. We apply these techniques to the risk-free yield curve and to the ABC bonds outstanding to derive a risk-free and risky yield curve on April 15, 2010, and July 15, 2010. As a substitute for a risky yield curve derived from all ABC bonds outstanding, we assume that, by coincidence, the yield curve for ABC is identical to the U.S. dollar LIBOR-swap curve.

For example, Exhibit 38.1 shows the zero-coupon bond yields for the risk-free and risky (ABC curve or swap curve) on April 15, 2010, as reported in Kamakura's "Friday Forecast."¹

Note that, on this day, the risk-free yield curve's zero-coupon yields were actually higher than the ABC Company/swap curve zero yields at the longer maturities. This is an increasingly common phenomenon that recognizes that the risk-free or sovereign yield curve is not, in fact, risk free.

In Exhibit 38.2, we also take the zero-coupon yields for the risk-free and risky ABC/swap yield curve on July 15, 2010.

Over this three-month period, we can see that the risk-free, zero-coupon yield curve has fallen substantially as shown in Exhibit 38.3.

Risky zero-coupon yields, with the exception of the short maturities, have also moved downward, as this graph in Exhibit 38.4 shows.

The zero-coupon credit spread for ABC company/swap curve has moved as shown in Exhibit 38.5.

Many researchers have found that it is common for credit spreads to rise when risk-free yields fall, and that is consistent with yield curve movements between these two dates. We extract the following zero-coupon bond prices from both yield curves on both dates and apply it for performance attribution to the cash flows on the ABC bond as shown in Exhibit 38.6.




EXHIBIT 38.1 U.S. Treasury and USD Interest Rate Swap Zero-Coupon Yields Derived from the Federal Reserve H15 Statistical Release Using Maximum Smoothness Forward Rate Smoothing, April 15, 2010

Sources: Kamakura Corporation; Federal Reserve.




EXHIBIT 38.2 U.S. Treasury and USD Interest Rate Swap Zero-Coupon Yields Derived from the Federal Reserve H15 Statistical Release Using Maximum Smoothness Forward Rate Smoothing, July 15, 2010




EXHIBIT 38.3 Risk-Free, Zero-Coupon Yields, April 15, 2010, and July 15, 2010




EXHIBIT 38.4 Risky Zero-Coupon Yields, April 15, 2010, and July 15, 2010




EXHIBIT 38.5 Zero-Coupon Spreads, April 15, 2010, and July 15, 2010

EXHIBIT 38.6 Actual Zero-Coupon Yields and Zero-Coupon Credit Spreads

Date	4/15/2010 Risk-Free, Zero- Coupon Yield	4/15/2010 Risky Zero- Coupon Yield	4/15/2010 Zero- Coupon Credit Spread	7/15/2010 Risk-Free Zero- Coupon Yield	7/15/2010 Risky Zero- Coupon Yield	7/15/2010 Zero- Coupon Credit Spread	Change in Risk- Free Zero Yields	Change in Risky Zero Yields	Change in Zero Spreads
20100430	0.163%	0.279%	0.116%						
20100531	0.160%	0.358%	0.198%						
20100630	0.159%	0.394%	0.235%						
20101231	0.310%	0.488%	0.178%	0.193%	0.729%	0.535%	-0.025%	0.280%	0.305%
20110630	0.537%	0.651%	0.114%	0.264%	0.631%	0.367%	-0.147%	0.085%	0.232%
20111231	0.850%	0.974%	0.124%	0.394%	0.641%	0.247%	-0.293%	-0.159%	0.134%
20120630	1.171%	1.321%	0.149%	0.593%	0.850%	0.257%	-0.420%	-0.300%	0.120%
20121231	1.468%	1.633%	0.165%	0.787%	1.048%	0.261%	-0.536%	-0.434%	0.102%
20130630	1.737%	1.910%	0.174%	0.967%	1.217%	0.250%	-0.637%	-0.557%	0.080%
20131231	1.995%	2.175%	0.180%	1.155%	1.407%	0.252%	-0.712%	-0.637%	0.075%
20140630	2.239%	2.420%	0.180%	1.352%	1.611%	0.259%	-0.766%	-0.688%	0.078%
20141231	2.476%	2.648%	0.172%	1.560%	1.813%	0.253%	-0.800%	-0.724%	0.076%
20150630	2.697%	2.848%	0.150%	1.763%	1.995%	0.232%	-0.825%	-0.755%	0.070%
20151231	2.908%	3.025%	0.117%	1.963%	2.163%	0.200%	-0.841%	-0.776%	0.065%
20160630	3.099%	3.177%	0.079%	2.149%	2.312%	0.163%	-0.856%	-0.791%	0.065%
20161231	3.274%	3.312%	0.039%	2.321%	2.448%	0.127%	-0.867%	-0.799%	0.068%
20170630	3.427%	3.431%	0.004%	2.474%	2.568%	0.093%	-0.878%	-0.805%	0.073%
20171231	3.564%	3.539%	-0.024%	2.613%	2.677%	0.065%	-0.885%	-0.809%	0.076%
20180630	3.682%	3.638%	-0.044%	2.733%	2.775%	0.041%	-0.891%	-0.815%	0.077%
20181231	3.788%	3.731%	-0.056%	2.842%	2.864%	0.023%	-0.895%	-0.821%	0.074%
20190630	3.879%	3.817%	-0.062%	2.936%	2.945%	0.009%	-0.898%	-0.829%	0.069%

EXHIBIT 38.7 Scenarios for Fixed Income Performance Measurement

Attribution Driver	Scenario Description				
	Actual 4/15/2010	Actual 7/15/2010	Risk-Free Shift Only	Spread Shift Only	Date Shift Only
Risk-Free Yields	4/15/2010 Risk-free	7/15/2010 Risk-free	7/15/2010 Risk-free	4/15/2010 Risk-free	4/15/2010 Risk-free
Risky Spreads	4/15/2010 Spreads	7/15/2010 Spreads	4/15/2010 Spreads	7/15/2010 Spreads	4/15/2010 Spreads
Cash Flow Dates	4/15/2010 Dates	7/15/2010 Dates	4/15/2010 Dates	4/15/2010 Dates	7/15/2010 Dates
Gross Bond Value	1532.60	1580.18	1623.37	1525.47	1499.56
Less Value of 6/30/ Coupon	−49.96	0.00	−49.96	−49.96	0.00
Net Bond Value	1482.64	1580.18	1573.41	1475.51	1499.56
Dollar Change in Bond Value		97.55	90.78	−7.12	16.92
% Change vs. Actual 4/15/2010	Not Applicable	6.579%	6.123%	−0.480%	1.141%

We want to understand the drivers of changes in bond value so we analyze these specific combinations of changes in the key macro factors that we have chosen to analyze as shown in Exhibit 38.7.

The first two scenarios use the real inputs for April 15 and July 15, 2010. The next three scenarios shift one of our three key drivers of returns, one at a time. The “risk-free shift only” scenario can be explained as follows:

Assume we are at April 15, 2010, and the risk-free curve took July 15, 2010, levels but the credit spread was the April 15 spread.

The “spread shift only” scenario means:

Assume we are at April 15, 2010, and risk-free curve is the April 15 curve, but the spread is the July 15 spread.

In the final scenario, the “date shift only” scenario applies the April 15 risk-free curve and spread on the July 15, 2010, value date, when all cash flows are three months closer than they were on April 15.

In each case, the ABC bond’s present value is the sum of the relevant zero-coupon bond times the cash flow on that particular date. Those zero-coupon bonds are shown in Exhibit 38.8.

EXHIBIT 38.8 Zero-Coupon Bonds

Date	Discount Factors Used		Risk-Free Shift Only	Spread Shift Only	Date Shift Only
	Actual 4/15/2010	Actual 7/15/2010			
4/30/10	0.999885	0.000000	0.999885	0.999858	0.000000
5/31/10	0.999549	0.000000	0.999551	0.999326	0.000000
6/30/10	0.999180	0.000000	0.999196	0.998725	0.000000
12/31/10	0.996530	0.996632	0.997060	0.994357	0.997924
6/30/11	0.992169	0.993972	0.994839	0.990181	0.994783
12/31/11	0.983462	0.990664	0.989529	0.981383	0.988359
6/30/12	0.971222	0.983458	0.981552	0.968789	0.977682
12/31/12	0.956632	0.974488	0.972106	0.954309	0.964121
6/30/13	0.940508	0.964612	0.961174	0.938243	0.948833
12/31/13	0.922381	0.952409	0.948210	0.919769	0.931598
6/30/14	0.903134	0.938171	0.933437	0.900173	0.912958
12/31/14	0.882634	0.922227	0.917207	0.879599	0.892903
6/30/15	0.862090	0.905735	0.900325	0.859106	0.872436
12/31/15	0.841234	0.888488	0.883094	0.838142	0.851572
6/30/16	0.820835	0.871166	0.865981	0.817459	0.830999
12/31/16	0.800504	0.853554	0.848926	0.796693	0.810561
6/30/17	0.780775	0.836253	0.832022	0.776583	0.790663
12/31/17	0.760962	0.818773	0.815034	0.756466	0.770766
6/30/18	0.741694	0.801735	0.798145	0.737103	0.751354
12/31/18	0.722329	0.784601	0.781114	0.717813	0.731909
6/30/19	0.703509	0.767940	0.764216	0.699266	0.712941

When we apply these scenarios, we get the components of performance shown in Exhibit 38.9.

Note that we first deduct from the actual April 15, 2010, gross bond value or net present value the present value of the June 30, 2010, coupon payment of \$50. This adjustment is necessary to insure we are comparing the value of two securities (ABC bond on April 15 and July 15) with the same attributes. This adjustment would not be necessary if we were using the daily return intervals that the CFA Institute recommends, since the coupon payment would just be an element of daily total return. In this case, we deduct the present value of that first coupon rather than speculate on reinvestment returns from June 30 to July 15.

The analysis shows that, after adjusting for the first coupon, our ABC bond returned a total of 6.579 percent from April 15, 2010, to July 15, 2010. The downward shift in the risk-free curve contributed 6.123 percent to this total return. The widening of credit spreads by itself reduced return by 0.480 percent. The passage of time, which moved all cash flows three months closer, contributed 1.141 percent to total return. A final component, the impact of all factors moving together, reduced total return by 0.204 percent.

EXHIBIT 38.9 Scenario Components

Attribution Driver	Scenario Description					Impact of All Factors' Joint Movement	Total
	Actual 4/15/2010	Actual 7/15/2010	Risk-Free Shift Only	Spread Shift Only	Date Shift Only		
Risk-Free Yields	4/15/2010 Risk-free	7/15/2010 Risk-free	7/15/2010 Risk-free	4/15/2010 Risk-free	4/15/2010 Risk-free		
Risky Spreads	4/15/2010 Spreads	7/15/2010 Spreads	4/15/2010 Spreads	7/15/2010 Spreads	4/15/2010 Spreads		
Cash Flow Dates	4/15/2010 Dates	7/15/2010 Dates	4/15/2010 Dates	4/15/2010 Dates	7/15/2010 Dates		
Gross Bond Value	1532.60	1580.18	1623.37	1525.47	1499.56		
Less Value of 6/30/2010 Coupon	-49.96	0.00	-49.96	-49.96	0.00		
Net Bond Value	1482.64	1580.18	1573.41	1475.51	1499.56		
Dollar Change in Bond Value		97.55	90.78	-7.12	16.92		
% Change vs. Actual 4/15/2010	Not Applicable	6.579%	6.123%	-0.480%	1.141%	-0.204%	6.579%

In this fashion, we have a complete and accurate understanding of the contribution of the risk factors we have chosen to analyze to total absolute return. Return relative to a benchmark index, or tracking error, is analyzed in the same way. We have avoided the serious performance attribution errors that can arise from the simplifying assumptions embedded in yield-to-maturity and simple credit spread calculations by using modern yield curve and credit spread smoothing on an exact day count basis. The process is simple and straightforward when using modern enterprise risk management software.

THE JARROW-MERTON PUT OPTION AND CAPITAL

We now turn back to more traditional measures of performance, using financial accounting as a basis, which have a long tradition in banking and (more recently) in insurance. In Chapters 36 and 37, we discussed the practical calculation of the Jarrow-Merton put option that is the comprehensive measure of integrated credit risk, market risk, liquidity risk, and interest rate risk that we first discussed in Chapter 2. In Chapters 19 through 35, we illustrated how we can value the assets and liabilities of any financial institution using the simulation techniques of Chapter 36. We can use this calculation as a basis for capital allocation to the various business units in a financial institution just as we can for the institution as a whole.

USING THE JARROW-MERTON PUT OPTION FOR CAPITAL ALLOCATION

Introduction

As we noted in Chapter 2, most of the business units of major financial institutions have either assets only or liabilities only. They rarely have both assets and liabilities at the same time. When this is the case, there is no need to allocate capital because the benchmark performance system we discuss above can be used. Nonetheless, many financial institutions allocate capital to achieve multiple objectives:

- They want all business units, after the allocation of debt and equity, to have the same relative degree of risk.
- They then feel comfortable, using either historical market returns or financial accounting returns, in expressing performance as the ratio of historical return to allocated capital.

As we noted in Chapter 2, this approach implicitly says that “good” business units have good returns on risk-adjusted capital and should get more emphasis. This portfolio measure of performance is thought of by many as a substitute for the buy low/sell high approach discussed in the introduction to this chapter.

Using the Jarrow-Merton Put Option Concept for Capital Allocation

In this section, we show how to use the Jarrow-Merton put option technology for capital allocation. We start with a one-period example and then generalize to a

multiperiod example that is consistent with the Federal Deposit Insurance Corporation (FDIC) Loss Distribution Model of Jarrow and colleagues (2003).²

Capital Allocation for a Single Period In Chapter 36, we outlined a careful process for simulating the impact of macroeconomic factors on the default probabilities and legal payment obligations of all counterparties from retail borrowers to sovereigns. We will use the identical process in this section. If we set a single time horizon as relevant for capital allocation, the process works like this:

1. Choose the time horizon T (one year is a common choice).
2. Obtain from market sources the default probability for the institution. This default probability can be derived in any method with which the institution is comfortable, as discussed in Chapters 16 and 17. Obtain the business unit's complete transaction-level listing of assets and liabilities.
3. Simulate the macro factors over N scenarios as of time T .
4. Calculate the default probabilities for all counterparties.
5. Calculate N values for the value of business unit assets less the amount of liabilities.
6. Assume the default probability for the business unit with no Jarrow-Merton put option is d^* . We need to solve for the value of the Jarrow-Merton put option, which reduces d^* to the institution's own default probability d .
7. Examine the distribution of negative net worths from the simulation.
8. Determine the amount K of a digital put option payable on default such that the percentage of negative net worth scenarios after payment of K is reduced from d^* to d . For example, if d is 1 percent, but 200 of the 10,000 scenarios produce negative net worth on a mark-to-market basis ($d^* = 2$ percent), we need to find K such that 100 of the negative scenarios would be erased. If the one-hundredth largest loss is \$72.27, then the digital option would pay \$72.27 on default, reducing d^* to 1 percent = d .
9. Calculate the value of the Jarrow-Merton put option, which pays K (\$72.27) upon default (i.e., negative mark to market before K is paid). Let us call the value of this put option X . This is the price of portfolio insurance for a default probability of d .
10. Adjust the capital and the assets of the business unit for the value of the put option. Before the put option, the capital of the business unit was $V - B$ and its capital ratio was $(V - B)/V$. After the put option, capital will be increased by X and assets will be increased by X . Assets will be $V + X$ and capital will be $V - B + X$. The ratio of capital to assets will be $(V - B + X)/(V + X)$.
11. This process equalizes the default probabilities of every business unit at the institutional level and explicitly incorporates the risk of each individual asset or liability on the books of the business unit.

There are other choices for the terms of the payoff on the Jarrow-Merton put option. The most obvious is one that equalizes the expected loss of each business unit. Even if default probabilities are equalized, the expected loss can differ and vice versa. This is a matter for institutional strategy.

In any event, the technology for the Jarrow-Merton put option in the first 37 chapters of this book allows us to explicitly price the value of portfolio insurance (both

credit risk and interest rate risk) on all assets and liabilities of each business unit so that the total default risk of each business unit after portfolio insurance is equalized.

Extending the Jarrow-Merton Capital Allocation to a Multiperiod Framework

While most financial institutions allocate capital with a one-year time horizon in mind, the 2006–2011 credit crisis makes it clear that a longer time horizon may be more appropriate. The fact that the credit default swap market has crystallized around a five-year maturity is indirect confirmation that the market believes that a longer time horizon has the most value from a credit insurance policy point of view.

For this reason, it may be appropriate to analyze the Jarrow-Merton put option as a multiperiod payout option that insures the survival over a longer time horizon just like the Loss Distribution Model of the FDIC does.

This changes the valuation of the Jarrow-Merton put option from the previous section in only modest ways:

- Scenario simulation has to deal with reinvestment of cash flows over multiple periods for correct assessment of the probability of business unit default.
- The default of a business unit can happen at multiple points in time, and multiple payments to the business unit may be necessary to insure survival over the entire time horizon. This is highly likely given the high correlation in defaults that we have discussed throughout this book.
- This will raise the cost of the Jarrow-Merton put option for the business unit, all other things being equal.

SUMMING UP

As noted in Chapter 2, backward looking measures of asset performance have minimal use as predictors of future performance. For that reason, we prefer forward looking measures of performance (buy low/sell high) for asset selection at both the transaction level and the portfolio level. Nonetheless, the Jarrow-Merton put option concept can be applied neatly for capital allocation for those institutions that have the need to apply the capital allocation concept. We explicitly price the portfolio insurance (integrating credit risk, market risk, liquidity risk, and interest rate risk) for each business unit such that we equalize the (1) probability of default or (2) expected loss for each business unit. The adjusted asset totals and capital totals for each business unit reflect the incremental cost of this credit insurance, the value of the Jarrow-Merton put option.

We turn next to a related topic, the safety and soundness of our own financial institution.

NOTES

1. Kamakura's Friday Forecasts are available on www.kamakuraco.com.
2. The Loss Distribution Model can be found at the FDIC's website, www.fdic.gov.

Managing Institutional Default Risk and Safety and Soundness

In 2003, the Basel Committee on Banking Supervision proposed the New Capital Accords as a way to insure the *safety and soundness* of financial institutions around the world. Since then, such a volume of documents has poured forth from the Basel Committee that few risk managers known to the authors admit to having read every page of even one Basel Committee document. Insiders at the Bank for International Settlements (BIS) have revealed an indirect confirmation of why the Basel II and III regulations have evoked such a yawn from real risk managers: inside the BIS, as of this writing, there is not a single piece of evidence that the Basel ratios (in any variation) are statistically significant in predicting financial institution failure. For that reason, we focus on methods for controlling default risk in this chapter in which both the statistical evidence and economic logic are strongly supportive of the risk “levers” we discuss.

In Chapters 36 through 38, we showed how modern risk management technology allows a much different approach to ensuring the safety and soundness of financial institutions. In this chapter, we discuss the implications of a more modern approach to risk management for the safety and soundness of a given institution from both a shareholder value perspective and from a regulatory perspective.

STEP 1: ADMITTING THE POSSIBILITY OF FAILURE

Step 1 in using a more modern approach to risk management is for management to admit the unthinkable—that the institution they manage has a probability of failure that is not zero. It is a rare management team that can make this admission, but a failure to acknowledge this fact is to deny reality and doom the risk management exercise at the institution to common practice at best. Even up through 1996, the authors were often told in the Tokyo market that “no Japanese bank has failed since the end of World War II,” a phrase intended to convey that the future was riskless as well. Within a few years, seven of the top 21 banks in Japan had failed. Even after the savings and loan crisis in the late 1980s and early 1990s in the United States, banking CEOs were in denial about the risk of their own firms. One of the authors was explicitly forbidden to discuss the bank’s own default risk by the CEO himself in the mid-1980s. After the author left the bank, it failed because of an overexposure to

commercial real estate loans. Prior to the credit crisis of 2006 to 2011, one major bank with more than \$500 billion in assets labeled its default database for the U.S. securities industry a “low default portfolio” that was so hard to model they used judgmental default probabilities of 2 to 3 basis points for U.S. securities firms, less than three years prior to the failure of Lehman Brothers and Bear Stearns in 2008 and the rescues of Merrill Lynch and Morgan Stanley by the U.S. government.

With the 2006–2011 credit crisis squarely in the rearview mirror, few objective analysts have to be convinced that the probability of financial institutions’ failure is not zero; even at the CEO level, those who proclaim their firms have zero risk are often subjected to shareholder lawsuits for inaccurate disclosure. Exhibit 39.1, provided by Kamakura Risk Information Services, shows that the one-year Kamakura Default Probabilities (KDP) for Bank of America (BAC) (lower line) and Citigroup (C) (upper line) peaked at about 20 and 40 percent respectively in March 2009, five months after the U.S. government injected \$45 billion into the merged Merrill Lynch and Bank of America and \$25 billion into Citigroup. Clearly, major financial institution default risk is not zero.

Exhibit 39.2 shows one-year default probabilities over the same period for the Royal Bank of Scotland Group (RBS) (upper line) and HBOS (lower line, until its rescue/merger early in 2009). The evidence is similar: even the largest and most prestigious financial institutions have not been too big to fail. In the recent crisis, they have simply been “too big not to bail out when they fail.”

Both graphs make it clear that default probabilities in troubled times can show a very high degree of correlation.




EXHIBIT 39.1 Bank of America (JC5 model) vs. Citigroup (JC5 model)




EXHIBIT 39.2 Royal Bank of Scotland Group (JC5 model) vs. HBOS PLC (JC5 model)

MANAGING THE PROBABILITY OF FAILURE

Are Ratings a Useful Guide?

In this section, we ask the question “Can an institution’s own credit rating be a proper risk index to be managed by the firm?” As we discussed at length in Chapter 18, one of the hardest aspects of using credit ratings for such a role is that ratings have no specific maturity and no specific default probability associated with them. Indeed, there has been a long debate in the risk management industry about whether traditional credit ratings are intended to be “point-in-time” ratings or whether they are intended to be “through-the-cycle” credit indicators. The reason that such a debate occurs is the vagueness of the maturity and default probabilities associated with a given rating for a specific company at any point in time. When viewed in the context of a modern quantitative default probability model, the debate about point in time versus through the cycle is a distinction without a difference.

We posed this question in Chapter 18: How does a modern quantitative default model differ from ratings? The differences are very great and very attractive from a practice use point of view:

- Each default probability has an explicit maturity.
- Each default probability has an obvious meaning. A default probability of 4.25 percent for a three-year maturity means what it says: there is a 4.25 percent (annualized) probability of default by this company over the three-year period starting today
- Each company has a full-term structure of default probabilities at maturities from one month to 10 years, updated daily.

What does “point in time” mean in this context? All of the default probabilities for Morgan Stanley, for example, prevailing on November 28, 2008 (about six weeks after the failure of Lehman Brothers) are default probabilities at that are point in time for Morgan Stanley.

Default probabilities at a different point in time, say January 25, 2009, will change if the inputs to the default models are different on January 25 than they were on November 28. What does “through the cycle” mean with respect to the default probabilities for Morgan Stanley on November 28, 2008? Through the cycle implies the longest default probability available on the term structure of default probabilities because this maturity does the best job of extending through as much of the business cycle as possible. For Morgan Stanley, the longest default probability available is the 10-year default probability, 0.83 percent (annualized basis). This is quite different from the high risk Morgan Stanley faced in the short run in the aftermath of the Lehman failure, because Morgan Stanley’s annualized one-month default probability at the same time was 8.96 percent. If the default probability for Morgan Stanley on November 28 is 0.83 percent at 10 years, it means that the through-the-cycle default probability for Morgan Stanley prevailing on November 28 is a 0.83 percent (annualized) default rate over the 10 years ending November 28, 2018. This is explicit and quite a contrast to the vagueness of the ratings concept.

To summarize, all of the default probabilities prevailing for Morgan Stanley on November 28, 2008, are the point-in-time default probabilities for Morgan Stanley at all maturities from one month to 10 years. The through-the-cycle default probability

for Morgan Stanley on November 28 is the longest maturity default probability, the 10-year default probability. The 10-year default probability is also obviously a point-in-time default probability because it prevails on November 28, 2008, the point in time we care about. On the next business day, all of the point-in-time default probabilities for Morgan Stanley will be updated, including the 10-year default probability, which has a dual role as the through-the-cycle default probability.

There is no uncertainty about these concepts: All default probabilities that exist today at all maturities are point-in-time default probabilities for Morgan Stanley, and the longest maturity default probability is also the through-the-cycle default probability.

How can these default probabilities be mapped to ratings, and what rating would be point in time and what rating would be through the cycle? An experienced user of quantitative default probabilities would ask in return “Why would you want to go from an explicit credit risk assessment with daily updates and with a known maturity and 10,000 grades (from 0 basis points to 10,000 basis points) to a vague credit assessment with no known maturity, only 20 grades, and updates that occur randomly at about one-year intervals?”

For these reasons, ratings are not a best practice risk lever for managing an institution’s own default risk.

Are CDS Spreads a Useful Guide?

As we discussed in Chapter 20, it was once thought that a competitive credit default swap market with a strong depth of trading would emerge to provide new market-based measures of the cost of default insurance, an alternative to the put option prices we used for Citigroup in Chapter 36. Enthusiasm for the CDS market was so high, in fact, that an academic, in the *Financial Times* of December 28, 2011, wrote, “How can the FDIC determine the correct price for its insurance? The answer is to use the large and active market in bank insurance via credit default swaps.”¹

In Chapter 20, we summarized the reasons why the credit default swap market cannot be used as suggested and why it is not a proper index of “own default risk” that should be an objective of risk management.

1. There is no large and active market in bank credit default swaps. Using data reported by Depository Trust & Clearing Corporation, during the 77 weeks ending December 30, 2011, credit default swaps were traded on only 13 reference names among U.S. banking firms
 - Ally Financial Inc.
 - American Express Company
 - Bank of America Corporation
 - Capital One Bank (USA), National Association
 - Capital One Financial Corporation
 - Citigroup Inc.
 - Citigroup Japan Holdings Corp.
 - Istar Financial Inc.
 - JPMorgan Chase & Co.
 - Metlife, Inc.
 - Morgan Stanley
 - The Goldman Sachs Group, Inc.
 - Wells Fargo & Company

2. Even on the firms listed above, there was an average of only 2.3 nondealer credit default swap trades per day during the 77 weeks ended December 30, 2011.
3. Six of the 11 firms listed are in a conflict of interest position as major dealers in the credit default swap market: Bank of America, Citigroup, JPMorgan Chase, Morgan Stanley, Goldman Sachs, and Wells Fargo. Dealer–dealer trades made up 81.68 percent of live trades in the DTCC over the 77-week period studied. The dealers would be setting deposit insurance rates for themselves if the proposal in the *Financial Times* quoted previously was adopted.
4. Credit default swaps, to the extent they trade, represent supply, demand, probability of default, and the probability of rescue, not the probability of failure alone.
5. Because of the thinness of nondealer trading and the very small number of dealing firms worldwide, there is high risk of collusion.

For these reasons, the proposal to use credit default swaps to price deposit insurance would not work given current market conditions and the lack of competition in the market for credit default swaps. For the same reasons, the CDS pricing on a given financial institution is not a proper risk management index that should be a focus of management. There are many superior tools available, and we turn to them now.

Using Quantitative Default Probabilities

Earlier in this book, we discussed the concept of the safety zone for interest rate risk, those interest rate risk positions where the probability of failure is zero. When we ignore credit risk, match funding all assets except those funded by equity leaves a financial institution free to invest the equity in securities of any maturity without leaving the safety zone. Once we allow for default of the financial institution's borrowers, there is no safety zone for any financial institution that borrows money to make loans. That means management has to deal practically with a nonzero probability of failure. Just as importantly, we should note that the probability of failure has a term structure just like interest rates, as we showed for Morgan Stanley (MS) in Exhibit 39.3, which shows the default probabilities out to 10 years.

In Chapter 36, we listed four primary questions that management and Boards of Directors should be able to answer in a simple and powerful “pass–fail test.” We also listed a supplemental list of 26 questions that add value to the risk management process. From that list of 30 questions, these questions very specifically focus on credit risk:

From the four-question pass/fail test:

- Question 2: Using an insider's knowledge of the assets and liabilities of the firm, both “on balance sheet” and “off balance sheet,” what is the best estimate, monthly for the next 10 years, of the probability that the firm will fail in each of these 120 monthly periods?
- Question 3: Using only information available to an outsider, what is the best estimate of the probability of the failure of the firm in both the short run and the long run?
- Question 4: If the firm is able to answer Questions 1, 2, and 3, what hedging position is necessary to insure that the macroeconomic factor sensitivity of the firm and default probability of the firm reach the target levels set by the Board of Directors?

From the list of 26 supplemental risk management questions:

- How do regulatory capital ratios like the Basel II risk measures vary over the business cycle and to what extent are they correlated with the firm's actual risk of default?
- What 10 macro factors represent the highest threat to the safety and soundness of the institution, from highest to lowest?
- What risk limits should the firm have in place with respect to these 10 macro factors?
- Looking backward in time, which macro factors had the largest historical impact on the firm's market capitalization and default risk?
- Is the firm's self-assessment of its own risk level consistent with the pricing on put options on its common stock and credit default swaps?
- If the answer is no, what is the reason for the inconsistencies?
- If the market prices of puts and CDS are overly pessimistic, what investor education campaign is necessary to more effectively communicate the true risk of the institution?

Note that, in the latter three questions, we have implicitly assumed that there is actual trading in the credit default swaps of the firm in question. CDS prices based on quotes, not traded prices, are unusable for all of the reasons listed in the previous section, and even traded prices have severe problems as we discussed above.

Management needs to set a target for the amount of failure risk that creates the maximum amount of shareholder value. For most financial institutions, this means targeting a one-year default probability in the 10 to 35 basis point range, a range like that of Goldman Sachs during the 2006 to 2011 credit crisis. This default probability level will ensure that the firm has continuous access to capital markets and the confidence of its policy holders (in insurance) and depositors (in banking). Unlike ratings, as we see in Exhibit 39.3, default probabilities for all maturities rise and fall during the course of the business cycle. Exhibit 39.4 shows the one-year (generally the lower line) and five-year (upper line) default probabilities for Wells Fargo (WFC) from January 1990 through December 2006, and Exhibit 39.5 shows the Wells Fargo credit crisis experience from January 2007 through August 2012.

Both the five-year default probability and the one-year default probability rise considerably during the worst part of the credit cycle, but the one-year default




EXHIBIT 39.3 Term Structure of Annualized Default Probabilities for Morgan Stanley (JC5 model), November 28, 2008




EXHIBIT 39.4 Wells Fargo & Company (JC5 model) One-Year and Five-Year Default Probabilities, 1990–2006




EXHIBIT 39.5 Wells Fargo & Company (JC5 model) One-Year and Five-Year Default Probabilities, 2007 to 2012

probability responds more sharply. This is identical to the behavior of one-year and five-year interest rates during the business cycle, and the analogy is strong.

CONTROLLING THE PROBABILITY OF FAILURE THROUGH THE CREDIT CYCLE

What do we do next once management recognizes that its own credit quality varies over the credit cycle because of the cyclical nature of the creditworthiness of its borrowers? A parallel question is this: “What do we do next once management recognizes that the cost of risk protection via the Jarrow-Merton put option prices (see Chapter 36) varies over the credit cycle because of the cyclical nature of the creditworthiness of its borrowers?” The next step in best practice enterprise risk management is to put a number of measurement and monitoring technologies to work regarding the credit quality of the financial institution itself. There is no reason for this not to be a daily discipline. A best practice report would include the following:

- Current reduced form default probabilities for the company, as discussed in Chapter 16.
- Current put option prices, expressed as a percent of current stock price, at all strike prices and maturities available, as discussed in Chapter 36.

- Current credit default swap quotes on the company (if and only if the bank is one of the rare institutions on which trading has occurred); quotes with no trading volume are worthless.
- Current primary (new issue) bond credit spreads for the company.
- Current secondary market bond credit spreads for the company.

For the most technical layer of senior management, the reporting should also include:

- The internally generated time series of monthly default probabilities that stem from a macro factor-driven Monte Carlo simulation as discussed in Chapter 36.
- The internally generated risk-neutral cost of the Jarrow-Merton put option that would eliminate integrated credit risk, market risk, foreign exchange risk, and interest rate risk, as discussed in Chapter 36.
- The internally generated risk-neutral cost of the line of credit that would eliminate liquidity risk, as discussed in Chapter 37.
- The “delta” of shareholder value and the Jarrow-Merton put option with respect to changes in:
 - Home prices
 - Commercial real estate prices
 - Stock price indices
 - Interest rates
 - Oil prices
 - Commodity prices
 - Foreign exchange rates
 - Other relevant macro factors

This kind of reporting discipline provides a framework for the practical action that is the ultimate measure of a risk management process—what’s the hedge?

HEDGING TOTAL RISK TO MAXIMIZE SHAREHOLDER VALUE

Once management has set a target range for the risk of failure and has measured the financial institution’s sensitivities to the macro factors that cause risk, management has three complementary choices for managing default risk/put option levels:

- Management can directly hedge the integrated risk caused by the macro factors using exchange traded derivatives and selected over-the-counter derivatives. Note: The disclosures of JPMorgan Chase hedging losses beginning in May 2012 are a classic example of what *not* to do. These revelations make it clear that the JPMorgan Chase positions were speculation, not a hedge, and that their overwhelming reliance on over-the-counter credit default index swaps left JPMorgan Chase subject to the kind of market manipulation that we warned of in earlier sections of this chapter and in Chapter 20.

- Management can adopt a more conservative liability structure to reduce the risk of failure.
- Management can issue (or buy back) common stock to change the risk of failure.

Since we know the deltas of default risk (see the logistic formula in Chapter 16), shareholder value, and the Jarrow-Merton put option with respect to the key macro factors driving risk, we can directly calculate how much hedging we need to do to put the risk of failure in the target range. Using the technology outlined in Chapter 36, we can show the high degree of correlation between the hedging instruments and the cash instruments being hedged, as required by Financial Accounting Standard (FAS) 133 in the U.S. and by International Accounting Standard (IAS) 39.

Liability structure can also be managed as long as management takes action early enough when credit quality starts to deteriorate. The safest incremental strategy is to issue more equity and use it to pay down short-term liabilities. This buys time to correct any problems the institution may have over the remainder of the business cycle. If all liabilities were paid down, obviously, the financial institution would find it almost impossible to fail.

In any event, management knows when it needs to take action and it knows which actions to take because of the risk management analysis outlined in Chapters 36 and 37. This knowledge and action orientation is good for shareholders and good for regulators.

IMPLICATIONS FOR BASEL II, BASEL III, AND SOLVENCY II

As van Deventer and Imai (2003) note, the New Basel Capital Accords were intended to improve the safety and soundness of financial institutions. The 2006–2011 credit crisis makes it clear that this objective was not achieved. These regulations were well intentioned and constrained by the lowest common denominator–level risk analytics that are typical of government regulations. The Basel II Accords, for example, do not specifically provide for, or recommend, Monte Carlo simulation of risks in spite of the fact that this can be done at minimal expense even in common spreadsheet software.

The Basel Accords suffer in another couple of dimensions when compared to the procedures that we have outlined in the first 38 chapters of this book:

- *Omitted variables.* The Basel Accords capital ratios effectively are a credit model that omits key explanatory variables that have been proven to be statistically important in predicting default: accounting ratios, common stock prices, macroeconomic factors, and so on.
- *No valuation framework.* The Basel Accords do not provide a consistent valuation framework that is essential to marking-to-market our current position and in simulating our potential default in the future.
- *No hedge.* The Basel Accords fail the most basic test of a risk management approach—they do not tell us the hedge we need to solve the problem.

For these reasons, the Basel Accords have to be regarded as a modest supplement to the tools of Chapters 1 through 38, not the core of a sophisticated risk management effort.

SIMULATING YOUR OWN PROBABILITY OF DEFAULT

The greatest advantage in analyzing the risk of an institution is a complete knowledge of all assets and liabilities that it owns now and a very good knowledge of what kind of assets and liabilities it is likely to own in the future. In that sense, you have an information advantage that Jarrow and Protter (2004) have noted is one of the key reasons why the reduced form modeling approach is most consistent with the behavior of market participants who do not have the same information you do.

By better understanding your own risk (Chapter 36) and the timing of a credit-related liquidity event (Chapter 37), which would reveal problems, you can take action promptly. At worst, you can take action as the market comes to similar realizations, and at best, you can avoid a problem and prevent it from becoming a public issue.

If you do not do this, you risk the needless destruction of your financial institution's franchise that has taken years or decades to build. The shareholders of Northern Rock PLC, Lehman Brothers, Bear Stearns, FNMA (Fannie Mae), and FHLMC (Freddie Mac) learned this lesson the hard way.

We turn back to this issue in Chapter 41, but first we turn to the information technology practices that are most likely to result in a timely, efficient, and cost-effective implementation of this risk management technology.

NOTE

1. www.ft.com/intl/cms/s/0/0a4cba22-1764-11e1-b00e-00144feabdc0.html#axzz296l46OQs.

Information Technology Considerations

In the first 39 chapters of this book, we reviewed the tools necessary to construct a fully integrated measure of interest rate risk, market risk, liquidity risk, foreign exchange risk, and credit risk. This measure, the Jarrow-Merton put option, produces an answer to the key question “What is the hedge?” Some of the tools we have discussed were focused on risk management strategy, such as usage of the Jarrow-Merton put option as a comprehensive measure of risk. Other tools have been more mechanical, such as how to value an American call option embedded in a bond when the default of the issuer is a possibility. In this chapter, we focus on another set of issues that have to be faced when implementing a comprehensive risk management system—the information technology aspects of an implementation.

The authors collectively have spent 105 years in the implementation of risk systems in major financial institutions in 23 countries around the world, and our comments below reflect the worldwide best practice in that regard.

COMMON PRACTICE IN RISK MANAGEMENT SYSTEMS: DEALING WITH LEGACY SYSTEMS

Many of the world’s financial institutions face a common situation when preparing to move forward in the risk management systems area. The move forward is usually not motivated by changes in regulatory risk measures, like the Federal Reserve’s Comprehensive Capital Analysis and Review 2012 and similar stress tests introduced in Europe. Particularly in countries where there has been a recent financial crisis, managers of the most sophisticated financial institutions now clearly realize that the common situation described in the bullet points below has left the institution unprepared for a crisis. These managers are rededicated to the comprehensive view of risk that we have described in this book, because their previous failure to take this approach nearly resulted in the failure of their institution. In a few cases, the institution did fail (as we saw in Chapter 37), but government capital injections preserved its role in the financial markets.

A common situation that senior management often faces looks like this:

- The risk management function is fragmented. Market risk, credit risk, liquidity risk, interest rate risk, and performance measurement are all handled by different groups with no central coordination at either the working level or the senior management level.

- There is at best only a partially completed enterprise-wide data warehouse.
- Risk management systems have been implemented on a compartmentalized basis:
 - Vendor A provides third-party default probabilities.
 - Vendor B provides credit risk loss simulation software that provides cumulative losses over a long, single period but no valuation, multiperiod simulation, or hedging.
 - Vendor C provides interest rate risk management simulation and market-to-market calculations, generally ignoring default risk, and using primitive one-factor interest rate models.
 - Vendor D provides traditional value-at-risk calculations for market risk.
 - Vendor E provides performance measurement for each business unit, either on a benchmark basis or on a transfer-pricing basis, depending on the nature of the financial institution.
 - Vendor F (often an in-house solution) provides capital allocation as a supplement to Vendor E's performance measurement data.
- These systems use different graphic user interfaces, different input database design, different financial mathematics, different reporting tools, different end-user staffs and are totally irreconcilable.

This common situation is totally wasteful, as 95 percent of the lines of source code in the six systems should be common. In effect, the user has bought the same system six times and then completely scrambled the picture even more with inconsistent data architecture and reporting. As wasteful as this situation is, for many years it was not the fault of software users. Legacy risk vendors as a group were incapable of producing risk systems that could analyze multiple risk silos on an accurate, consistent, and transparent basis. Over the past 15 years, however, that situation has dramatically changed.

Even more than the wastefulness of this typical situation is its danger—the systems infrastructure obscures the view of the organization's risk and return and defeats the very objectives of the risk management process that we outlined in Chapter 1. Most importantly, there is no link at all between random movements in macroeconomic factors, which drive defaults and valuation for every item on the balance sheet of the financial institution. How can there be an integrated view of risk when each system is looking at a different piece of the puzzle? This is perhaps the biggest reason for the failures and near failures due to home price collapses of firms like Bank of America, Citigroup, Royal Bank of Scotland, HBOS, Northern Rock PLC, Wachovia, Lehman Brothers, Bear Stearns, Washington Mutual, Federal National Mortgage Association (Fannie Mae), Federal Home Loan Mortgage Corporation (Freddie Mac), and many others during the 2006–2011 credit crisis. A similarly long list of failures stems from the macroeconomic factor influences on the Japanese banking and life insurance system after the burst of the Japanese bubble beginning in 1990.

With a typical siloed risk system infrastructure, the problems are too numerous to list them all. Here is just one critical example. A rise in interest rates will never increase default probabilities and change credit-adjusted valuation in an integrated way in a fragmented risk architecture. The default probability vendor ignores the impact of interest rates on default probabilities and does not take inputs from

simulations in other systems. The interest rate risk system ignores default. The financial institution is completely blind with respect to simulations and stress tests of rising interest rates and their impact on default.

This architecture has extremely serious flaws that are of grave concern to many. If senior management has ended up in this situation prior to the developments in risk technology that we have outlined in this book, how can we escape this untenable situation?

It is to that task that we now turn.

UPGRADING THE RISK INFRASTRUCTURE: THE REQUEST FOR PROPOSAL PROCESS

The first step in the process of reforming the risk infrastructure starts with a request for proposal process. In a perfect world, there would be a single risk management staff group that manages market risk, credit risk, interest rate risk, liquidity risk, and performance measurement on a fully integrated basis. There are many fine books on the subject of the organization and execution of enterprise-wide risk management, so we won't detour in that direction here. We need to remain focused on the more micro tasks of execution. Even if the staff of these groups remains separate with a different managerial reporting hierarchy for each group, a joint project team should be organized to design and implement the request for proposal process. If the team is not a joint team, each group will fight to preserve its right to choose the vendor, rather than fighting to choose the best integrated risk vendor. Senior management intervention to fix this is often inevitable.

In order for the request for proposal process to clearly reveal the skills, strengths, and liabilities of the potential vendors, it has to be both comprehensive and very detailed. This means that the project team members have to be both relatively senior and with considerable hands-on experience in each of the areas of risk, we want to cover—market risk, credit risk, interest rate risk, liquidity risk, and performance measurement. A typical best practice request for proposal covers all four areas with equal emphasis, and has 200 to 1,000 very specific questions on the product functionality along the lines of the approach outlined in this book.

Prior to 1995, it was uncommon to see such integrated risk requests for proposals because, frankly, the technology was still in its early stages. Now it is obvious to most market participants that the technology outlined in this book exists and has been implemented by many of the world's most successful financial institutions. For that reason, the authors estimate that 80 percent of the requests for proposal for risk management systems reflect this integrated approach to risk. The only area that seems to be lagging is the large American banks where the legacy of many mergers has left high walls between risk management groups of the merged banks. Even in the United States, however, the trend toward integrated risk management is very strong and distinct. Legacy risk groups, however, often resist being merged to improve risk management with a tenacity that demonstrates how unfit they are for a senior role in risk management.

One of the authors learned to his chagrin when he resigned from Lehman Brothers that the only profession in the world with a lower perceived integrity than investment bankers is that of software salesman. This is a sad testimony to the

industry and one issue that the Request for Proposal (RFP) process has to be explicitly designed to deal with.

The level of detail in the question of the RFP should be high for two reasons. First it will test the comprehensiveness of both the vendors' knowledge and completeness of the vendors' solutions. Second it will allow for a detailed audit of the vendors' solutions after the RFP responses have been gathered. There is a third benefit to the RFP process. It is common for working-level risk staff to prefer a vendor who does not have the best risk system. One reason is prior use of a legacy risk system, allowing the risk analyst to avoid getting better. Another reason, candidly, is bribery in spite of the increasing enforcement of the Foreign Corrupt Practices Act via fines and jail terms. A detailed RFP and its response make it very clear when the inferior vendor has been selected and will immediately trigger an investigation in a sophisticated firm.

We believe that best practice in the RFP process is for short-listed vendors (typically two or three firms) to do more than just demonstrate the product. Even a half-day demonstration can be scripted by a vendor in such a way that deficiencies in the product will remain out of view of the potential client. In addition to the demonstration, the potential client should go through every response to every question provided by the vendor and insure that the responses provided by the vendor are accurate. Much to the authors' surprise, one very large international bank that followed this process (i.e., warning the vendors in advance) revealed that only one vendor had been truthful in their responses to the RFP.

By following this process, risk managers will insure that they have a good understanding of the competing products and will help to root out unethical practices and unethical people from the software industry.

Most importantly, no product is perfect and no product has all features that all potential clients might want. Working with a key vendor for many years on the best future development path for the product is the very best way for a financial institution to achieve best practice in integrated risk management.

PAID PILOTS AS FINAL PROOF OF CONCEPT

The authors believe that once a vendor has been tentatively selected, best practice leads to a paid pilot with that vendor's system while the number two vendor waits on the sidelines. This is essentially a "mini-installation" and insures, in as complete a fashion as possible, that the vendor's solution works as advertised. While a few vendors resist the pilot concept, the authors would not recommend purchasing any risk management system without one if there is any doubt about the vendor selected. Paid pilots reveal the strengths and weaknesses of both the system and the management team of the vendor in short order. Since everyone in the world would be interested in a free pilot, potential clients should realize that vendors have to ration their scarce resources by charging for the pilot. We strongly recommend that the key team members on the pilot from the user's perspective lead the user's installation team.

In the past few years, mergers and acquisitions have reduced the number of significant risk vendors to a very small, single-digit number. By "risk vendor," we mean a vendor who can analyze all 700 million assets and liabilities on the balance

sheet of one of the largest banks in the world, not just the 2,000 traded products with a CUSIP or an ISIN number. Because the number of vendors is so small, it is very easy to do detailed due diligence on each vendor by calling their existing clients on an unsolicited basis. This will reveal, in the case of a few vendors, very high installation costs, high failures rates, a lack of innovation, and antiquated product architecture. By making these phone calls, one may be able to skip the costs of a paid pilot and proceed immediately with a full installation of the right vendor.

KEYS TO SUCCESS IN SOFTWARE INSTALLATION

The success rate on software installations in general is surprisingly low. The word “surprisingly” is the key word here. There should be only one reason that an installation fails and that is the sad situation where the software does not work—this should never happen if the preceding process is carefully followed. As a result, the failure of an installation of software that works is truly a failure of the management process and can be avoided 100 percent of the time if management is diligent at both the user and the vendor level.

The key to success from the user perspective is to have a dedicated project team 100 percent committed to the installation as their full-time job. The installation team should have a mix of business experts who represent the ultimate end user internally, and IT experts who are experts at getting the necessary input data and changing its format using standard industry tools. These tools include both data mapping and reporting of database (both input files and output files) contents.

Turnover on this project team is the greatest risk to the success of the project, because it is quite difficult for the vendor to replace the institution-specific knowledge that is lost when someone on the user’s project team resigns. This knowledge normally can come only from within the user’s institution itself.

The vendor’s installation team normally has three types of experts. The first, the project manager is a generalist whose expertise involves both clear and organized planning and reporting. The second type of expert is the excellent business user of the software from the vendor’s perspective. The third type of expert is the IT expert who works closely with his peers on the user’s project team.

The vendor’s project team members typically come from three sources: from the vendor’s staff itself, from a formal distributor of the software, or from a third-party systems integrator. The authors believe that the success rate is highest if the project team on the vendor side comes either from the vendor itself or from a long-term distributor of the product. If the only relationship with the vendor is on a project basis, a third-party systems integrator is in a complex position. The short-term revenue from the project is maximized by dragging out the installation, and the systems integrator suffers relatively little from the long-term reputational damage to the vendor from projects where installation was too slow, too expensive, or failed. From an end-user’s perspective, the user should want an installation team very focused on how important a successful installation is to the reputation of the vendor.

When this long-term reputational perspective is properly in place, a project should never fail unless the end user for some reason doesn’t staff the user side of the installation correctly.

VENDOR SIZE: LARGER VENDOR OR SMALL VENDOR?

From a conceptual point of view, every head of information technology wants a software vendor with a 10-year default probability of 0.03 percent or less, a 50-year history of innovation, and a project team of 25-year veterans that works for minimum wage. Unfortunately, the nature of the risk management software industry makes this dream highly unlikely to come about. As we look at the industry, a number of trends are quite clear:

- The innovation in the risk management business has consistently come from new, small firms.
- Many of these small firms have been able to accomplish only one major innovation and end users are forced to switch vendors in order to purchase the next innovation.
- As a result, small firms that can do multiple innovations and that allow changes in risk management models in the software with just a mouse-click by the user are very highly prized.
- This increases the size of the system that must be constructed by a new entrant, “raising the bar” and discouraging new entrants to the business.
- Large software firms generally add to their risk management products by acquisition of other firms.
- After an acquisition by a large firm, culture clashes between the large firm/small firm mentality often lead to high turnover among the most valuable staff of the acquired firm.
- In response, the large firm ceases innovation on the product and puts it into “cash cow” mode until it dies.
- When a risk product does die because of lack of innovation, the large software firm simply buys another vendor.

How does a sophisticated potential buyer of integrated risk management software deal with these issues? The most intelligent IT experts at major financial institutions look at the stability and ability of senior management of the vendor and of the core development team. This measure of vendor risk is almost flawless in indicating the chances of a successful installation and a long-term relationship, because it says a lot about the skill of the management of the vendor at all levels. It is a much better indicator of the probability of success than the size of the vendor, because firms of all sizes can score well on this measure if the firm is well run. At a large firm, compensation schemes designed to preserve the intellectual property of a small firm that is acquired are simple to structure and execute and yet relatively few large firms have been good at it.

A couple of incidents from the recent past illustrate how risky a large software vendor can be from an end-user perspective if the low turnover strategy is not pursued:

- A major U.K. software house bought an innovative risk software firm with a 15-person development team based in Los Angeles. Shortly after the acquisition, each member of the Los Angeles development team was ordered to move to London. As attractive as London is as a place to live, 14 of the 15 team members resigned rather than relocating. This essentially destroyed the product.

- A major New York-listed software company purchased an innovative risk management firm with substantial revenue and more than 300 employees. Within a year, the CEO of the acquired firm had left, and revenue in the product line fell by more than 50 percent.
- Another New York-listed firm acquired a company in the credit risk area with revenues of more than \$50 million a year and a staff of almost 200 people. Within two years, 12 of the 13 most senior people in the acquired company had resigned.
- A New York-listed firm in the credit risk area acquired a small but highly innovative software firm with expertise in interest rate risk management. Within one year, the CEO of the acquired firm had resigned and within two years, the acquirer notified software clients of the acquired firm that the software product of the acquired firm would no longer be supported.

From the end-user perspective, assessing the maturity, longevity, and commitment of the management team of the vendor is the single best way to assess the long-term risk of the vendor. As the examples above show, large size can increase risk if the managerial compensation systems are not handled skillfully.

BEING A BEST PRACTICE USER

Once an installation is successfully accomplished, the vendor-user relationship is just beginning. There are a number of best practices that we feel will enable end users to get the maximum benefit from an enterprise-wide software installation:

- Build mutual knowledge with the vendor through a regular exchange of visits and information. These relationships should span many levels of management, not just be concentrated at the working level.
- Encourage continued innovation by the vendor by adopting new modules as they emerge and by being an active participant in the design of expanded functionality.
- Be an active participant in the user group.
- Invest in the vendor by sponsoring new functionality with financial support to speed the degree of innovation.
- Establish a test database for regular testing of new versions from the vendor and make this test database available to the vendor.
- Regularly send staff to the vendor for advanced training.

In addition to these steps, we see an increasing trend toward long-term consulting contracts between the senior management team of the vendor and the most senior end users at major financial institutions. This kind of relationship insures that, above and beyond the day-to-day working-level relationship, senior management of the financial institution is fully exploiting an enterprise-wide risk management solution to achieve the objectives that we have outlined in this book.

We turn to that discussion in the next and final chapter in this book.

Shareholder Value Creation and Destruction

Risk management is the discipline that clearly shows management the risks and returns of every major strategic decision at both the institutional level and the transaction level. Moreover, the risk management discipline shows how to change strategy in order to bring the risk-return trade-off into line with the best long- and short-term interests of the institution.

In Chapter 1, we began this book with the preceding definition of risk management. In the 39 chapters after that, we showed, step-by-step, how to take advantage of existing financial and information technology to achieve these objectives. The 2006–2011 credit crisis has been the venue for the greatest mass destruction of shareholder value in the financial services industry in North America and Europe since the Great Depression. How can a serious and highly skilled risk manager help ensure that his or her financial institution is back on track, pursuing best practice, risk-adjusted shareholder valuation creation? That is the subject of this chapter. There are three keys to success, which we cover in turn:

1. Do no harm.
2. Measure the need to change.
3. Master the politics and exposition of risk management.

DO NO HARM

In the introduction to this book, we related the biggest mistakes and fallacies in the risk management discipline. The human tendency to repeat these mistakes is pervasive, and perhaps the biggest part of one's career as a best practice risk manager is to prevent these mistakes from happening. In other words, as a risk manager, ensure that the risk managers of the firm “do no harm” like the shareholder value destruction we have all witnessed during the 2006–2011 credit crisis. For that reason, we reemphasize the importance of “outing” those who are about to make these common mistakes:

- If it hasn't happened to me yet, it won't happen to me, even if it's happened to someone else.
- Silo risk management allows my firm to choose the "best of breed" risk model for our silo.
- I don't care what's wrong with the model. Everyone else is using it.
- I don't care what's wrong with the assumptions. Everyone else is using them.
- Mathematical models are superior to computer simulations.
- Big North American and European banks are more sophisticated than other banks around the world and we want to manage risk like they do.
- Goldman says they do it this way and that must be right.

If none of the largest financial institutions in Europe and the United States had made these mistakes, the 2006–2011 credit crisis would not have happened.

MEASURE THE NEED TO CHANGE

Having resolved to do no harm, for a risk manager at a financial services firm which has survived the crisis (either as a public firm or as an institution that has been at least temporarily nationalized), the next step is to measure the need to change. In other words, we need to ask "How good or bad is the current risk infrastructure of the firm and what do we need to do to fix it?" Surprisingly, risk managers at a number of firms that survived only through the largesse of the taxpayers suffer from hallucinations that manifest themselves in statements like "no risk system could have prevented what happened to my firm" or "all models failed in the credit crisis." A financial services firm CEO who hears such statements should immediately dismiss the speaker, who is an ignorant fool incapable of protecting the shareholders' (and the CEO's) interests.

In the wake of the 2006–2011 credit crisis, it is more obvious to the management, boards, and shareholders of major financial institutions that the legacy silo approach to risk management systems and risk management organization was simply ineffective in preventing the failure of some of the world's most prominent financial institutions. A fully integrated enterprise risk management capability, which combines credit risk, market risk, asset and liability management, liquidity risk, performance management, and capital allocation, exists today. This integrated enterprise risk management technology operates at the individual transaction level to assess risk accurately from the bottom up for every organizational unit and for the firm as a whole. This section presents a simple 10-point quiz that indicates whether the firm's existing primary risk system is capable of playing this role. In short, we are measuring the need to change.

1. Is your primary risk system capable of simulating default/no default for every reference name on the balance sheet at the transaction level for all types of reference names (consumer, nonpublic firms, public firms, sub-sovereigns, and sovereigns)?
 - Yes
 - No

2. Is your primary risk system capable of simulating random movements in default probabilities for every reference name as a function of time and any user-defined macroeconomic factors, such as home prices, stock index levels, interest rates, unemployment, growth in gross domestic product, foreign exchange rates, and gasoline and oil prices?
 - Yes
 - No
3. Is your primary risk management system capable of calculating a default-adjusted market value for every asset and liability, on balance sheet or off balance sheet, at the transaction level, regardless of whether that transaction is traded (i.e., has an ISIN or CUSIP number) or nontraded?
 - Yes
 - No
4. Is your primary risk system capable of calculating these market values at a large (up to 99,999 for example) user-defined set of future dates in a random interest rate environment where interest rates are driven by six or more risk factors (the minimum number of interest rate risk factors required by the Basel market risk rules)?
 - Yes
 - No
5. Is your primary risk system capable of calculating cash flows produced by every transaction in every scenario for every one of the user-defined cash flow periods for liquidity risk assessment?
 - Yes
 - No
6. Is your primary risk system capable of calculating GAAP interest income, interest expense, and amortization (if any) for every transaction in every scenario for every one of the user-defined accounting periods?
 - Yes
 - No
7. Is your primary risk system capable of a dynamic Monte Carlo simulation of alternative value-at-risk measures, recognizing that the balance sheet is constantly evolving instead of being static (the legacy VaR assumption)?
 - Yes
 - No
8. Is your primary risk system capable of aggregating these calculations to produce (for every user-defined business unit, portfolio subset, and for the institution as a whole) the mark-to-market value of the organization, its GAAP net income, its GAAP balance sheet, its dynamic value at risk, and its aggregate net cash flow? (*Note: These calculations should be available in every scenario in every user-defined period*)
 - Yes
 - No
9. Is your primary risk system capable of predicting the probability that the institution will fail to meet its target capital levels (whether on a regulatory or mark-to-market basis) in every one of the user-defined calculation periods?
 - Yes
 - No

10. Is your primary risk system capable of calculating the probability of default of the institution in any one of the user-defined calculation periods?
- Yes
 - No

RATING YOUR PRIMARY RISK SYSTEM

A firm's primary risk system can be ranked by how many of the previous 10 questions were answered with yes. The ranks are as follows:

- *5 or Less: Unacceptable.* The primary risk system is not capable of meeting even a limited number of the objectives now demanded as part of regulatory capital guidelines and stress tests, and it would not meet the basic corporate governance standards at a sophisticated organization.
- *6–7: Good.* The primary risk system being used is not bad, but substantial improvements to the risk system are necessary and these improvements are likely to be time consuming and costly with no guarantee of success.
- *8–9: Very good.* The primary risk system being used is close to the state of the art.
- *10: Congratulations.* You are using a state-of-the-art risk management system.

This ranking of risk systems cuts to the heart of what distinguishes integrated enterprise risk management best practice from legacy silo risk applications. There are hundreds of other questions that can be added to this list.¹

MASTER THE POLITICS AND EXPOSITION OF RISK MANAGEMENT: SHAREHOLDER VALUE CREATION

Each of the authors works closely with major hedge funds and other institutional investors on a daily basis. Their objectives from a risk management perspective are ruthlessly simple: Given equivalent levels of risk, which bond should I go long and which bond should I go short? Day after day, if they do this skillfully, their mark-to-market returns versus their benchmark will be revealed to clients and potential clients each month. All will be able to see clearly whether the hedge funds are achieving their objectives.

Of all the financial institutions that the authors have worked with, hedge funds have the greatest sense of urgency and the greatest link between their success in achieving best practice in risk management and their personal compensation. Their own shareholder value created is manifestly clear—the excess of their management fees in excess of costs.

Among major traditional financial institutions, there are some very clear “stars” and some rapid progress being made.

At a recent gathering of 20 senior risk managers (just prior to the 2006 to 2011 credit crisis), nineteen of the 20 members surveyed reported that their institution was

regularly using mark-to-market technology to quantify their interest rate risk using the techniques like those we outlined in Chapters 6 through 10. Thirty years ago, perhaps only one or two of the institutions surveyed would have given the same answer, with the others relying primarily on the simulation of financial accrual-based net income as their primary focus for risk management effects. That focus, while clearly of interest to the CEO and Wall Street analysts, was certainly not good enough to prevent the \$1 trillion collapse of the savings and loan industry from rising rates in the 1980s and 1990s. In that sense, in the past 30 years, much progress has been made.

At the same time, the survey of the risk managers gathered at that session was striking in a “so close, but oh so far” sense. When the risk managers marked-to-market the balance sheet of their organizations, much to our surprise, most of them used only a single yield curve (usually the LIBOR curve, occasionally the risk-free curve) to calculate the market value of every obligation on this balance sheet. The choice of the LIBOR-swap curve is particularly ironic in the wake of the 2011–2012 LIBOR manipulation scandal. The institutions completely ignored credit risk in the mark-to-market process, and that lack of preparation says a lot about the destruction of shareholder value in the credit crisis. In most large financial institutions, the interest rate risk system is the primary risk system because it alone holds information on most of the assets and liabilities of the firm. In a sense, the risk managers were universally admitting that their firms failed the 10-point quiz of the previous section. Many of them are moving rapidly now to fix the problem.

Clearly, no hedge fund manager would last long by marking-to-market the present value of an IBM bond with present value factors from the LIBOR curve. It makes no sense in either theory or practice to use only one yield curve if your mandate is to create shareholder value on a risk-adjusted basis.

The survey response indicates that almost all of the risk managers surveyed were still “held prisoner” in a tightly compartmentalized “interest rate risk management” box that falls far short of the risk management definition above. Their interest was not in the mark-to-market calculation per se, but the sensitivity of the mark-to-market calculation to changes in the level of interest rates as captured by the single yield curve.

With a bit of fine-tuning of policies and procedures and with better risk management software, the institutions surveyed could have been achieving a lot more value for their shareholders using the full complement of tools and techniques in this book, not just the interest rate analytics of Chapters 6 through 10. We can see from the examples of fund managers, hedge funds, insurance companies, banks, and securities firms exactly what to do. From a senior management perspective, we need to demand the same accuracy and insights from a total balance sheet perspective that we do on the trading floor on a transaction-by-transaction basis.

Using modern integrated risk management technology in a daily production environment, management can achieve this objective at a fraction of the cost of running the disparate systems for compartmentalized risk management that we outlined in Chapter 40. Many of the world’s best financial institutions have already taken this approach, and they are very successfully exploiting their peers whose risk management practices remain bogged down in the approaches used in the 1980s and early 1990s. In the next section, we summarize what senior management of these institutions can see on a daily basis that their peers cannot yet perceive.

DAILY MANAGEMENT REPORTING OF TOTAL RISK

A state-of-the-art set of integrated risk management reports would contain the following information on the institution's risk and shareholder value added:

- Default probabilities for the institution compared to target default probabilities:
 - Derived from historical defaults
 - Derived from bond prices
 - Derived from credit derivatives prices (if and only if based on actual trades)
- Default probabilities for the five most important competitors of the institution.
- Default probabilities for the 10 largest mark-to-market credit exposures of the institution.
- Sensitivity of all of these default probabilities to changes in:
 - The economy via a stock price index like the S&P 500
 - Home prices
 - Commercial real estate prices
 - Interest rates
 - Foreign exchange rates
 - Oil prices
 - Other commodity prices
- Cost of the Jarrow-Merton put option that would eliminate all of the institution's
 - Interest rate risk
 - Liquidity risk
 - Credit risk
 - Total risk
- Sensitivity of the Jarrow-Merton put option for each of the objectives above with respect to changes in:
 - The economy via a stock price index like the S&P 500
 - Home prices
 - Commercial real estate prices
 - Interest rates
 - Foreign exchange rates
 - Oil prices
 - Other commodity prices
- Year-to-date shareholder value creation for the institution's
 - 10 largest business units
 - 10 most important product units
 - 10 most important geographical units

On these bases:

 - Daily mark to market of the relevant portfolio versus a risk-free benchmark with identical interest rate characteristics.
 - Financial accounting basis using matched maturity transfer pricing versus the appropriate yield curve (risk-free yield curve for assets, the institution's marginal cost of funds for liabilities).
- Probability of not meeting net income targets for each of the next five years

- Probability of needing “excessive” provisions for loan losses in each of the next five years
- Probability of missing managerial target capital ratios in each of the next five years.
- Probability of missing regulatory capital ratios in each of the next five years.

Among the readers of this book, there are probably many who would like to have this information and do not have it. It is available today, and many of your competitors already have access to it. Getting this information for your own institution simply involves going through the IT process that we outlined in Chapter 40. The Federal Deposit Insurance Corporation, through its Loss Distribution Model, announced on December 10, 2003, that it is already going through the process of generating information just like this for every financial institution in the United States with deposit insurance. Similarly, the Office of the Comptroller of the Currency has made it public that it is using the reduced form probability approach of Chapter 16 to evaluate the risk of every publicly held banking firm in the United States, the public firms they lend to, and 183 sovereign entities. If the regulators already have information like this about your institution, how can you not be looking at it as well?

MOVING FROM COMMON PRACTICE TO BEST PRACTICE

In Chapters 1 and 2, the authors related how long it took for commonly accepted risk management tools of today to become accepted as conventional wisdom. In the case of many management tools, this has taken two or three decades, but the pace of innovation is increasing rapidly because competition in the financial services industry is rapidly intensifying and the consequences of being wrong have been devastating for many large financial services firms, their shareholders, and their CEOs as we saw during the 2006–2011 credit crisis. Note that the first reduced form credit model was introduced only in 1995 (by Jarrow-Turnbull), that it was first offered commercially in 2002, and that the FDIC had already adopted the technology on the largest portfolio of credit insurance in the world by 2003. Black and Scholes would have been ecstatic if their options model was as widely adopted in twice that time.

Financial institutions face many barriers to innovation. As a group, the employees of financial institutions (including the authors in their prior lives) are a conservative group. Of these conservative people, risk managers tend to be super conservative. Within the group of risk managers, one risk expert likes to categorize the breed as being either “place keepers” or “bomb throwers.” Place keepers maintain the risk management discipline using existing technology with modest innovation. Bomb throwers want to make great leaps forward.

For this reason, the people involved in risk management play a critical role in the pace of innovation. For various levels of management, there are some simple ways of moving from common practice to best practice, because the regulatory agencies are moving in that direction rapidly, as the FDIC Loss Distribution Model shows. We summarize each in turn.

THE SENIOR MANAGEMENT PERSPECTIVE

Senior managers in major financial institutions should be able to readily access the daily risk reports that we outlined previously—they should be demanded from staff.

For an institution with a traditional risk management structure and software infrastructure, the cost of these reports is literally negative. By gradually eliminating all of the six types of vendors discussed in Chapter 40, there will be cash left over even after an enterprise-wide risk management solution has been put in place. The shareholder value created from this process will be very large, but even from a short-term dollars-and-cents perspective it simply reduces the direct and indirect cost of risk management and provides much clearer strategic insights.

THE MIDDLE MANAGEMENT PERSPECTIVE

Middle managers are the key to success of every financial institution from a shareholder value-added perspective. They have both the hands-on knowledge of the working level and the strategic perspective of senior management.

What if senior management is pushing to move from common practice to best practice, but the working level does not have the knowledge or skill set to get there? The answer is simple—replace the working level with better people.² More likely, the working level will be interested if pushed, and a change in risk vendors will provide a welcome challenge and a well-justified sense of contribution to the well-being of the organization.

What if senior management is the barrier? It is so ironic that you need a license to drive a car, but you don't need a license to be president of a major financial institution. The personality of the CEO of a financial institution has a tremendous impact on shareholder value creation and risk management best practice. How can a CEO not want to know the risk that his own institution will default? Especially if the regulators, shareholders, and rating agencies are all looking at those numbers? The quotations in the Introduction showed how many of the world's largest financial institutions were run by CEOs who had little hands-on knowledge of their own business. It's not likely that CEOs of little knowledge will be occupying the CEO chairs in the near future.

When, in spite of the CEO casualty rate of 2006 to 2011, the CEO has little or no interest in risk management, there is still a way forward. With luck, a little pushing will create the interest, at least somewhere on the senior risk management committee. If not, just wait—your time will come in the senior spot and you will be able to change things quickly. Worst case, you can use cost reduction as the rationale for putting an enterprise-wide risk management infrastructure in place. If the cost-reduction benefits are not apparent, you are not talking to the right vendors.

THE WORKING-LEVEL PERSPECTIVE

As the authors related early in this book, the working level has the longest wait before they can make a major impact on how the institution looks at risk-adjusted shareholder value creation and destruction. One of us tried to sell the mark-to-market risk

management concept that is now standard at the major banks around the world for three years at a major bank without success. When faced with that kind of barrier, what should you do as a working-level staffer?

One of the most effective motivators for senior management is to prove to them that they are at risk of no longer achieving common practice in risk management, let alone best practice. Peer bank actions are the most effective of convincing people that the financial institution is being left behind. Another opportunity to effect change comes in the aftermath of a crisis in either the financial markets you operate in (like the 2006–2011 credit crisis or Korea after the 1997–1998 Asia crisis), or at the institution where you work after a risk management incident. Being part of the solution is very good for your career.

Many times, unfortunately, the messenger who conveys the message that a crisis is coming is not welcomed before the fact and not thanked afterward when they prove to be right. This is a situation to be avoided at all costs. When you find yourself employed by a financial institution that lacks the institutional will to create risk-adjusted shareholder value with the best tools available, quit! The firm is destined to be acquired or to get into trouble, so why sit around and wait for it to happen, ruining your own reputation as a risk manager? If you need any proof, the blog-a-sphere is filled with tales of risk managers who raised the red flag of warning at their firms but stuck around to be the fall guys.

GETTING HELP TO CREATE SHAREHOLDER VALUE

The tools and techniques in this book are fully implemented and in place at some of the most sophisticated financial institutions around the world. The concepts are tried and true, although in some cases they are being applied in different ways from the ways they might have been used before.

Since the techniques are already at work, getting help to speed your own implementation of these ideas is as simple as sending us an e-mail.³ We look forward to hearing from you, and we thank you for making it all the way to the last sentence of this book.

POSTSCRIPT

This chapter was written during the Tanabata Festival in Sendai, Japan, in August 2012. We write in the memory of those whose lives were lost in the great earthquake and tsunami of March 11, 2011, and in honor of those who celebrate the lives of their lost loved ones. —Donald R. van Deventer, Kenji Imai, and Mark Messler

NOTES

1. For assistance in preparing a request for proposal that will clearly identify a state-of-the-art risk system, contact our colleagues at e-mail: info@kamakuraco.com.
2. If you need some recommendations, contact us at e-mail: info@kamakuraco.com.
3. Donald R. van Deventer, Kenji Imai, and Mark Messler can be reached at e-mail: info@kamakuraco.com.

Bibliography

- Adams, K. J., and D. R. van Deventer. 1993. "Monte Carlo Simulation: The Pros and the Cons." *Balance Sheet* 2, no. 3 (Autumn): 18–24.
- Adams, K. J., and D. R. van Deventer. 1994. "Fitting Yield Curves and Forward Rate Curves with Maximum Smoothness." *Journal of Fixed Income* (June): 52–61.
- Adrian, T., R. K. Crump, and Emanuel Moench. 2012. "Pricing the Term Structure with Linear Regressions." Working Paper, Federal Reserve Bank of New York, May 9.
- Allen, L., and A. Saunders. 2003. "A Survey of Cyclical Effects in Credit Risk Measurement Models. Bank for International Settlements Working." Paper 126.
- Altman, E. 1968. "Financial Ratios, Discriminant Analysis, and the Prediction of Corporate Bankruptcy." *Journal of Finance* 23, no. 4: 589–609.
- Altman, E., R. Haldeman, and P. Narayanan. 1977. "ZETA Analysis: A New Model to Identify Bankruptcy Risk of Corporations." *Journal of Banking and Finance* 1, no. 1: 29–55.
- Amin, K., and R. A. Jarrow. 1991. "Pricing Foreign Currency Options under Stochastic Interest Rates." *Journal of International Money and Finance* 10: 310–329.
- Amin, K., and R. A. Jarrow. 1992. "Pricing American Options on Risky Assets in a Stochastic Interest Rate Economy." *Mathematical Finance* 2, no. 4 (October): 217–237.
- Andersen, L. 2007. "Discount Curve Construction with Tension Splines." *Review of Derivatives Research* 10, no. 3: 227–267.
- Andersen, L., and D. Buffum. 2004. "Calibration and Implementation of Convertible Bond Models." *Journal of Computational Finance*, published online only (Winter).
- Arnold, L. 1974. *Stochastic Differential Equations: Theory and Applications*. New York: John Wiley & Sons.
- Ayache, E., P. A. Forsyth, and K. R. Vetzal. 2003. "Valuation of Convertible Bonds with Credit Risk." *Journal of Derivatives* 11, no. 1 (April): 9–29.
- Basel Committee on Banking Supervision. 2001. *Consultative Document: The New Capital Accords*. Basel: Bank for International Settlements.
- Basel Committee on Banking Supervision. 2004. "Basel II: International Convergence of Capital Measurement and Capital Standards: a Revised Framework." June.
- Basel Committee on Banking Supervision. 2010. "Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems." December. Revised June 2011.
- Berk, J., and R. Roll. 1988. "Adjustable Rate Mortgages: Valuation." *Journal of Real Estate Finance and Economics* 1, no. 2 (June): 163–184.
- Bermudez, A., and N. Webber. 2003. "An Asset Based Model of Defaultable Convertible Bonds with Endogenous Recovery." Working Paper, City University, November 26.
- Bharath, S., and T. Shumway. 2008. "Forecasting Default with the Merton Distance to Default Model." *Review of Financial Studies* (May): 1339–1369.
- Black, F. 1976. "The Pricing of Commodity Contracts." *Journal of Financial Economics* 3, no. 1–2 (March): 167–179.
- Black, F. 1995. "Interest Rates as Options." *Journal of Finance* 50, no. 5: 1371–1376.
- Black, F., E. Derman, and W. Toy. 1990. "A One-Factor Model of Interest Rates and Its Application to Treasury Bond Options." *Financial Analysts Journal* (January–February): 33–39.

- Black, F., and P. Karasinski. 1991. "Bond and Option Pricing when Short Rates Are Log-normal, Capital Standards: Interest Rate Risk." *Financial Analysts Journal* (July–August): 52–59.
- Black, F., and M. Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81, no. 3 (May–June): 637–654.
- Blum, M. 1974. "Failing Company Discriminant Analysis." *Journal of Accounting Research* 12, no. 1 (Spring): 1–25.
- Bohn, J., N. Arora, and I. Korablev. 2005. "Power and Level Validation of the EDF™ Credit Measure in the U.S. Market." Moody's KMV memorandum.
- Boyle, P. 1977. "Options: A Monte Carlo Approach." *Journal of Financial Economics* 4 (May): 323–328.
- Breeden, D. T. 1994. "Complexities of Hedging Mortgages." *Journal of Fixed Income* 4, no. 3 (December): 6–42.
- Brennan, M. J., and E. Schwartz. 1979. "A Continuous Time Approach to the Pricing of Bonds." *Journal of Banking and Finance* 3, no. 2: 133–155.
- Buono, M., R. B. Gregory-Allen, and U. Yaari. 1992. "The Efficacy of Term Structure Estimation Techniques: A Monte Carlo Study." *Journal of Fixed Income* (March): 52–59.
- Campbell, J. Y., J. Hilscher, and J. Szilagyi. 2008. "In Search of Distress Risk." *Journal of Finance* (December): 2899–2939.
- Campbell, John Y., J. Hilscher, and J. Szilagyi. 2011. "Predicting Financial Distress and the Performance of Distressed Stocks." *Journal of Investment Management* 63, no. 6: 1–21.
- Campbell, J. Y., and G. B. Taksler. 2003. "Equity Volatility and Corporate Bond Yields." *Journal of Finance* 58, no. 6: 2321–2350.
- Chan, K. C., G. A. Karolyi, F. A. Longstaff, and A. B. Sanders. 1992. "An Empirical Comparison of Models of the Short Term Interest Rate." *Journal of Finance* 47, no. 3 (July): 1209–1228.
- Chava, S., and R. A. Jarrow. 2003. "Modeling Loan Commitments and Liquidity Crisis." Cornell University Working Paper.
- Chava, S., and R. A. Jarrow. 2008. "Modeling Loan Commitments." *Finance Research Letters* 5, no. 1 (March): 11–20.
- Chen, D. H., H. H. Huang, R. Kan, A. Varikooty, and H. N. Wang. 1998. "Modelling and Managing Credit Risk." In *Asset & Liability Management: A Synthesis of New Methodologies*, ed. M. Ong. London: Risk Publications.
- Chen, L. 1994. *Interest Rate Dynamics, Derivatives Pricing, and Risk Management*. Berlin: Springer.
- Chen, R. 1992. "Exact Solutions for Futures and European Futures Options on Pure Discount Bonds." *Journal of Financial and Quantitative Analysis* 27 (March): 97–108.
- Chen, R., and L. Scott. 1992. "Pricing Interest Rate Options in a Two-Factor Cox-Ingersoll-Ross Model of the Term Structure." *Review of Financial Studies* 5, no. 4: 613–636.
- Collin-Dufresne, P., R. S. Goldstein, J. S. Martin. 2001. "The Determinants of Credit Spread Changes." *Journal of Finance* (December): 2177–2207.
- Cox, D. R. 1972. "Regression Models and Life Tables." *Journal of the Royal Statistical Society* 34, no. 2: 187–220.
- Cox, J. C., J. E. Ingersoll, Jr., and S. A. Ross. 1980. "An Analysis of Variable Rate Loan Contracts." *Journal of Finance* 35, no. 2: 389–403.
- Cox, J. C., J. E. Ingersoll, Jr., and S. A. Ross. 1981. "The Relationship between Forward Prices and Futures Prices." *Journal of Financial Economics* 9, no. 4 (December): 321–346.
- Cox, J. C., J. E. Ingersoll, Jr., and S. A. Ross. 1985. "A Theory of the Term Structure of Interest Rates." *Econometrica* 53, no. 2: 385–407.
- Cox, J. C., S. A. Ross, and M. Rubinstein. 1979. "Option Pricing: A Simplified Approach." *Journal of Financial Economics* 7: 229–263.

- Dai, Q., and K. J. Singleton. 2000. "Specification Analysis of Affine Term Structure Models." *Journal of Finance* 55, no. 5: 1943–1978.
- Dai, Q., and K. J. Singleton. 2003. "Term Structure Dynamics in Theory and Reality." *Review of Financial Studies* 16, no. 3: 631–678.
- Dash, E., and J. Creswell. "The Reckoning: Citigroup Saw No Red Flags Even as It Made Bolder Bets." *New York Times*, November 22, 2008.
- Dickler, D. T., R. A. Jarrow, and D. R. van Deventer. 2011a. "Inside the Kamakura Book of Yields: A Pictorial History of 50 Years of U.S. Treasury Forward Rates." Kamakura Corporation memorandum, September 13.
- Dickler, D. T., R. A. Jarrow, and D. R. van Deventer. 2011b. "Inside the Kamakura Book of Yields: A Pictorial History of 50 Years of U.S. Treasury Zero Coupon Bond Yields." Kamakura Corporation memorandum, September 26.
- Dickler, D. T., R. A. Jarrow, and D. R. van Deventer. 2011c. "Inside the Kamakura Book of Yields: A Pictorial History of 50 Years of U.S. Treasury Par Coupon Bond Yields." Kamakura Corporation memorandum, October 5.
- Dickler, D. T., and D. R. van Deventer. 2011a. "Inside the Kamakura Book of Yields: An Analysis of 50 Years of Daily U.S. Treasury Forward Rates." Kamakura blog, www.kamakuraco.com, September 14, 2011.
- Dickler, D. T., and D. R. van Deventer. 2011b. "Inside the Kamakura Book of Yields: An Analysis of 50 Years of Daily U.S. Treasury Zero Coupon Bond Yields." Kamakura blog, www.kamakuraco.com, September 26, 2011.
- Dickler, D. T., and D. R. van Deventer. 2011c. "Inside the Kamakura Book of Yields: An Analysis of 50 Years of Daily U.S. Par Coupon Bond Yields." Kamakura blog, www.kamakuraco.com, October 6, 2011.
- Dothan, L. U. 1978. "On the Term Structure of Interest Rates." *Journal of Financial Economics* 6, no. 1: 59–69.
- Duffee, G. 1999. "Estimating the Price of Default Risk." *Review of Financial Studies* 12, no. 1: 197–226.
- Duffie, D. *Dynamic Asset Pricing Theory*. Princeton, NJ: Princeton University Press, 1996.
- Duffie, D., and R. Kan. 1994. "Multi-Factor Term Structure Models." *Philosophical Transactions: Physical Sciences and Engineering* 347, no. 1684 (Mathematical Models in Finance, June): 577–586.
- Duffie, D., and K. J. Singleton. 1997. "An Econometric Model of the Term Structure of Interest-Rate Swap Yields." *Journal of Finance* 52, no. 4: 1287–1321.
- Duffie D., and K. Singleton. 1999. "Modelling Term Structures of Defaultable Bonds." *Review of Financial Studies* 12, no. 4: 197–226.
- Dunn, K. B., and C. Spatt. 1988. "Private Information and Incentives: Implications for Mortgage Contract Terms and Pricing." *Journal of Real Estate Finance and Economics* 1 (April): 47–60.
- Ellis, D. M., and J. V. Jordan. 2001. "The Evaluation of Credit Union Non-Maturity Deposits." National Economic Research Associates.
- Elton, E. J., M. J. Gruber, D. Agrawal, and C. Mann. 2001. "Explaining the Rate Spread on Corporate Bonds." *Journal of Finance* 56, no. 1 (February): 247–277.
- Falkenstein, E., and A. Boral. 2000. "RiskCalc™ for Private Companies: Moody's Default Model." Moody's Investors Service memorandum.
- Federal Reserve System. 1993. "Risk-Based Capital Standards: Interest Rate Risk." Notice of Proposed Rulemaking, September 14.
- Federal Reserve System. 2009a. "The Supervisory Capital Assessment Program: Design and Implementation." April 24.
- Federal Reserve System. 2009b. "The Supervisory Capital Assessment Program: Overview of Results." May 7.

- Federal Reserve System. 2011. "Comprehensive Capital Analysis and Review: Summary Instructions and Guidance." November 22.
- First Boston Corporation. 1984. *Handbook of Securities of the United States Government and Federal Agencies and Related Money Market Instruments*. New York.
- Fishman, G. S. 1995. *Monte Carlo: Concepts, Algorithms, and Applications*. Berlin: Springer.
- Flannagan, C. T., M. D. Herskovitz, and H. T. Loy. 1989. "Understanding and Modelling Cost of Funds ARMs Prepayments." Mortgage-Backed Securities Research, Merrill Lynch, February.
- Flesaker, B. 1993. "Testing the Heath-Jarrow-Morton/Ho-Lee Model of Interest Rate Contingent Claims Pricing." *Journal of Financial and Quantitative Analysis* 28, no. 4 (December): 483–495.
- Friedman, A. 1975. *Stochastic Differential Equations and Applications*, vol. 1. New York: Academic Press.
- Garman, M. B., and S. W. Kohlhagen. 1983. "Foreign Currency Option Values." *Journal of International Money and Finance* 2, no. 3: 231–237.
- Geske, R. 1979. "The Valuation of Compound Options." *Journal of Financial Economics* 7, no. 1: 63–81.
- Green, J., and J. B. Shoven. 1986. "The Effect of Interest Rates on Mortgage Prepayments." *Journal of Money, Credit and Banking* 18, no. 1: (February): 41–59.
- Heath, D., R. A. Jarrow, and A. Morton. 1990a. "Bond Pricing and the Term Structure of Interest Rates: A Discrete Time Approximation." *Journal of Financial and Quantitative Analysis* 25, no. 4: 419–440.
- Heath, D., R. A. Jarrow, and A. Morton. 1990b. "Contingent Claims Valuation with a Random Evolution of Interest Rates." *Review of Futures Markets* 9, no. 1: 54–76.
- Heath, D., R. A. Jarrow, and A. Morton. 1992a. "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claim Valuation." *Econometrica* 60, no. 1: 77–105.
- Heath, D., R. A. Jarrow, and A. Morton. 1992b. "Easier Done than Said." *RISK Magazine*, October, 77–80.
- Hildebrand, F. B. 1987. *Introduction to Numerical Analysis*. New York: Dover.
- Hilliard, J. E., J. Madura, and A. L. Tucker. 1991. "Currency Option Pricing with Stochastic Domestic and Foreign Interest Rates." *Journal of Financial and Quantitative Analysis* 26, no. 2: 139–151.
- Hilscher, J., and M. Wilson. 2011. "Credit Risk and Credit Ratings." Brandeis University working paper.
- Ho, T. S. Y., and S.-B. Lee. 1986. "Term Structure Movements and Pricing Interest Rate Contingent Claims." *Journal of Finance* 41, no. 5 (December): 1011–1029.
- Hosmer, D. W., and S. Lemeshow. 2000. *Applied Logistic Regression*. New York: John Wiley & Sons.
- Hu, G. n.d. "Mortality Assessment Technology: A New Tool for Life Insurance Underwriting." BioSignia, Inc. memorandum.
- Huang, J., and W. Kong. 2003. "Explaining Credit Spread Changes: Some New Evidence from Option-Adjusted Spreads of Bond Indexes." *Journal of Derivatives* (Fall): 30–44.
- Hull, J. 1992. *Options, Futures, and other Derivative Securities*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall.
- Hull, J., and A. White. 1990a. "Valuing Derivative Securities Using the Explicit Finite Difference Method." *Journal of Financial and Quantitative Analysis* 25, no. 1: 87–100.
- Hull, J., and A. White. 1990b. "Pricing Interest-Rate Derivative Securities." *Review of Financial Studies* 3, no. 4: 573–592.
- Hull, J., and A. White. 1993a. "One Factor Interest Rate Models and the Valuation of Interest Rate Derivative Securities." *Journal of Financial and Quantitative Analysis* 28, no. 2: 235–254.

- Hull, J., and A. White. 1993b. "Efficient Procedures for Valuing European and American Path-Dependent Derivatives." *Journal of Derivatives* 1, no. 1: 21–31.
- Hull, J., and A. White. 1994a. "Numerical Procedures for Implementing Term Structure Models I: Single-Factor Models." *Journal of Derivatives* 2, no. 1: 7–16.
- Hull, J., and A. White. 1994b. "Numerical Procedures for Implementing Term Structure Models II: Two-Factor Models." *Journal of Derivatives* 2, no. 2: 37–48.
- Hull, J., and A. D. White. 1995. "A Note on the Models of Hull and White for Pricing Options on the Term Structure: Response." *Journal of Fixed Income* 5, no. 2: 97–102.
- Hutchinson, D., and G. Pennacchi. 1996. "Measuring Rents and Interest Rate Risk in Imperfect Financial Markets: The Case of Retail Bank Deposits." *Journal of Financial and Quantitative Analysis* 31, no. 3 (September): 399–417.
- Ingersoll Jr., J. E. 1987. *Theory of Financial Decision Making*. Savage, MD: Rowman & Littlefield.
- Ingersoll, Jr., J. E., J. Skelton, and R. L. Weil. 1978. "Duration Forty Years Later." *Journal of Financial and Quantitative Analysis* 13, no. 4 (November): 627–648.
- Jamshidian, F. 1989. "An Exact Bond Option Formula." *Journal of Finance* 44, no. 1 (March): 205–209.
- Jamshidian, F. 1990. "Bond and Option Evaluation in the Gaussian Interest Rate Model." Financial Strategies Group, Merrill Lynch.
- Janosi, T., and R. A. Jarrow. 2002. "Programmers Notes on Yield Curve Smoothing." Kamakura Corporation memorandum.
- Jarrow, R. A. 1999. "Technical Guide: Default Probabilities Implicit in Debt and Equity Prices." Kamakura Corporation Technical Guide (revised 2001).
- Jarrow, R. A. 2001. "Default Parameter Estimation Using Market Prices." *Financial Analysts Journal* 57, no. 5 (September–October): 1–18.
- Jarrow, R. A. 2002. *Modelling Fixed Income Securities and Interest Rate Options*, 2nd ed. Stanford CA: Stanford University Press.
- Jarrow, R. A. 2004. "Risky Coupon Bonds as a Portfolio of Zero-Coupon Bonds." *Finance Research Letters* 1, no. 2 (June): 100–105.
- Jarrow, R. A. 2009. "The Term Structure of Interest Rates." *Annual Review of Financial Economics* 1: 69–96.
- Jarrow, R. A. 2011. "Risk Management Models: Construction, Testing, Usage." *Journal of Derivatives* 18, no. 4 (Summer): 89–96.
- Jarrow, R. A. 2012. "Problems with Using CDS Spreads to Infer Default Probabilities." *Journal of Fixed Income* (forthcoming).
- Jarrow, R. A., R. Bennett, M. Fu, D. Nuxoll, and H. Zhang. 2003. "A General Martingale Approach to Measuring and Valuing the Risk to the FDIC Deposit Insurance Funds." Federal Deposit Insurance Corporation memorandum.
- Jarrow, R. A., and S. Chava. 2002. "A Comparison of Explicit versus Implicit Estimates of Default Probabilities." Working Paper, Cornell University.
- Jarrow, R. A., and S. Chava. 2004. "Bankruptcy Prediction with Industry Effects." *Review of Finance* 8, no. 4: 537–569.
- Jarrow, R. A., T. Janoski, and F. Zullo. 1999. "An Empirical Analysis of the Jarrow-van Deventer Model for Valuing Non-Maturity Deposits." *Journal of Derivatives* 7, no. 1 (Fall): 8–31.
- Jarrow, R. A., D. Lando, and S. Turnbull. 1997. "A Markov Model for the Term Structure of Credit Risk Spreads." *Review of Financial Studies* 10, no. 2: 481–523.
- Jarrow, R. A., D. Lando, and F. Yu. 2005. "Default Risk and Diversification: Theory and Applications." *Mathematical Finance* (January): 1–26.
- Jarrow, R. A., and J. Levin. 1996. "The Pricing of Non-Maturity Deposits." *American Banker* (March 8).

- Jarrow, R. A., L. Li, L. Matz, M. Mesler, and D. R. van Deventer. 2006. "Kamakura Risk Information Services Technical Guide." Version 4.1, Appendix D: Mapping Default Probabilities and Ratings to Implied Spreads. Kamakura Corporation memorandum, May 25.
- Jarrow, R. A., L. Li, L. Matz, M. Mesler, and D. R. van Deventer. 2007. "The Determination of Corporate Credit Spreads." *RISK Magazine* (September).
- Jarrow, R. A., and P. Protter. 2004. "Structural versus Reduced Form Models: A New Information Based Perspective." *Journal of Investment Management*, 2, no. 2: 1–10.
- Jarrow, R. A., and A. Purnanandam. 2005. "Generalized Coherent Risk Measures: The Firm's Perspective." *Finance Research Letters*, 2 (March): 23–29.
- Jarrow, R. A., and A. Purnanandam. 2007. "The Valuation of a Firm's Investment Opportunities: A Reduced Form Credit Risk Perspective." *Review of Derivatives Research* 10, no. 1 (January): 39–58.
- Jarrow, R. A., and S. Turnbull. 1995. "Pricing Derivatives on Financial Securities Subject to Credit Risk." *Journal of Finance* 50, no. 1: 53–85.
- Jarrow, R. A., and S. Turnbull. 1996. *Derivative Securities*. Cincinnati, OH: South-Western.
- Jarrow, R. A., and D. R. van Deventer. 1996. "Power Swaps: Disease or Cure?" *RISK Magazine* (February).
- Jarrow, R. A., and D. R. van Deventer. 1998. "The Arbitrage-Free Valuation and Hedging of Demand Deposits and Credit Card Loans." *Journal of Banking and Finance* 22, no. 3 (March): 249–272.
- Jarrow, R. A., and D. R. van Deventer. 1998. "Integrating Interest Rate Risk and Credit Risk in ALM." In *Asset & Liability Management: A Synthesis of New Methodologies*, ed. M. Ong. London: Risk Publications.
- Jarrow, R. A., and D. R. van Deventer. 2004. "Practical Use of Credit Risk Models in Loan Portfolio and Counterparty Exposure Management: An Update." In *Credit Risk: Models and Management*, 2nd ed., ed. David Shimko. London: Risk Publications.
- Jarrow, R. A., and D. R. van Deventer. 2005. "Estimating Default Correlations Using a Reduced Form Model." *Risk Magazine* (January): 83–88.
- Jarrow, R. A., D. R. van Deventer, and X. Wang. 2003. "A Robust Test of Merton's Structural Model for Credit Risk." *Journal of Risk* 6, no. 1 (Fall): 39–58.
- Jarrow, R. A., and Yildirim, Yildiray. 2003. "Pricing Treasury Inflation Protected Securities and Related Derivatives Using an HJM Model." *Journal of Financial and Quantitative Analysis*, 38, no. 2 (June): 337–359.
- Jones, E. P., S. P. Mason, and E. Rosenfeld. 1984. "Contingent Claims Analysis of Corporate Capital Structure: An Empirical Investigation." *Journal of Finance* 39: 611–626.
- Jorion, P. 2006. *Value at Risk: The New Benchmark for Management Financial Risk*, 3rd ed. New York: McGraw-Hill.
- Kamakura Corporation. 2003a. "Kamakura Risk Information Services Technical Guide." Version 2.2, May.
- Kamakura Corporation. 2003b. "Kamakura Risk Information Services Technical Guide." Version 3.0, December.
- Kamakura Corporation. 2006. "Kamakura Risk Information Services Technical Guide." Version 4.1, January.
- Kamakura Corporation. 2011. "Kamakura Risk Information Services Technical Guide." Version 5.0, March.
- Karatzas, I., and S. E. Shreve. 1991. *Brownian Motion and Stochastic Calculus*. New York: Springer.
- King, G., and S. Soneji. 2011. "Statistical Security for Social Security." Harvard University Working Paper.
- Lando, D. 2004. *Credit Risk Modeling: Theory and Applications*. Princeton, NJ: Princeton University Press.

- Lewis, M. 2011. *Moneyball: The Art of Winning an Unfair Game*. New York: W. W. Norton.
- Longstaff, F. A. 1989. "A Nonlinear General Equilibrium Model of the Term Structure of Interest Rates." *Journal of Financial Economics* 23, no. 2: 195–224.
- Longstaff, F. A., and E. S. Schwartz. 1992a. "A Two Factor Interest Rate and Contingent Claims Valuation." *Journal of Fixed Income* 2, no. 3: 16–23.
- Longstaff, F. A., and E. S. Schwartz. 1992b. "Interest Rate Volatility and the Term Structure: A Two-Factor General Equilibrium Model." *Journal of Finance* 47, no. 4: 1259–1282.
- Longstaff, F. A., and E. S. Schwartz. 1995. "A Simple Approach to Valuing Risky Fixed and Floating Rate Debt." *Journal of Finance* 50, no. 3: 789–820.
- Loomis, Carl. "Robert Rubin on the Job He Never Wanted." *Forbes*, Nov. 28, 2007.
- Lynch Jr., J. J., and J. H. Mayle. 1986. *Standard Securities Calculation Methods: Fixed Income Securities Formulas*. New York: Securities Industry Association.
- Macaulay, F. R. 1938. *Some Theoretical Problems Suggested by Movements of Interest Rates, Bond Yields, and Stock Prices in the United States since 1856*. New York, Columbia University Press.
- Marcus, A. J., and I. Shaked. 1984. "The Valuation of FDIC Deposit Insurance Using Option-pricing Estimates." *Journal of Money, Credit, and Banking* 16, no. 4: 446–460.
- Matten, C. *Managing Bank Capital: Capital Allocation and Performance Measurement*. Singapore: John Wiley & Sons, 1996.
- McConnell, J. J., and M. Singh. 1994. "Rational Prepayments and the Valuation of Collateralized Mortgage Obligations." *Journal of Finance* 49, no. 3 (July): 891–922.
- McCulloch, J. H. 1975. "The Tax Adjusted Yield Curve." *Journal of Finance* 30, no. 3: 811–829.
- McDonald, L. G., and P. Robinson. 2009. *A Colossal Failure of Common Sense: The Inside Story of the Collapse of Lehman Brothers*. New York: Crown Business.
- Meissner, G., ed. 2008. *The Definitive Guide to CDOs*. New York: Incisive Media.
- Merton, R. C. 1970. "A Dynamic General Equilibrium Model of the Asset Market and Its Application to the Pricing of the Capital Structure of the Firm." Working Paper No. 497–70, A. P. Sloan School of Management, Massachusetts Institute of Technology. Note: Reprinted in *Continuous Time Finance*, Robert Merton. Cambridge, MA: Blackwell, 1993.
- Merton, R. C. 1973a. "An Intertemporal Capital Asset Pricing Model." *Econometrica* 41 (September): 867–887.
- Merton, R. C. 1973b. "Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science* 4 no. 1: 141–183.
- Merton, R. C. 1974. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance* 29, no. 2: 449–470.
- Merton, R. C. 1977. "An Analytic Derivation of the Cost of Deposit Insurance and Loan Guarantees: An Application of Modern Option Pricing Theory." *Journal of Banking and Finance* 1, no. 1: 3–11.
- Merton, R. C. 1978. "On the Cost of Deposit Insurance When There Are Surveillance Costs." *Journal of Business* 51, no. 3: 439–452.
- Mitsubishi Finance Risk Directory 1990/1991*. 1990. London: Risk Publications. Researched and compiled by *Risk Magazine*.
- Modigliani, F., and M. H. Miller. 1958. "The Cost of Capital, Corporation Finance, and the Theory of Investment." *American Economic Review* 48, no. 3: 261–297.
- Nelson, C. R., and A. F. Siegel, A.F. 1987. "Parsimonious Modeling of Yield Curves." *Journal of Business*, 60, no 4: 473–489
- Nielsen, L., J. Saa-Requejo, and P. Santa-Clara. 1993. "Default Risk and Interest Rate Risk: The Term Structure of Default Spreads." Working paper. INSEAD.

- O'Brien, J. M. 2000. "Estimating the Value and Interest Rate Risk of Interest-Bearing Transaction Deposits." Board of Governors of the Federal Reserve System, November.
- Office of the Special Inspector General of the Troubled Asset Relief Program. 2009. "Emergency Capital Injections Provided to Support the Viability of Bank of America, Other Major Banks, and the U.S. Financial System" (October 5).
- Patruno, G. N. 1994. "Mortgage Prepayments: A New Model for a New Era." *Journal of Fixed Income* 4, no. 3 (December): 42–57.
- Pearson, N. D., and T. Sun. 1994. "Exploiting the Conditional Density in Estimating the Term Structure: An Application to the Cox, Ingersoll and Ross Model." *Journal of Finance* 49, no. 4: 1279–1304.
- Pennacchi, G. G. 1987a. "Alternative Forms of Deposit Insurance: Pricing and Bank Incentive Issues." *Journal of Banking and Finance* 11, no. 2: 291–312.
- Pennacchi, G. G. 1987b. "A Reexamination of the Over- (or Under-) Pricing of Deposit Insurance." *Journal of Money, Credit and Banking* 19, no. 3: 341–360.
- Penter, P. M. 1989. *Splines and Variational Methods*. New York: John Wiley & Sons.
- Quigley, J. M. 1987. "Interest Rate Variations, Mortgage Prepayments and Household Mobility." *Review of Economics and Statistics* 69, no. 4 (November): 636–643.
- Ramaswamy, K., and S. Sundaresan. 1986. "The Valuation of Floating Rate Instruments: Theory and Evidence." *Journal of Financial Economics* 17, no. 2 (December): 251–272.
- Ranson, B. J. 2003. *Credit Risk Management*. Austin, TX: Sheshunoff Information Service.
- Rappaport, Liz. 2011. "New Fee to Bank Cash." *Wall Street Journal*, August 5.
- Raynes, S., and A. Rutledge. 2003. *The Analysis of Structured Securities: Precise Analysis and Capital Allocation*. New York: Oxford University Press.
- Rendleman, R., and B. Bartter. 1979. "Two-State Option Pricing." *Journal of Finance* 34, no. 5 (December): 1093–1110.
- Richard, S. F., and R. Roll. 1989. "Prepayments on Fixed-Rate Mortgage-Backed Securities." *Journal of Portfolio Management* 15, no. 3 (Spring): 73–82.
- Schwartz, E., and W. Torous. 1989. "Prepayment and the Valuation of Mortgage-Backed Securities." *Journal of Finance* 44, no. 2 (June): 375–392.
- Schwartz, H. R. 1989. *Numerical Methods: A Comprehensive Introduction*. New York: John Wiley & Sons.
- Selvaggio, R. D. 1996. "Using the OAS Methodology to Value and Hedge Commercial Bank Retail Demand Deposit Premiums." In *Handbook of Asset and Liability Management*, rev. ed., ed. F. Fabozzi and A. Konishi. New York: John Wiley & Sons.
- Shea, G. S. 1985. "Term Structure Estimation with Exponential Splines." *Journal of Finance* 40, no. 1: 319–325.
- Shimko, D. C. 1992. *Finance in Continuous Time: A Primer*. Miami, FL: Kolb.
- Shimko, D. C., N. Tejima, and D. R. van Deventer. 1993. "The Pricing of Risky Debt when Interest Rates are Stochastic." *Journal of Fixed Income* 3, no. 2 (September): 58–66.
- Shumway, T. 2001. "Forecasting Bankruptcy More Accurately: A Simple Hazard Model." *Journal of Business* 74, no. 1: 101–124.
- Slattery, M., and D. R. van Deventer. 2011. "Model Risk in Mortgage Servicing Rights." Kamakura blog, www.kamakuraco.com, December 5.
- Sobehart, J., S. Keenan, and R. Stein. 2000. "Validation Methodologies for Default Risk Models." *Credit* (May): 51–56.
- Standard & Poor's. 2002. "Ratings Performance 2002: Default, Transition, Recovery and Spreads."
- Stanton, R. 1995. "Rational Prepayments and the Valuation of Mortgage-Backed Securities." *Review of Financial Studies* 8, no. 3 (fall): 677–708.
- Stigum, M. 1991. *Money Market Calculations*. Homewood, IL: Irwin.

- Tavakoli, J. M. 2003. *Collateralized Debt Obligations and Structured Finance: New Developments in Cash and Synthetic Securitization*. Hoboken, NJ: John Wiley & Sons.
- Uyemura, D., and D. R. van Deventer. 1993. *Financial Risk Management in Banking*. Chicago/Cambridge, MA: Bankers/Probus.
- van Deventer, D. R. 1993. "The Valuation of Fixed and Floating Rate Mortgages." Kamakura Corporation research memorandum, December.
- van Deventer, D. R. 2011. "A Credit Crisis Chronology, Part 1 Through February 2008: This Time Isn't Different." Kamakura blog, www.kamakuraco.com (May 13). Redistributed on www.riskcenter.com, May 16, 2011.
- van Deventer, D. R. 2011. "Pitfalls in Asset and Liability Management: One Factor Term Structure Models." Kamakura blog, www.kamakuraco.com (November 7.) Reprinted in Bank Asset and Liability Management Newsletter, January.
- van Deventer, D. R. 2011. "Pitfalls in Asset and Liability Management: One Factor Term Structure Models and the Libor-Swap Curve." Kamakura blog, www.kamakuraco.com, November 23.
- van Deventer, D. R., and K. Imai. 1996. *Financial Risk Analytics: A Term Structure Model Approach for Banking, Insurance, and Investment Management*. Chicago: McGraw Hill.
- van Deventer, D. R., and K. Imai. 2003. *Credit Risk Models and the Basel Accords*. Singapore: John Wiley & Sons.
- van Deventer, D. R., and J. Outram. 2002. "The New Capital Accord and Internal Bank Ratings." In *Credit Ratings: Methodologies, Rationale and Default*, ed. M. K. Ong. London: Risk Publications.
- van Deventer, D. R., and X. Wang. 2002. "Basel II and Lessons from Enron: The Consistency of the Merton Credit Model with Observable Credit Spreads and Equity Prices." Working Paper, Kamakura Corporation.
- van Deventer, D. R., and X. Wang. 2003. "Advanced Credit Model Performance Testing to Meet Basel Requirements." In *The Basel Handbook: A Guide for Financial Practitioners*, ed. M. K. Ong. London: Risk Publications.
- van Deventer, D. R., and X. Wang. 2003. "Measuring Predictive Capability of Credit Models Under the Basel Capital Accords: Consecro and Results from the United States, 1963–1998." Working paper, Kamakura Corporation, 210.
- van Deventer, D. R. and S. Yasunaga. 2011. "A Credit Crisis Chronology, Part 2 from March 2008 Through March 2009: This Time Isn't Different." Kamakura blog, www.kamakuraco.com, May 14.
- Vasicek, O. A. 1977. "An Equilibrium Characterization of the Term Structure." *Journal of Financial Economics* 5: 177–188.
- Vasicek, O. A., and H. G. Fong. 1982. "Term Structure Modeling Using Exponential Splines." *Journal of Finance* 37, no. 2: 339–356.

- Academic term structure model. *See*
 - Term structure models
- Accrued interest, 60–61
- Accuracy, ratings, 422, 429–431
- Accuracy ratio, 379–380, 381
- Actual/360 basis, 549
- Actual/365 basis, 549
- Actuarial tables, 710
- Actuaries, 8
- Adams and van Deventer, 54, 74, 75,
 - 80, 101, 103, 112, 116, 122, 123, 142, 615
- Adjustable-rate mortgages (ARMs), 647
- Affine term structure model, 239
- Agency ratings, 419
- Agius, Marcus, 548
- AIG (American International Group),
 - xxxi, xxxii, 47, 428, 729–730, 737–744
- Allowance for loan losses, 724
- Alt-A loans, 727
- Ambac, xxxi
- American fixed income options. *See*
 - Fixed income options
- American International Group (AIG).
 - See* AIG (American International Group)
- American swaptions, 579
- Ameriquest Mortgage, xxix
- Amin and Jarrow model, 675–678, 680,
 - 683–684, 703–704
- Amortizing loans/bonds, 63–64, 461,
 - 613–615, 626–636
- Analytical approximations, 597
- Andersen, Leif, 75, 706
- Annualized default probabilities, 392
- Araten, Michel, 489
- Arbitrage free value, 584
- Arrears swaps, 580–586
- Asia crisis, 758
- Asset and liability management (ALM),
 - 5, 7–8, 640–643
- Asset-backed securities (ABSs), 480. *See also* Mortgage-backed securities
- Asset classes, 3
- Average spread per dollar loan, 417
- AXA, xxx
- Ayache, E., 706
- Bailouts, xxvii, 13, 45, 566, 736, 754.
 - See also* Credit crisis of 2006–2011
- Balance sheets, 8–9, 12
- Bank bill futures. *See* Sydney Futures Exchange Bank Bill Contract
- Bank deposit insurance, 596
- Bank deposits. *See* Nonmaturity deposits
- Bank failures, 31
- Bank for International Settlements (BIS),
 - 39, 783
- Banking industry, overview, 3–6
- Bank of America, xxvii, xxx, xxxii,
 - xxxiv, 4, 13, 21–24, 369, 370, 649, 694, 759, 784, 786
- Bank of America Merrill Lynch, 341
- Bank of England, xxx, 6, 39, 669–670
- Bankruptcy, 17–18, 596
- Bankruptcy risk, 700, 704–707
- Barclays, xxx, 341, 548, 650
- Barings, 759
- Basel Accords, 39, 40, 375, 425, 791
- Basel capital regulations, 374
- Basel Committee on Banking
 - Supervision, 7, 40, 231, 243, 374, 783
- Basis risk, 30
- BB&T, 649
- Bear Stearns, xxvii, xxix, xxxi, xxxii, 4,
 - 47, 742–746

- Beck, David, 742
- Benchmarking, 17–18
- Bermudan option, 600
- Bermudez, A., 706
- Bharath, S., 11, 367, 375, 378, 381, 382, 438, 439
- Bias in default probabilities, 379, 383–385
- The Big Short* (Lewis), 655, 736
- Binomial lattices, 619
- Black, Derman, and Toy model, 139, 146, 162, 239, 305, 306
- Black, Fischer, 9
- Black and Karasinski model, 178, 239, 306
- Black model, xxvi, 549–550
- Black-Scholes model, 9, 28, 285, 317, 360, 440, 445, 510–511
- Blum, M., 375
- BNP Paribas, xxx, 341
- Board of directors, Citigroup, 7
- Board of Governors of the Federal Reserve, xxiv, 31, 38, 39
- Bond market liquidity, 364
- Boral, A., 375, 376. *See also* Falkenstein and Boral test
- Boston, xxviii
- Breeden, Douglas T., 643, 645, 646
- Breeden, Smith, 643
- Brennan, Michael J., 306
- Brennan and Schwartz model, 306–307
- British Army, 4
- British banking industry, xxxi
- Brownian motion, 363, 394
- Bubbles:
- housing, 437
 - Japanese, xxiii, xxiv, 10
- Bullet bond, 461
- Burffum, Dan, 706
- Burnout, 645
- Bush, George W., xxxiii, xxxiv
- Bushy trees, 125–126, 596
- in credit risky bond valuation, 455–461, 467–470
 - examples, 143, 164, 170
 - with foreign currency derivatives, 678
 - in one-factor models, 129–137, 150–153, 155–156
 - in three-factor models, 206–217
 - in two-factor models, 173–178
- Buy low/sell high, 31, 766–767
- Callable bonds, 505–508, 613
- Calpine, 768
- Campbell, John Y., 11
- Canadian banking industry, xxx
- Capital:
- primary, 38
 - risk-adjusted, 34–35
 - role in capital and business strategy, 32–35
- Capital allocation, xxvi, 7, 9, 728–731
- Capital asset pricing model, 700, 702
- Capital management, regulation of, 37–40
- Capital Purchase Program (CPP), xxxiii, 751
- Capital regulation, 26
- Capital structure, 449
- Caplet, 550
- Caps, 398
- background on, 548–550
 - common uses, 549
 - as European options on forward rate agreements, 550
 - forming cap and floor derivatives, 550
 - measuring counterparty credit risk on, 565–566
 - valuing, 550–554
 - valuing a floating-rate loan with a cap, 557–563
 - variations on, 565
- Case-Shiller home price index, xxix, 653, 683
- Cash-flow CDOs, 480–481
- Cash flow simulation, 258
- Cathcart, Ron, 742
- CCAR. *See* Comprehensive Capital Analysis and Review 2012 (CCAR)
- CDOs. *See* Collateralized debt obligations (CDOs)
- CDSs. *See* Credit default swaps (CDSs)
- CDS spreads, 786–787
- Certificates of deposit (CDs), 12
- Charities, 6

- Chava, Sudheer. *See* Jarrow-Chava model
- Cheapest-to-deliver option, 528–529, 546
- Chen, Derek H., 528, 530
- Chen, Ren-raw, 306–307, 308, 310
- Chen's model, 307–308
- Chicago Mercantile Exchange, 513
- Chief executive officer, 3
- Chief financial officer, 4, 5
- China, and credit crisis of 2006–2010, xxxi
- Chinese walls, 8, 714
- Chrysler, xxxiv
- Chu, Ting-heng, 650, 655
- CIR (Cox, Ingersoll, and Ross) model, 239, 304–306
- Citibank, 341
- Citigroup, xxvii, xxx–xxxiv, 4, 5, 6–8, 474, 488, 649, 700, 701, 704, 705, 720–723, 784, 786
- Clean price, 61
- Closed form, xxvii
- Collars, 550
- Collateral, in CDOs, 480
- Collateralized debt obligations (CDOs), xxvi, 338, 480–490
- cash flow, 480–481
- copula method, 487–490
- disappearance of, 639
- as example of reduced form modeling, 483–486
- expected losses on, 19–20
- Jarrow-Merton put option framework, 490
- market concentration and, 479
- regulation in, 478–480
- risk drivers, 481–482
- steps in valuation, 481–483
- synthetic, 480–481
- waterfall, 481, 482
- Collateralized mortgage obligations (CMOs), 647
- Collateral values, 682–693
- Collin-Dufresne, Pierre, 418
- Colossal Failure of Common Sense*, A (McDonald and Robinson), 4
- Commercial banks, 6, 8
- interest rate risk management, 253–254
- performance measurement, 9
- reasons for failure of, 31
- risk management in, 29–35
- transfer pricing, 20–26
- yes/no risk measures, 30
- Commercial paper, 694
- Commercial Paper Funding Facility, 694, 736, 758
- Commercial paper purchases, 736
- Commercial real estate price index, 482
- Commodities, 251
- Common stock. *See* Equities; Equity returns
- Company assets valuation, 446–447, 476
- Compensation packages, 6
- Compensation systems, 725
- Components of performance, 778, 779
- Compound credit risk, 359–360
- Compound default risk, 360
- Compound interest, 66–68
- Comprehensive Capital Analysis and Review 2012 (CCAR), 4
- Comptroller of the Currency, 4
- Computer simulations, xxvi–xxvii
- Computer storage capacity, 10
- Concentration of risk, 481
- Conflicts of interest, 7, 422
- Constant interest rate volatility, 142
- Constant prepayment rate, 640–641
- Constant prepayment speed models, 644–645
- Consumer behavior modeling. *See* Irrational behavior
- Continental Illinois, 335
- Continuously compounded interest rate, 66–67
- Conventional mortgages, 650
- Conventional wisdom, 648
- Convertible bonds, 706–707
- Convexity, 267–268, 278–281
- Cooperative financial institutions, 6
- Copula model, xxvi, 448–449, 473, 736
- Core deposits, 725
- Corporate strategy clarity, 339–340

- Correlation:
 in credit default swaps, 476–477
 in default events, 448–449
 in default intensities, 476
 in default probabilities, 372–373
- Corruption, in LIBOR swap curve, 414–415
- Costless collars, 550
- Costs of origination, 766
- Counterparty credit risk, 373, 477, 565–566
- Countrywide, xxix, xxx, xxxii, 694
- Coupon-bearing bonds, 116, 140
 with embedded European call option, 503–508
 European call option on, 501–503
 futures on, 528–529
 with one-factor models, 156–158
 options on, 511
 with three-factor models, 225–226
 with two-factor models, 184, 187
- Coupon calculation, 62–63
- Cox, D. R., 708
- Cox, Ingersoll, and Ross (CIR) model, 239, 304–306
- Cox process, 708
- Craig, Susanne, 14
- Credit crisis of 2006–2011, xxiii, 10, 13, 473–474. *See also* Collateralized debt obligations (CDOs); Credit default swaps (CDSs)
- American International Group (AIG), 738–744
- and CDO market, 486–487
- chronology of events, xxviii–xxxiv, 742–744
- consolidated JPMorgan Chase, Bear Stearns, and Washington Mutual, 744–746
- contributing risk fallacies, 735–736
- data on top borrowers, 738–739
- global corporate average cumulative default rates, 432–434
- home prices during, 639, 682–693
- implications for liquidity risk management and analysis, 758–763
- legacy ratings problems, 431–437
- lessons from, 487–490
- Morgan Stanley, 749–751
- mortgage prepayments, 625, 640–641, 646
- regulations stemming from, 340–341
- State Street, 746–749
- stress testing, 425
- Credit default swap (CDS) spreads, 415, 786–787
- Credit default swaps (CDSs), 247, 350–351, 474–480
 auction protocol, 452
 as collateral in CDOs, 480
 comparing performance with, 391
 correlation in, 476–477
 digital default swap, 475–478
 and life insurance premiums, 709–710
 market manipulation and, 341–347
 municipals and sub-sovereigns, 352
 prices of, 337, 339
 as self-insurance, 729–730
 single-name, 478–479, 480
 sovereign credits, 353–357
 steps in valuation, 481–483
 super-senior tranches, 489
 trading volume implications, 357–358
 traditional structure, 474–475
- Credit derivatives, 473–492. *See also* Collateralized debt obligations (CDOs); Credit default swaps (CDSs)
- Credit derivatives pricing, 421
- Credit facilities. *See* Revolving credit
- Credit model testing, 375–391
 consistency of estimated and actual defaults, 383–386
 Falkenstein and Boral test, 383–386
 and interest rates, 388–391
 Jarrow-Chava model, 380–381
 with market data, 388
 misunderstandings about, 376–378
 vs. naïve models of risk, 386–388
 performance components, 378–380
 reduced form model *vs.* Merton model, 381–382
 ROC accuracy ratios, 387–388
 use of market data, 388
- Credit protection buyer/provider, 474–477

- Credit ratings, 390. *See also* Legacy ratings
accuracy of, 429–431
usefulness of, 785–786
- Credit risk, xxvi, 3–13, 335–358, 648.
See also Credit risk management;
Market prices for credit risk analysis
compound, 359–360
counterparty, 373, 477, 565–566
fluctuating, 696
impact on foreign exchange risk
formulas, 681
and irrational prepayment behavior, 637–638
and liquidity risk linkages, 759–760
market data sources for, 336–339
ordinal measure of, 422
- Credit risk management, 7. *See also*
Credit model testing; Legacy ratings
- Credit risk museum, 359
- Credit risk premium, 416–420
- Credit risky bond valuation, 453–472
amortizing loan, 461
bonds with credit risk, 465–472
bonds with no credit risk, 454
bullet bond, 461
floating rate loans, 465
simulating future values, 454–455, 471–472
- Credit scoring, 367
- Credit spread levels, 415
- Credit spreads, 388, 477–478, 770
comparing performance with, 391
data problems with risky issuers, 413–415
fitting and modeling, 396–420
fitting with cubic splines, 409–410
market convention for, 397–398
maximum smoothness forward rate technique, 410–411
model-independent, 398–404
risk factor limits (RFLs), 5
- Credit spread smoothing, 396–397, 404–409
- Credit Suisse, xxx, 341
- Credit supply and demand, 416–420
- Crump, Richard K., 327, 332
- Cubic credit spread smoothing, 409–410, 411–413
- Cubic forward rate function, 73
- Cubic splines, 94–101, 115, 409–410
- Cubic yield function, 73
- Cumulative probability of local maximums and minimums, 52
- Cyclical, 422
- Dai, Qiang, 315
- Dana Corporation, 768
- Data aggregation, 266
- Declining interest rate volatility, 142
- Defaultable European call options, 361
- Defaultable interest rate swaps and swaptions, 579
- Defaultable options on forward and futures contracts, 546–547
- Default consistency, 385
- Default correlation, 336
- Default intensities, 475
converting to discrete default probabilities, 391–394
correlation in, 476
- Default probabilities, 769, 785–786. *See also* Credit ratings; Merton default probabilities; Quantitative default probabilities
bias in, 379, 383–385
conversions, 391–394
in copula approach, 489–490
correlations in, 372–373
estimating, in Merton model, 445–446
in life insurance, 708–712
mapping to actual defaults, 447
sufficient statistic for, 438–439
- Default probability of the institution, 57, 732
- Default rates by ratings grade, 428–429
- Default risk:
in fixed income options, 620
and performance measurement, 768–769
- The Definitive Guide to CDOs* (Meissner), 488, 490
- Delinquency rates, 5
- Delivery option, 528, 529
- Dell, 479–480

- Delphi, 768
- Deposit insurance, 12
- Depository Trust and Clearing Corporation (DTCC), 341–342, 346, 347
- Deposit run-off, 670
- Derivatives, 567. *See also* Credit derivatives; Exotic derivatives
- Derman, Emanuel. *See* Black, Derman, and Toy model
- Deutsche Bank, 341
- Dexia Credit Local, 751
- Dexia SA, 751–758
- Dickler, Daniel T. *See* Dickler, Jarrow, and van Deventer
- Dickler, Jarrow, and van Deventer, 46, 47, 54, 123, 125, 140, 142, 162, 190, 191, 231, 236, 238, 239, 269, 302, 304
- Digital default swap, 475–478
- Digital options:
with exotic derivatives, 586
in one-factor models, 158
on one-year U.S. Treasury rate, 140
in three-factor models, 227–229
in two-factor models, 185, 187–189
- Digital range notes, 588, 589–591
- Dirty price, 61
- Discount factors, 62
- Discrete default probabilities, 391–394
- Distance to default, 428–429, 441–442
- Diversification, 5
- Dividend payments, 597
- Dodd-Frank, 425
- “Do no harm,” 800–801
- Dothan, Uri, 305
- Drift, 239
- Drum, Kevin, 491
- Due diligence, transaction-level, 639
- Duffie, Darrell. *See* Duffie and Singleton
- Duffie and Singleton, 11, 286, 361, 362, 380, 394, 396, 708
- Duration, 3, 39, 267–278
bond math market convention, 273–276
conventional vs. Macaulay’s, 278
Macaulay’s duration, 258, 268–270, 273, 278, 283
modified, 276–277
overview, 267–268
as term structure model, 288–292
- Duration hedging, 238, 242, 270–273, 278–281
- Dynamic programming, 600
- Earthquake, 12
- Effective maturity of equity, 259
- 80/20 rule, 257
- Ellis, David M., 673, 674
- Embedded call option, 613
- Embedded European call option, 503–508
- Emergency Economic Stabilization Act (EESA), xxxiii
- Empirical probabilities of interest rate movements, 235–236
- Enterprise risk management, 801
- Enterprise risk system, 653
- Equities, 251
default risk and performance measurement, 768–769
prices of, as source of market data, 338–339
risk management, 28
- Equity returns, 700–707
convertible bonds, 706–707
credit-risk adjusted approach, 703–704
derivatives approach, 702–703
options, 704–706
traditional fund management approach, 701–702
- Eurodollar futures, 513, 529–530
- Eurodollar futures-type forward contracts, 524–527, 540–545
- European call options, 361, 367
- European options, 454–455, 625–626
- European options on bonds, 495–512
call option on coupon-bearing bond, 501–503
coupon-bearing bond with embedded European call option, 503–508
Jarrow-Merton put option, 512
one-factor Vasicek model, 509–511
options on coupon-bearing bonds, 511
options on defaultable bonds, 509

- European options on forward and futures contracts, 531–547
defaultable options, 546–547
Eurodollar futures-type forward contracts, 540–545
forward contracts on zero coupon bonds, 532–538
forward rate agreements, 538–540
futures on coupon-bearing bonds, 546
importance of, 531
money market futures contracts, 546
notations and useful examples, 531–532
European swaptions, 578–579
Euroyen futures, 529–530
Event of default, 427, 448
Exact day count, 267
Exchange rates, 482
Existing futures positions, 527–528
Exotic derivatives, 580–595
 arrears swaps, 580–586
 credit risk with, 594–595
 digital options, 586
 digital range notes, 588, 589–591
 range floaters, 588, 592–593
Explanatory variables, 373, 380, 387, 392, 394, 395
Extended Merton model. *See* Ho and Lee model
Extended Vasicek model. *See* Hull and White model
Extensions of credit, 737, 740–741, 744–746, 748, 751, 752, 755
F9 Model Monkeys, 473, 481, 486, 488, 509, 532
Failed ratings, 437
Failure:
 admitting possibility of, 783–785
 controlling probability of, 789–790
 managing probability of, 785–789
Falkenstein, Eric, 376, 386. *See also* Falkenstein and Boral test
Falkenstein and Boral test, 377–379, 383–386, 447
Fallacies in risk management, xxiii–xxviii
False VaR, 730
Fannie Mae (FNMA), xxxii, 47
FDIC Loss Distribution Model, 31, 40, 359
Federal Reserve, xxiv, xxxi–xxxiii, 38
 and credit crisis (*see* Credit crisis of 2006–2011)
Federal Reserve Bank of St. Louis, xxviii
Federal Reserve Board, 657
FHLMC, 47. *See* Freddie Mac
Financial accounting, 4
Financial Accounting Standards (FAS), 60
Financial agency regulation, 340–341
Financial bailout. *See* Bailouts
Financial Crisis Inquiry Commission, 6
Financial Deregulation and Monetary Control Act, 38
Financial Institutions Monitoring System, 31
Financial Services Authority, 513
Finite difference methods, 618–619
Firm-Wide Risk Committee, xxix
First Chicago, 335
First Horizon, 649
First Interstate Bancorp, 24–26, 388–391
First to default swaps, 337–338, 481–483
Fisher-Weil duration, 273
Fishman, George S., 565
Fitting interest rate volatility, 317–332
Fixed coupon bonds, 62–63, 70–71
Fixed income derivatives, 398. *See also* Fixed income options
Fixed income mathematics, 59–72
Fixed income options, 596–621
 background, 596–597
 bank deposit insurance, 596
 Bermudan option, 600
 binomial lattices, 619
 callable bond with a three-factor HJM bushy tree, 598–613, 614–615
 finite difference methods, 618–619
 HJM valuation when default risk is present, 620
 Monte Carlo simulation, 615–618

- Fixed income options (*Continued*)
 rationally prepaid amortizing loan, 613–615
 trinomial lattices, 619–620
 valuation approaches, 597
- Fixed income performance attribution, 772–780
- Fixed income securities, 251. *See also*
 Fixed income options
- Fixed rate mortgages, 12
- Floating rate bonds, 64
- Floating rate loan, 465, 557–563
- Floating rate mortgages, 3, 21, 59
- Floorlets, 550
- Floors, 398
 background on, 548–550
 common uses, 549
 forming cap and floor derivatives, 550
 measuring counterparty credit risk on, 565–566
 valuing, 554–557
 variations on, 565
- Fluctuating credit risk, 696
- Ford Motor Company, 488
- Foreign currency derivatives:
 forwards, 676–677
 implications of term structure model-based formula, 680–681
 legacy valuation approaches, 678–680
 numerical valuation methods, 677–678
- Foreign exchange markets, 675–681
- Foreign exchange risk, 453
- The Formula That Killed Wall Street* (Salmon), 491
- Forsyth, P. A., 706
- Forward bond prices, 62
- Forward contracts, 513–527
 Eurodollar futures-type, 524–527
 European options on (*see* European options on forward and futures contracts)
 forward rate agreements, 520–524
 margin requirement, 514, 520
 notations and useful examples, 531–532
 open interest in, 513–514
 on zero coupon bonds, 514–524
- Forward interest rates, 69–70
- Forward rate agreements, 520–524
 caps as European options on, 550
 European options on, 538–540
- Forward rates, 47, 49, 50, 51, 53–55, 62, 80
- Forward rate shifts. *See* Empirical probabilities of interest rate movements
- Forward rate smoothing. *See also* Yield curve smoothing
 maximum smoothness forward rate technique, 101–111
 smoothness vs. curve length, 112–116
- Forward rate volatility, 142–146
- Forward zero coupon bond prices, 70–71
- Foundations, 6
- Freddie Mac, xxx, xxxii, 428, 749, 792, 794
- Fukushima nuclear disaster aftermath, 714
- Fuld, Dick, 4
- Funaro, Steve, 742
- Funding costs, 415
- Futures contracts, 527–530
 cheapest-to-deliver option, 528–529, 546
 Eurodollar and Euroyen futures contracts, 529–530
 European options on, 546
 existing futures positions, 527–528
 notations and useful examples, 531–532
 on zero coupon bonds, 527–528
- Future values, 66–67
- GAAP (generally accepted accounting principles), 234, 235, 236, 567, 653, 801
- Gap analysis, 257, 258–263
- Garman, Mark B. *See* Garman-Kohlhagen model
- Garman-Kohlhagen model, 678–679, 680
- Gaussian random walk. *See* Random walk
- GDP growth, 482, 485
- Gelband, Mike, 4

- General Growth Properties, 768
Generally accepted accounting principles (GAAP). *See* GAAP (generally accepted accounting principles)
General Motors (GM), xxxiv, 488
Germany, xxxiii
Global Association of Risk Professionals, 452
“Global CDO Issuance,” 486–487
Goldman Sachs, xxix, xxxiii, 342, 786
Government of Singapore Investment Corporation, xxxi
Granularity of ratings, 422
Greece, 418
Greenspan, Alan, xxiv
- Hardware, 453
Hazard rate, 475
Hazard rate modeling, 367
HBOS PLC, xxvii, xxxiii, 784
Heath, David, 9–10, 123
Hedge funds, subprime, xxix
Hedging, 3, 8
 duration hedging, 238, 242, 270–273, 278
 foreign currency derivatives, 680
 implied irrationality and, 636–637
 of interest rate risk, 250–256
 macro factor risk, 364
 of mortgage servicing rights, 653–654
 performance hedging, 389–390
Hedging instruments, 725, 733
Hewlett-Packard (HP), 424–425
High default risk equity portfolios, 701
Hilliard, Jimmy E. *See* Hilliard, Madura, and Tucker (HMT) model
Hilliard, Madura, and Tucker (HMT) model, 678–680
Hilscher, Jens, 11, 429–431
Historical VaR, 18–19, 726
Historical volatility, 317
HJM model, 9–10
 bushy trees, 129–137 (*see also* Bushy trees)
 coupon-bearing bond valuation, 140, 156–158, 184, 187, 225–226
 digital option valuation, 140, 158, 185, 187–189, 227–229
 duration in, 288–292
 with European options of forward and futures contracts (*see* European options on forward and futures contracts)
 with European options on bonds (*see* European options on bonds)
 in fixed income options, 596, 620
 and foreign exchange derivatives, 675–681
 forward rate volatility assumptions, 142–146
 hedging, 271–273
 introduction to, 123–141
 Ito’s lemma in, 286–288
 key implications and notation, 129–131, 146–147, 167–170, 201–205
 Merton’s model and, 293–300
 model assumptions and key input data, 124–129, 142–146, 162–163, 191–192
 and Monte Carlo simulation, 236–238
 for nonmaturity deposits (*see* Nonmaturity deposits)
 with one factor (*see* One-factor models)
 overview, 308
 probability-weighted discount factors, 138–139
 pseudo probabilities, 131–132, 147, 171–178, 205–217
 with rate- and maturity-dependent volatility, 143–160
 in risk-free bond valuation, 454–455
 special cases, 283–313
 with three risk factors (*see* Three-factor models)
 with two risk factors (*see* Two-factor models)
 valuation in, 139, 153–156, 178–186, 217–229
 Vasicek model and, 292–293
Ho, Thomas S. Y. *See* Ho and Lee model
Ho and Lee model, 128, 143, 239, 285, 298–300, 308–310
Holding company transfer pricing system, 25

- Home mortgage loans, 615. *See also* Mortgages
- Home prices, 13
- housing bubble, 437
- impact on collateral values, 682–693
- role in defaults and prepayments, 652–653
- U.S. vs. Japan, xxv
- Hosmer, D. W., 369, 394
- Housing bubble, 437
- HSBC, xxix, xxxi, 342
- Hu, Guizhou, 710
- Huang, Jing-zhi, 418
- Hull, John, 440, 549, 616. *See also* Hull and White lattice approach; Hull and White model
- Hull and White lattice approach, 619–620
- Hull and White model, 143, 239, 303, 313–314, 596
- Hybrid credit models, 381
- Hypo Real Estate AG, xxxiii
- Hypo Real Estate Nationalization Bill, xxxiv
- IKB Deutsche Industriebank AG, xxx
- Implied forward interest rates, 69–70
- Implied irrationality, 636–637
- Implied value, 446–447
- Implied volatility, 317
- Inaccurate rating models, 436
- Industrial Bank of Japan, 335
- Industrial corporations, 6
- IndyMac Bank, xxxii, 743
- Inflation rates, 481
- Information technology, 7, 793–799
- legacy systems, 793–795
- Request for Proposal (RFP) process, 795–796
- software installation, 797
- software vendors, 798–799
- Infrastructure upgrading, 795–796
- Ingersoll, Jonathan E., Jr. *See* Cox, Ingersoll, and Ross model; Ingersoll, Skelton, and Weil
- Ingersoll, Skelton, and Weil, 268, 282
- ING Group, 759
- Innovation, pace of, 806
- Institutional barriers to progress, 8–11
- Institutions, defined, 250
- Insurance. *See* Life insurance; Property and casualty insurance
- Insurance companies, 6
- interest rate risk management, 252, 266–267
- risk management in, 8
- Insurance premiums, 475
- Integrated enterprise risk management, 801
- Integrated risk, 640, 714
- Interest compounding, 66–68
- Interest rate caps. *See* Caps
- Interest rate crisis, 3, 10, 673
- Interest rate derivatives, 10, 548, 550, 565
- Interest rate floors. *See* Floors
- Interest rate modeling. *See* HJM model
- Interest rate movements, empirical probabilities of, 235–236
- Interest rate risk, xxvi, 30
- analyzing, step-by-step, 55–58
- hedging, 250–256
- measuring, 26–27, 258
- overview, 45–58
- safety zone, 58, 250
- Interest rate risk management: commercial banks, 253–254
- insurance companies, 252
- legacy approaches to (*see* Legacy approaches to interest rate risk management)
- obvious decisions, 255
- pension funds, 251–252
- pitfalls in, 238–246
- political factions in, 251–254
- Interest rates: actual 360/365 basis, 549
- and forward rate volatility, 142–146
- implied, 69–70
- mismatching, 250–256
- parallel shifts in, 269–271
- risk management, 28
- upshifts/downshifts in, 126, 129–131
- Interest rate sensitivity, 258
- Interest rate sensitivity gaps, 27, 258–263

- Interest rate swaps, 414, 567–579
 basics, 567–568
 daily margin requirements, 568
 defaultable, 579
 observable fixed rate, 574
 prices, 338
 swaptions, 574–579
 valuing, 568–574
- Interest rate volatility, 126–129, 142. *See also* Forward rate volatility
- Interest rate volatility fitting, 317–332
- Interest rate volatility function, 126
- Interest rate volatility surface, 200
- International Association of Credit Portfolio Managers, xxvi
- International Financial Reporting Standards (IFRS), 60
- Interpolated swap rates, 650
- Intertemporal capital asset pricing model, 445
- Irrational behavior, 622–637
 amortizing loan example, 626–636
 and credit risk, 637–638
 historical research into, 624, 637
 implied irrationality and hedging, 636–637
 transactions cost approach, 624–625
- Irrevocable line of credit, 758, 760, 761
- Iteration, 635
- Ito's lemma, 286–288, 289, 450–451
- ITT Corp., 5
- Jamshidian, Farshid. *See* Jamshidian formulation
- Jamshidian formulation, 510–511
- Janosi, Tibor, 76, 91, 103, 106, 658, 672
- Japanese bubble, xxiii, xxiv, 10
- Japanese government bond (JGB)
 contract, 528–529
- Jarrow, Robert, 9–10, 31, 76, 91, 103, 106, 123, 364, 426, 428, 491, 658, 672, 709–710.
 See also Dickler, Jarrow, and van Deventer
- Jarrow, van Deventer, and Wang, 238, 239, 241, 378, 390
- Jarrow and Turnbull model, 359–361, 373–374, 440, 476, 703–706
- Jarrow-Chava model, 380–381, 382–386, 419–420, 744, 747
- Jarrow-Merton put option, 251, 255, 258, 364–365, 391, 453, 490, 548, 596, 714, 731–733
 for capital allocation, 780–782
 example, as measure of total risk, 719–723
 and legacy ratings, 437–439
 and liquidity risk, 760–761
 valuing, 472
- Jarrow model, 361–372, 475. *See also* Credit model testing
 fitting to credit derivative prices, 366
 fitting to current and historical price data, 366
 fitting to debt prices, 365–366
 summary, 373–374
 zero coupon bond prices in, 363–365
- Jenkins, Huw, xxx
- Jones, Sam, 491
- Jordan, James, 673, 674
- Jorion, Philippe, 726
- Journal of Investment Management*, 438
- JPMorgan Chase, xxxi, xxxiii, 246–247, 339, 342, 477–478, 546, 649, 742–746, 786
- Kamakura Corporation, xxviii, 46, 49, 50, 51, 375, 424, 736–737
- Kamakura default probabilities (KDP). *See* KDP models
- Kamakura Risk Information Services (KRIS), 347, 367, 426–428
- Kamakura Risk Manager, 123, 124, 143, 404, 672
- Karasinski, Piotr, 146. *See also* Black and Karasinski model
- KDP models, 383–385, 784
- Keenan, S., 376
- Keynes, John Maynard, xxiii–xxiv, xxvi
- Killinger, Kerry, 744
- King, Gary, 715
- King, Mervyn, 6
- Klein, Sean Patrick, 144

- Kohlhagen, Steven W. *See* Garman-Kohlhagen model
- Korean Development Bank, 675
- Lando, David, 389, 394, 447, 448, 699
- Las Vegas, xxix
- Lattice methods, 618, 619–620
- Lee, Sang-Bin. *See* Ho and Lee model
- Legacy approaches:
- to credit risk management, 421–452 (*see also* Legacy ratings)
 - with foreign currency derivatives, 678–680
 - option-adjusted spread, 645–647
- Legacy approaches to interest rate risk management, 256, 257–282
- convexity, 267–268
 - duration, 267–278
 - gap analysis, 257, 263–264
 - interest rate sensitivity gap analysis, 258–259
 - maturity structure modeling, 267–268
 - multi-period simulation, 264–267
 - safety zone, 258, 259–263
- Legacy prepayment analysis, 643, 644
- Legacy ratings:
- accuracy of, 422, 429–431
 - and credit crisis of 2006–2010, 431–437
 - default rates, 422, 428–429
 - D rating, 427
 - failed ratings, 437
 - granularity, 422
 - and Jarrow-Merton put option, 437–439
 - legal pressure for AAA ratings, 437
 - moral hazard in self-assessment, 426–429
 - overview, 421–422
 - point-in-time vs. through-the-cycle, 423–425
 - SD rating, 427
 - transition matrices, 426
- Legacy systems, 793–795
- Lehman Brothers, xxxii, 4, 32, 38, 47, 54, 254–255, 367–368, 428, 594–595, 700, 744–751, 758, 763, 768, 769, 784, 795
- Lehman Brothers Government Bond Index, 18
- Lemeshow, S., 369, 394
- Level 3 assets, 654
- Levin, Carl (Senator), xxviii
- Levin Report, xxviii, 431, 435–437
- Lewis, Michael, 655, 736, 768
- Li, David, 488, 490
- LIBOR, 10, 123–124, 346
- and bond valuations, 64
 - swap curve, 22
- LIBOR manipulation, 341, 346, 413–415, 513, 546, 548, 650–651, 804
- LIBOR market model, 548
- LIBOR panel, 649
- Life insurance, 708–712
- cancellation privileges, 711
 - credit default swaps (CDSs) and life insurance premiums, 709–710
 - mortality table, 710
 - valuing policies, 711–712
- Life insurance companies, 252
- Life insurance funds, 6, 8
- Life insurance policies, 624
- Lin, Che-chun, 650, 655
- Linear forward rate function, 73
- Linear yield function, 73
- Lines of credit. *See* Revolving credit
- Liquidity:
- in bond market, 364 (*see also* Liquidity risk)
 - case studies in liquidity risk, 735
- Liquidity analysis, 735–764. *See also* Liquidity risk
- Liquidity crises:
- AIG, 737–744
 - consolidated JPMorgan Chase, Bear Stearns, and Washington Mutual, 744–746
 - Dexia SA, 751–758
 - Morgan Stanley, 749–751
 - State Street, 746–749
- Liquidity events, 758–759
- Liquidity premium, 396–397, 415
- Liquidity puts, 6, 7
- Liquidity risk, xxvi, 3–13
- and credit risk linkages, 759–760

- managerial behavior and market funds supply, 761–763
- measuring as a line of credit, 760–761
- optimal liquidity strategy, 763
- top five liquidity crises (*see* Liquidity crises)
- Liquid market assumption, 449
- Lloyds, xxxiii
- Logarithm, 418
- Logistic functions, 483–484
- Logistic regression, 367–372, 641–642, 644, 648, 710
- Lognormal distributions, 702–704
- London interbank offered rate (LIBOR). *See* LIBOR
- London Whale, 247
- Long Beach Mortgage Corporation, 742
- Long-Term Capital Management (LTCM), 13
- Loss distribution model, 31, 40, 359
- Loss given default, 448, 642
- Macaulay, Frederick R., 39, 268
- Macaulay's duration, 39, 238, 258, 268–270, 273, 283
- MacKenzie, Donald, 491
- Macro factors, 4, 6, 364, 725, 728, 730, 732, 733
- Madura, Jeff. *See* Hilliard, Madura, and Tucker (HMT) model
- Maiden Lane I, 743
- Mapping default probabilities, 447
- Marginal cost, 417
- Marginal revenue, 417
- Margin requirements:
 - in forward contracts, 514, 520
 - in swap markets, 568
- Market-based equity ratio, 57, 732
- Market capitalization, 723, 725
- Market concentration measures, 479
- Market data, in credit model testing, 387–388
- Market prices for credit risk analysis, 335–336
 - and clarity in corporate strategy, 336, 339–340
 - credit default swaps (CDSs), 337, 339, 341–347, 350–351
 - data sources for, 336–339
 - and institutional safety and soundness measurement, 336, 340–341
 - pricing accuracy, 336, 339–340
 - and risk management sophistication, 336, 340
- Market risk, xxvi, 3–13
- Market segmentation, 416
- Market value of portfolio equity, 57, 732
- Markowitz, Harry, 438
- Markowitz Award, 438, 701
- Mark-to-market approach, 3–6
- Mark-to-market performance, 769, 771–772
- Matched maturity transfer pricing system, 22–24
- Mathematical models, xxvi–xxvii
- Matten, Chris, 35, 765
- Maturity structure modeling, 267
- Maximum smoothness forward credit spread smoothing, 411–413
- Maximum smoothness forward rates, 101–111, 117–120, 400–403, 410–411
- McConnell, John J., 624, 647
- McDonald, Larry, 4, 513
- Mean reversion, 239
- Measuring returns, 20
- Meissner, Gunter, 488, 490
- Merit Financial, Inc., xxix
- Merrill Lynch, xxvii, xxx–xxxii, 4–5, 474, 488, 726, 768
- Merton, Robert, 15–16
- Merton default probabilities, 376, 381, 391, 447
- Merton model, 361–362, 441–444. *See also* Credit model testing
 - and early default, 447–448
 - loss given default in, 448
 - and random interest rates, 447
 - vs.* reduced form models, 381–382
 - vs.* selective naïve models, 387
 - summary, 449–450
- Merton model of risky debt, 390, 438–441, 442–451, 488, 703
- estimating default probabilities in, 445–446

- Merton model of risky debt (*Continued*)
 implied value and return volatility, 446–447
 multipayment bonds and, 444–445
 Merton's intertemporal capital asset pricing model, 445
 Merton's structural model, 238
 Merton term structure model, 286, 293–300. *See also* Ho and Lee model
 Merton-type structural model, 419
 Mesler, Mark, 415, 419, 430
 MF Global, 530
 Miller, M. H., 440, 442
 Mismatching, 250–256
 Model arbitrage, xxviii
 Model calibration, 377–379
Modeling Fixed Income Securities and Interest Rate Options (Jarrow), 123
 Model risk alerts, 257, 421
 Modern fixed income performance attribution, 772
Modern Pricing of Interest-Rate Derivatives (Rebonato), 548
 Modified duration, 276–277
 Modigliani, Franco, 440, 442
 Moench, Emanuel, 327, 332
Money Ball (Lewis), 768
 Money market funds, 38
 Money market futures contracts, 546
 Monte Carlo simulation, 10–11, 547
 advantages, 617–618
 in fixed income options, 596, 615–618
 with foreign currency derivatives, 677–678
 interpolated results from, 243–249
 limitations, 616–617
 in MBSs, 646
 using HJM modeling, 236–238
 Monte Carlo VaR, 726
 Moral hazard, 426–429
 Morgan Stanley, xxx, xxxi, xxxiii, 342, 749–751, 752–753, 785–786, 788
 Mortality rates, 708–711
 Mortality tables, 710
 Mortgage-backed securities, 639
 asset and liability management (ALM), 640–643
 legacy prepayment analysis of, 643–647
 mortgage servicing rights, 648–654
 option-adjusted spread on, 643
 Mortgage defaults, 13, 652–653. *See also* Credit crisis of 2006–2011;
 Mortgage-backed securities
 Mortgage prepayments, 640–642
 best practices, 641
 path-dependent, 646
 rationally prepaid mortgages, 684–693
 refinancing incentives, 645–647
 role of home prices in, 652–653
 Mortgages:
 conventional, 650
 fixed-rate, 21
 floating rate, 3, 21, 59
 prepayment of, 622, 640
 Mortgage servicing rights, 648–654
 best practice vs. common practice in, 650–654
 cash flow sources in, 653
 incorrect hedging of, 653–654
 role of home prices in defaults and prepayments, 652–653
 simulation of random movements in yields, 651–652
 valuation of, 649–650
 yield curves use, 650–651, 652
 Morton, Andrew, 9–10, 123
 Mozilo, Angelo, xxix
 Multinomial logistic regression, 642–643, 644, 648, 699
 Multinomial logit, 640, 642
 Multipayment bonds, 444–445
 Multiperiod simulation, 27, 264–267
 Multiple models approach, xxiii
 Multithreading, 10
 Municipals, 346, 352
 Naïve models of risk, 386–387
 National government treasuries, 6
 Natural logarithm, 418
 Need for change, measuring, 801–803
 Negative alpha, 769, 770–772
 Negative duration, 673

- Nelson-Siegel yield curve smoothing technique, 73, 230, 404
vs. cubic splines, 101
vs. maximum smoothness forward rate approach, 108, 110
vs. quadratic spline approach, 91–92
vs. stepwise constant yields approach, 77–85
- Net assets, 453
- Net income, 8, 56, 732
- Net income simulation, 3–4, 258
- Net present value, 495
- New Capital Accords, 375, 783, 791
- New Century, xxix
- NINJAs, 639
- “No arbitrage,” 231–234, 316
- Nominal interest rates, 55
- Nomura, 342
- Nonmaturity deposits, 656–674
bank credit risk impact on, 669–672
case study, 672–673
deposit franchise, 657
rate and balance movement formulas, 659–669
three-month savings deposits analysis, 672–673
total cash flow of, 658–669
- Normally distributed returns, 700–702, 727
- Northern Rock, xxx, xxxi, 669–670
- O’Brien, James, 673, 674
- O’Neal, Stanley, xxx, 488
- One-factor models, 161, 230–231. *See also* Black, Derman, and Toy model; Cox, Ingersoll, and Ross (CIR) model; Ho and Lee model; Hull and White model; Vasicek model
bushy trees, 129–137
coupon-bearing bond valuation, 140, 156–158
digital option valuation, 140, 158
forward rate volatility assumptions, 142–146
Ito’s lemma in, 287–288
key implications and notation, 129–131, 146–147
model assumptions and key input data, 124–129, 142–146
pitfalls in use of, 238–243
probability-weighted discount factors, 138–139
pseudo probabilities, 131–132, 147
with rate and maturity-dependent volatility, 143–160
valuation in, 139, 153–156
- Open interest, 513–514
- Operational risk, 7
- Optimal call, 627
- Option-adjusted spread (OAS), 643, 645–647
- Option-adjusted value (OAV) spread, 647
- Option ARMs, 742–743
- Option on assets of firm, 440, 447–448, 450
- Options. *See also* European options;
Fixed income options
Bermudan option, 600
on coupon-bearing bonds, 501–508
digital, 586
risk management, 28
- Ordinal measure of credit risk, 422
- Ordinal ranking, 379–380
- Origination costs, 766
- Ornstein-Uhlenbeck process, 286, 287, 300
- Over-the-counter derivatives, 567
- Pairwise correlations, 489
- Pandit, Vikram, 5
- Parallel rate shock model, 293
- Parallel shifts in interest rates, 269–271
- Parallel yield curve shifts, 288–298
- Parameter fitting, 454
- Par coupon, 613
- Path-dependent prepayment, 646
- Patruno, Gregg, 647
- Pension funds, 6
interest rate risk management, 251–252
portfolio return assumption, 710
risk management in, 8–9, 35–36
- Pension obligations, 712–713
- Perceived risk, xxiv, 15

- Perfect hedge, 680
- Performance hedging, 389–390
- Performance measurement, 7, 765–782
 - and capital regulation, 26
 - commercial banks, 9
 - components of performance, 778, 779
 - default risk in, 768–769
 - fixed income performance attribution, 772–780
 - fixed income portfolio, 770–772
 - insurance and banking, 769–770
 - Jarrow-Merton put option, 780–782
 - pension plans, 8
 - plus alpha, 767–768, 769–770
 - transaction level *vs.* portfolio level, 766–767
 - transfer pricing, 767–768
- Periodicity of analysis, 267
- Personal computers, 10
- Pitfalls in interest rate risk management, 238–243
- Plus alpha performance measurement, 769–772
- Point-in-time ratings, 423–425
- Pollack, Lisa, 491
- Pool rate, 21
- Portfolio return assumption, 710
- Prather, Larry, 650, 655
- Predictive ROC accuracy ratio, 380–381
- Prepayment behavior, 622, 624–637, 652–653
 - amortizing loan example, 626–636
 - historical research into, 624, 637
 - irrational, and credit risk, 637–638
 - mortgages (*see* Mortgage prepayments)
 - transactions cost approach, 624–625
- Prepayment models, 641
- Prepayment of mortgages. *See* Mortgage prepayments
- Prepayment option, 597, 616
- Prepayment rate, 641
- Prepayment risk, 648
- Prepayment speeds, 643–645
- Prepayment tables, 641, 644, 646
- Present value, 59–60, 70
 - calculating, 61–62
 - convexity and, 280
 - in risky bond valuation, 453–454, 465
- Present value of a basis point (PVBP), 726
- Price, 60, 61
- Price shifts, 132–135
- Primary capital, 38
- Primary Dealer Credit Facility, 736, 737
- Prince, Charles “Chuck,” xxvii, xxx, 488
- Principal, on fixed coupon bonds, 62–63
- Probability-weighted discount factors, 138–139, 456, 459–461
- Property and casualty insurance, 252, 713
- Prussia, Leland, 21, 22
- Pseudo probabilities:
 - vs.* empirical probabilities, 235–236
 - in one-factor models, 131–132, 147
 - in three-factor models, 205–217
 - in two-factor models, 171–178
- Put option prices, 57
- Put options:
 - and interest rate risk, 55–58
 - Jarrow-Merton framework (*see* Jarrow-Merton put option)
- Quadratic forward rate function, 73
- Quadratic splines, 85–94
- Quadratic yield function, 73
- Quantitative credit models, 11
- Quantitative default probabilities, 422, 424–425, 787–789
- Quartic forward rate function, 73, 104–111, 115
- Quartic spline, 101, 103
- Random walk, 285
- Range floater, 588, 592–593
- Ranson, Brian J., 474
- “The Rating Chernobyl” (Jarrow), 426, 428
- Rating primary risk system, 803
- Ratings. *See* Legacy ratings
- Ratings agencies, 422
 - default history, 429
 - moral hazard in self-assessment, 426–429
 - third-party audit, 429
- Rational call behavior, 608

- Rationally prepaid mortgage, 684–693
- Rational prepayment, 640. *See also*
- Prepayment behavior
- Raynes, Sylvain, 474
- Real estate, 251
- Rebonato, Riccardo, 548
- Receiver operating characteristics curve, 378
- Recovery, 362
- Recovery rates, 475
- Reduced form credit models
- correlations in default probabilities, 372–373
 - in credit risky bond valuation, 471–472
 - Jarrow and Turnbull model, 359–361, 373–374
 - Jarrow-Chava model, 380–381
 - Jarrow model, 361–372, 373–374
 - vs.* Merton model, 381–382
 - tests using historical data, 373–374
- Reduced form default probabilities, 391–394, 429–431
- Reduced form modeling, 483–486
- Reese, Ann, 5
- Refinancing incentives, 645–647
- Regression coefficients, 163, 166, 196
- Regulation, 340–341. *See also* Basel
- capital regulations
 - in capital management, 37–40
 - in CDOs, 478–480
- Regulation Q, 38
- Regulatory capital calculations, 7
- Regulatory capital ratios, 724
- Replicating portfolio, 8
- Reporting, 789–790, 805–806
- Request for Proposal (RFP) process, 795–796
- Reserve for loan losses, 724
- Reserve Primary Money Fund, xxxiii
- Residential real estate, 5. *See also*
- Mortgages
- Resolution Trust Corporation, 674
- Retail client default modeling, 371
- Return on the market, 702
- Returns, measuring, 20
- Return volatility, 446–447
- Revolving credit, 694–699
- cash flow needs of lines of credit, 697–699
 - credit quality and line usage, 697
 - credit quality of borrower, 696
 - drawdowns, 696
 - fluctuating credit risk, 696
 - seasonal cash needs/revenues, 696
- Richard, Scott, 350–351
- Risk:
- concentration of, 481
 - integrated, 640, 714
 - interest rate (*see* Interest rate risk)
 - perceived, xxiv, 15
 - practical quantification of, 15
 - prepayment, 648
 - sovereign, 405
 - total, 453, 719–723, 790–791
- Risk-adjusted return on capital, 34–35
- Risk and return trade-offs, 11
- Risk-based capital, 39
- Risk factor limits, 5
- Risk factor lookup tables, 168–169
- Risk-free curve, 404–409
- Risk-free yield curve, 771–773
- Risk management:
- advent of personal computers and, 10
 - averting failures, 783–792
 - credit risk, 28
 - credit risk *vs.* market risk, 5, 8
 - definitions and objectives, 6–8, 15, 30, 719
 - equities, 28
 - historical developments in, 9–11
 - interest rates, 26–27
 - key fallacies in, xxiii–xxviii
 - mark-to-market approach, 4–6
 - modern approach to, 783–792
 - multiple models approach, xxiii
 - questions for CEOs and boards of directors, 723–725, 787–788
 - role of capital in, 32–35
 - silos approach to, xxiv–xxvi, 8–9, 10
 - tools/software, 257
- Risk management synthesis, 3–14
- Risk-neutral probabilities. *See* Pseudo probabilities
- Risk-neutral valuation, 132
- Risk systems, xxvi, 5

- Risk vendors, 796–797
- Risk-weighted assets, 39
- Risk-weighted capital ratio, 731
- Risky coupon bonds, 509
- Risky debt. *See* Merton model of risky debt
- Risky yield smoothing, 411–413
- Risky zero yield, 407
- Robinson, Patrick, 4
- ROC accuracy ratio, 379–380, 381, 387
- Rose, Charlie, 5
- Ross, Stephen A. *See* Cox, Ingersoll, and Ross (CIR) model
- Royal Bank of Scotland, xxvii, xxx, 4, 342, 694–695, 784
- Rubin, Robert, xxvii, 6, 487
- S&P 500 index, 17–18
- Safety and soundness, managing, 783–792
- Safety zone, 250, 254–255, 258, 259–263, 787
- Salmon, Felix, 481, 486, 488, 491, 509, 532
- Salomon Brothers, xxviii
- Sampling error, 565, 616–617
- Savings and loan associations, 3, 12
- Savings and loan crisis, 3–4, 12
- Savings banks, 6
- Scholes, Myron, 9
- Schwab, Charles, 10, 346
- Schwartz, E. *See* Brennan and Schwartz model
- Scott, Louis, 306, 307
- Seasonal revenues/cash needs, 696
- Securities firms, 6
- Securities Industry and Financial Markets Association, 486
- Security Pacific Corporation, 759
- Security Pacific National Bank, 759
- Segmentation, 416
- Selection bias, 16–19, 377
 - in historical VaR, 18–19
 - return *vs.* relative benchmark, 17–18
- Self-assessment, moral hazard in, 426–429
- Self-insurance, 729–731
- Seniority level of debt, 362
- Senior tranches, 6
- September 11 terrorist attacks, 759
- Shareholder value added, 37
- Shareholder value creation, 800–808
 - best practice *vs.* common practice in, 806
 - daily management reporting, 805–806
 - “do no harm,” 800–801
 - management perspectives, 807
 - measuring need for change, 801–803
 - pace of innovation, 806
 - rating primary risk system, 803
 - working-level perspective, 807–808
- Shareholder value maximization, 790–791
- Sharpe ratio, 31, 295
- Shimko, David, 116, 287. *See also* Shimko, Tejima, and van Deventer
- Shimko, Tejima, and van Deventer, 24, 287, 304, 306, 375, 447, 449, 703
- Shimko test, 112, 116
- Shumway, Tyler, 11, 375, 378, 381, 382, 387, 438, 439
- Silo risk management, xxiv–xxvi, 8–9, 10, 11, 640, 731, 735, 794, 801
- Simulation, of net income, 3–4, 8
- Simulation modeling, 264–267
 - data aggregation in, 266
 - key assumptions in, 264–266
- Simulation models, 257
- Singapore Exchange, 528–529
- Singh, Manoj, 624, 647
- Single monthly mortality rate, 640
- Singleton, Kenneth J. *See* Duffie and Singleton
- 64-bit operating system, 11, 257
- S.K. Securities, 546
- Slattery, Mark, 655
- Small business default modeling, 370–371, 378
- Smith, Bill, 342
- Smith, John, 342
- Smith, Randall, 14
- Smoothing, 396–397. *See also* Yield curve smoothing
- Smoothness, 43, 53–55, 246
- Sobehart, Jorge, 376, 379, 380, 382
- Society of Actuaries RP-2000 tables, 711

- Software installation, 797
- Software vendors, 798–799
- Solvency II, 791
- Sovereign credits, 353–357
- Sovereign risk, 405
- Sparks, Dan, xxix
- Spears, Taylor, 491
- Special purpose vehicles, 594
- Speed of mean reversion, 286
- Spot exchange rates, 676, 677
- Spot rates, 148, 150, 152–153, 154, 687–688
- Spread pricing, 23
- Spreadsheet software:
 - advent of, 10
 - HJM replication in, 185–186
- Standard & Poor's, xxix
- Standard deviation, of forward rate changes, 144, 145
- Stanford Institute for Economic Policy Research, 715
- Stanford University, 4
- Stanton, Richard, 647
- State Street Global Advisors, 747
- Statistical probabilities. *See* Empirical probabilities of interest rate movements
- Stepwise constant yields, 77–85
- Stochastic interest rates, 675, 678
- Stochastic processes, 243, 284–285. *See also* Random walk
- Strategic default, 684
- Strategies, 6, 8
- Stress tests, 425, 727, 733
- Structural models, 238, 381
- Subprime loans, 727, 742–744.
 - See also* Credit crisis of 2006–2011
- Sub-sovereigns, 352
- Sufficient statistic, 438–439
- Sun Trust, 649
- Swap rates, 650
- Swaps, 398, 580–586
- Swaptions, 398, 574–579
- Switzerland, xxxiii
- Sydney Futures Exchange Bank
 - Bill Contract, 515, 520, 525, 527–528, 529
- Synthesis, 3–14
- Synthetic CDOs, 480–481
- Synthetic replication, 702
- Szilagyi, Jan, 11
- Taleb, Nassim, 726, 729–730
- Targeted Investment Program (TIP), xxxiv
- Tavakoli, Janet M., 474
- Term Auction Facility, xxxi
- Term structure models:
 - alternative models, 303–308
 - credit model testing (*see* Credit model testing)
 - duration as, 288–292
 - estimating parameters, 316–332
 - in forward and futures contracts, 515, 519, 524
 - Ho and Lee model, 285
 - Merton's model, 286, 293–300
 - use of Ito's lemma, 286–288, 289
 - Vasicek model, 286, 292–293
 - vocabulary of, 284–286
- Term structures:
 - of default probabilities, 124
 - forward rates, 136
 - interest rates, 126
 - zero coupon bond prices, 136
- Terry, William “Mack,” 21, 22, 267
- Thain, John, 4–5
- Third-party audit, 429
- Third-party insurance, 729–731
- Three-factor models, 190–229, 230–231. *See also* Chen's model
 - coupon-bearing bond valuation, 225–226
 - digital option valuation, 227–229
 - with foreign currency derivatives, 678
 - formula for zero coupon bond price shifts, 205–206
 - interest rate volatility surface, 200
 - key implications and notation, 201–205
 - model assumptions and key input data, 191–192
 - pseudo probabilities, 205–217
 - risk factors, 192–200
 - valuation in, 217–229

- Through-the-cycle ratings, 423–425
Tokyo Electric Power, 708, 709, 710, 714
“Too big to fail,” 751
Total risk, 453
 hedging, 790–791
 Jarrow-Merton put option as measure of, 719–723
Toy, W., 146. *See also* Black, Derman, and Toy model
TRACE (Trade Reporting and Compliance Engine), 337
Tracking error, 17
Trade volume, 417
Trading, 4, 7
Trading floor, 8–9
Tranches, 6, 7, 12
Transactions cost approach, 624–625, 643, 647
Transfer pricing, 7, 15, 765, 767–768
 at Bank of America, 21–24
 at First Interstate Bancorp, 24–26
 holding company, 25
 matched maturity, 22–24
 selective history of, 20–26
Transition matrices, 426
Treasury bills, 12
Trinomial lattices, 619–620
Troubled Asset Relief Program (TARP), xxxiii, xxxiv, 736, 746
Tsunami aftermath, 714
Tucker, Alan L. *See* Hilliard, Madura, and Tucker (HMT) model
Turnbull, Stuart. *See* Jarrow and Turnbull model
Two-factor models, 230–231
 Brennan and Schwartz model, 306–307
 coupon-bearing bond valuation, 184, 187
 digital option valuation, 185, 187–189
 Ito’s lemma in, 287 (*see also* Brennan and Schwartz model)
 key implications and notation, 167–170
 model assumptions and key input data, 162–163
 pseudo probabilities, 171–178
 valuation in, 178–186
 zero coupon bond valuation, 183, 186
UBS, xxx, xxxi, xxxiii, 342, 488
UK mortgage losses, xxx
UK short-term loan scheme, xxxiii
UniCredit, 342
United States Senate, Levin Report, 431, 435–437
University of Iowa, xxviii
Upshifts/downshifts in interest rates, 125, 126, 129–131
US Bancorp, 649
U.S. Treasury forward rates, 53–55, 125
U.S. Treasury yield curve, 45–57, 190
Uyemura, Dennis, 37, 38, 41, 257
Value, 60–61. *See also* Present value
 amortizing loans, 63–64
 fixed coupon bonds, 62–63
Value at risk (VaR), 719–720
 false, 730
 historical, 18–19, 726–728
 Monte Carlo, 19, 726, 727
 problems of, 726–728
 vs. put option approach, 728–731
 PVBP, 726
 selection bias in, 16–17
 variance/covariance approach, 18–19, 726
 VaR technology, 726, 727
Vanasek, Jim, xxviii, 742
Van Deventer, Donald R., 37, 38, 41, 257. *See also* Adams and van Deventer; Dickler, Jarrow, and van Deventer; Jarrow, van Deventer, and Wang
Van Deventer and Imai, 19, 34, 36, 37, 38, 40, 41
Van Deventer and Wang test, 378, 383, 390
Variance-covariance VaR, 18–19, 726, 727
Vasicek, Oldrich, 9, 75, 80, 96, 102, 122. *See also* Vasicek model

- Vasicek model, 128, 143, 238–239, 286, 292–293, 300–302, 310–313, 509–511
- Vega equivalents, 594
- Vetzal, K. R., 706
- Volatilities. *See* Forward rate volatility
- Volatility of equity returns, 446
- Volatility smile, 28, 389
- Wachovia, xxxii, xxxiii, 4, 47, 670, 768
- Wall Street Journal*, 14, 487–488
- Wang, Xiaoming. *See* Jarrow, van Deventer, and Wang
- Washington Mutual, xxviii, xxxiii, 4, 47, 369–370, 371, 742–746, 768
- Waterfall, 481, 482
- Webber, Nick, 706
- Wells Fargo, xxxiii, 649, 786, 788, 789
- “When issued” bond, 71
- White, Alan. *See also* Hull and White lattice approach; Hull and White model
- Whitehouse, Mark, 488
- Wiener process, 284–285, 287, 293, 299, 300, 305, 314
- Wilson, Mungo, 421, 429–431, 438
- Winn-Dixie, 768
- WMI Holdings. *See* Washington Mutual
- World War I, 4
- Wuffi, Peter, xxx
- Yield curves, 27, 62, 64, 212–217, 388
continuous, 495
in mortgage servicing rights, 650–651, 652
risk factors and, 230–234
- Yield curve simulation. *See* HJM model
- Yield curve smoothing, 73–122, 495. *See also* Forward rate smoothing
with coupon-bearing bond prices, 116, 399
with credit spreads, 398, 404–409
cubic spline approach, 94–101, 115
curve length, 112–116
maximum smoothness forward rate technique, 101–111, 117–120
- Nelson-Siegel approach (*see* Nelson-Siegel yield curve smoothing technique)
- quadratic spline approach, 85–94
- quartic forward rate function, 104–111, 115
- revised inputs, 399–400
- in risk-free bond valuation, 455
- Shimko test, 116
- steps to defining best methodology, 73–77
- stepwise constant yields approach, 77–85
- techniques comparison, 111–112
- Yield curve twists, 161–162, 190–191, 241, 278
- Yield step function, 73
- Yield to maturity, 62, 68–69, 273–276. *See also* Credit spreads
- Yildirim, Yildiray, 364
- Zero coupon bonds
building bushy trees, 132–137
fitting prices jointly with volatility parameters, 317–330
forward contracts on, 514–524, 532–538
forward rate agreements, 520–524
futures contracts on, 527–528
price shifts, 124, 132–135, 147–149
yields, 77–85
- Zullo, Ferdinando, 658