

CHAPMAN & HALL/CRC FINANCIAL MATHEMATICS SERIES

An Introduction to

CREDIT RISK MODELING

Christian Bluhm

*HypoVereinsbank Group
Munich, Germany*

Ludger Overbeck

*Deutsche Bank Group
Frankfurt, Germany*

Christoph Wagner

*Allianz Group
Munich, Germany*



CHAPMAN & HALL/CRC

A CRC Press Company

Boca Raton London New York Washington, D.C.

Library of Congress Cataloging-in-Publication Data

Bluhm, Christian.

An introduction to credit risk modeling / Christian Bluhm, Ludger Overbeck, Christoph Wagner.

p. cm. — (Chapman & Hall/ CRC financial mathematics series)

Includes bibliographical references and index.

ISBN 1-58488-326-X (alk. paper)

1. Credit—Management—Mathematical models. 2. Risk management—Mathematical models. I. Overbeck, Ludger. II. Wagner, Christoph. III. Title. IV. Series.

HG3751 .B58 2002

658.8'8'.0151—dc21

2002073623

CIP

This book contains information obtained from authentic and highly regarded sources. Reprinted material is quoted with permission, and sources are indicated. A wide variety of references are listed. Reasonable efforts have been made to publish reliable data and information, but the author and the publisher cannot assume responsibility for the validity of all materials or for the consequences of their use.

Neither this book nor any part may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, microfilming, and recording, or by any information storage or retrieval system, without prior permission in writing from the publisher.

All rights reserved. Authorization to photocopy items for internal or personal use, or the personal or internal use of specific clients, may be granted by CRC Press LLC, provided that \$.50 per page photocopied is paid directly to Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923 USA. The fee code for users of the Transactional Reporting Service is ISBN 1-58488-326-X/01/\$0.00+\$1.50. The fee is subject to change without notice. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

The consent of CRC Press LLC does not extend to copying for general distribution, for promotion, for creating new works, or for resale. Specific permission must be obtained in writing from CRC Press LLC for such copying.

Direct all inquiries to CRC Press LLC, 2000 N.W. Corporate Blvd., Boca Raton, Florida 33431.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation, without intent to infringe.

Visit the CRC Press Web site at www.crcpress.com

© 2003 by Chapman & Hall/CRC

No claim to original U.S. Government works

International Standard Book Number 1-58488-326-X

Library of Congress Card Number 2002073623

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0

Printed on acid-free paper

Preface

In banking, especially in risk management, portfolio management, and structured finance, solid quantitative know-how becomes more and more important. We had a two-fold intention when writing this book:

First, this book is designed to help mathematicians and physicists leaving the academic world and starting a profession as risk or portfolio managers to get quick access to the world of credit risk management. Second, our book is aimed at being helpful to risk managers looking for a more quantitative approach to credit risk.

Following this intention on one side, our book is written in a *Lecture Notes* style very much reflecting the keyword “introduction” already used in the title of the book. We consequently avoid elaborating on technical details not really necessary for understanding the underlying idea. On the other side we kept the presentation mathematically precise and included some proofs as well as many references for readers interested in diving deeper into the mathematical theory of credit risk management.

The main focus of the text is on *portfolio* rather than *single obligor* risk. Consequently correlations and factors play a major role. Moreover, most of the theory in many aspects is based on probability theory. We therefore recommend that the reader consult some standard text on this topic before going through the material presented in this book. Nevertheless we tried to keep it as self-contained as possible.

Summarizing our motivation for writing an introductory text on credit risk management one could say that *we tried to write the book we would have liked to read before starting a profession in risk management* some years ago.

Munich and Frankfurt, August 2002

Christian Bluhm, Ludger Overbeck, Christoph Wagner

Acknowledgements

Christian Bluhm would like to thank his wife Tabea and his children Sarah and Noa for their patience during the writing of the manuscript. Without the support of his great family this project would not had come to an end. Ludger Overbeck is grateful to his wife Bettina and his children Leonard, Daniel and Clara for their ongoing support.

We very much appreciated feedback, support, and comments on the manuscript by our colleagues.

Questions and remarks of the audiences of several conferences, seminars and lectures, where parts of the material contained in this book have been presented, in many ways improved the manuscript. We always enjoyed the good discussions on credit risk modeling issues with colleagues from other financial institutions. To the many people discussing and sharing with us their insights, views, and opinions, we are most grateful.

Disclaimer

This book reflects the personal view of the authors and not the opinion of HypoVereinsbank, Deutsche Bank, or Allianz. The contents of the book has been written for educational purposes and is neither an offering for business nor an instruction for implementing a bank-internal credit risk model. The authors are not liable for any damage arising from any application of the theory presented in this book.

About the Authors

Christian Bluhm works for HypoVereinsbank's group portfolio management in Munich, with a focus on portfolio modeling and risk management instruments. His main responsibilities include the analytic evaluation of ABS transactions by means of portfolio models, as introduced in this book.

His first professional position in risk management was with Deutsche Bank, Frankfurt. In 1996, he earned a Ph.D. in mathematics from the University of Erlangen-Nuernberg and, in 1997, he was a post-doctoral member of the mathematics department of Cornell University, Ithaca, New York. He has authored several papers and research articles on harmonic and fractal analysis of random measures and stochastic processes. Since he started to work in risk management, he has continued to publish in this area and regularly speaks at risk management conferences and workshops.

Christoph Wagner works on the risk methodology team of Allianz Group Center. His main responsibilities are credit risk and operational risk modeling, securitization and alternative risk transfer. Prior to Allianz he worked for Deutsche Bank's risk methodology department. He holds a Ph.D. in statistical physics from the Technical University of Munich. Before joining Deutsche Bank he spent several years in postdoctoral positions, both at the Center of Nonlinear Dynamics and Complex Systems, Brussels and at Siemens Research Department in Munich. He has published several articles on nonlinear dynamics and stochastic processes, as well as on risk modeling.

Ludger Overbeck heads the Research and Development team in the Risk Analytics and Instrument department of Deutsche Bank's credit risk management function. His main responsibilities are the credit portfolio model for the group-wide RAROC process, the risk assessment of credit derivatives, ABS, and other securitization products, and operational risk modeling. Before joining Deutsche Bank in 1997, he worked with the Deutsche Bundesbank in the supervision department, examining internal market risk models.

He earned a Ph.D. in Probability Theory from the University of Bonn. After two post-doctoral years in Paris and Berkeley, from 1995 to 1996, he finished his Habilitation in Applied Mathematics during his affiliation with the Bundesbank. He still gives regular lectures in the mathematics department of the University in Bonn and in the Business and Economics Department at the University in Frankfurt. In Frankfurt he received a Habilitation in Business and Economics in 2001. He has published papers in several forums, from mathematical and statistical journals, journals in finance and economics, including RISK Magazine and practioners handbooks. He is a frequent speaker at academic and practioner conferences.

Contents

1 The Basics of Credit Risk Management

- 1.1 Expected Loss
 - 1.1.1 The Default Probability
 - 1.1.1.1 Ratings
 - 1.1.1.2 Calibration of Default Probabilities to Ratings
 - 1.1.2 The Exposure at Default
 - 1.1.3 The Loss Given Default
- 1.2 Unexpected Loss
 - 1.2.1 Economic Capital
 - 1.2.2 The Loss Distribution
 - 1.2.2.1 Monte Carlo Simulation of Losses
 - 1.2.2.2 Analytical Approximation
 - 1.2.3 Modeling Correlations by Means of Factor Models
- 1.3 Regulatory Capital and the Basel Initiative

2 Modeling Correlated Defaults

- 2.1 The Bernoulli Model
 - 2.1.1 A General Bernoulli Mixture Model
 - 2.1.2 Uniform Default Probability and Uniform Correlation
- 2.2 The Poisson Model
 - 2.2.1 A General Poisson Mixture Model
 - 2.2.2 Uniform Default Intensity and Uniform Correlation
- 2.3 Bernoulli Versus Poisson Mixture
- 2.4 An Overview of Today's Industry Models
 - 2.4.1 CreditMetrics™ and the KMV-Model
 - 2.4.2 CreditRisk⁺
 - 2.4.3 CreditPortfolioView
 - 2.4.3.1 CPV Macro
 - 2.4.3.2 CPV Direct
 - 2.4.4 Dynamic Intensity Models

- 2.5 One-Factor/Sector Models
 - 2.5.1 The CreditMetrics™/KMV One-Factor Model
 - 2.5.2 The CreditRisk⁺ One-Sector Model
 - 2.5.3 Comparison of One-Factor and One-Sector Models
- 2.6 Loss Distributions by Means of Copula Functions
 - 2.6.1 Copulas: Variations of a Scheme
- 2.7 Working Example: Estimation of Asset Correlations
- 3 Asset Value Models**
 - 3.1 Introduction and a Small Guide to the Literature
 - 3.2 A Few Words about Calls and Puts
 - 3.2.1 Geometric Brownian Motion
 - 3.2.2 Put and Call Options
 - 3.3 Merton's Asset Value Model
 - 3.3.1 Capital Structure: Option-Theoretic Approach
 - 3.3.2 Asset from Equity Values
 - 3.4 Transforming Equity into Asset Values: A Working Approach
 - 3.4.1 Itô's Formula "Light"
 - 3.4.2 Black-Scholes Partial Differential Equation
- 4 The CreditRisk⁺ Model**
 - 4.1 The Modeling Framework of CreditRisk⁺
 - 4.2 Construction Step 1: Independent Obligor
 - 4.3 Construction Step 2: Sector Model
 - 4.3.1 Sector Default Distribution
 - 4.3.2 Sector Compound Distribution
 - 4.3.3 Sector Convolution
- 5 Alternative Risk Measures and Capital Allocation**
 - 5.1 Coherent Risk Measures and Conditional Shortfall
 - 5.2 Contributory Capital
 - 5.2.1 Variance/Covariance Approach
 - 5.2.2 Capital Allocation w.r.t. Value-at-Risk
 - 5.2.3 Capital Allocations w.r.t. Expected Shortfall
 - 5.2.4 A Simulation Study
- 6 Term Structure of Default Probability**
 - 6.1 Survival Function and Hazard Rate
 - 6.2 Risk-neutral vs. Actual Default Probabilities

- 6.3 Term Structure Based on Historical Default Information
 - 6.3.1 Exponential Term Structure
 - 6.3.2 Direct Calibration of Multi-Year Default Probabilities
 - 6.3.3 Migration Technique and Q-Matrices
- 6.4 Term Structure Based on Market Spreads

7 Credit Derivatives

- 7.1 Total Return Swaps
- 7.2 Credit Default Products
- 7.3 Basket Credit Derivatives
- 7.4 Credit Spread Products
- 7.5 Credit-linked Notes

8 Collateralized Debt Obligations

- 8.1 Introduction to Collateralized Debt Obligations
 - 8.1.1 Typical Cash Flow CDO Structure
 - 8.1.1.1 Overcollateralization Tests
 - 8.1.1.2 Interest Coverage Tests
 - 8.1.1.3 Other Tests
 - 8.1.2 Typical Synthetic CLO Structure
- 8.2 Different Roles of Banks in the CDO Market
 - 8.2.1 The Originator's Point of View
 - 8.2.1.1 Regulatory Arbitrage and Capital Relief
 - 8.2.1.2 Economic Risk Transfer
 - 8.2.1.3 Funding at Better Conditions
 - 8.2.1.4 Arbitrage Spread Opportunities
 - 8.2.2 The Investor's Point of View
- 8.3 CDOs from the Modeling Point of View
 - 8.3.1 Multi-Step Models
 - 8.3.2 Correlated Default Time Models
 - 8.3.3 Stochastic Default Intensity Models
- 8.4 Rating Agency Models: Moody's BET
- 8.5 Conclusion
- 8.6 Some Remarks on the Literature

References

List of Figures

- 1.1 Calibration of Moody's Ratings to Default Probabilities
- 1.2 The Portfolio Loss Distribution
- 1.3 An empirical portfolio loss distribution
- 1.4 Analytical approximation by some beta distribution
- 1.5 Correlation induced by an underlying factor
- 1.6 Correlated processes of obligor's asset value log-returns
- 1.7 Three-level factor structure in KMV's Factor Model
- 2.1 Today's Best-Practice Industry Models
- 2.2 Shape of Ganima Distributions for some parameter sets
- 2.3 CreditMetrics™/KMV One-Factor Model: Conditional default probability as a function of the factor realizations
- 2.4 CreditMetrics™/KMV One-Factor Model: Conditional default probability as a function of the average one-year default probability
- 2.5 The probability density f_{ps}
- 2.6 Economic capital EC_α in dependence on α
- 2.7 Negative binomial distribution N with parameters $(\alpha, \beta) = (1, 30)$
- 2.8 $t(3)$ -density versus $N(0, 1)$ -density
- 2.9 Normal versus t -dependency with same linear correlation
- 2.10 Estimated economic cycle compared to Moody's average historic default frequencies
- 3.1 Hedging default risk by a long put
- 3.2 Asset-Equity relation
- 5.1 Expected Shortfall
- 5.2 Shortfall contribution versus var/covar-contribution
- 5.3 Shortfall contribution versus Var/Covar-contribution for business units
- 6.1 Cumulative default rate for A-rated issuer
- 6.2 Hazard rate functions
- 7.1 Total return swap
- 7.2 Credit default swap
- 7.3 Generating correlated default times via the copula approach
- 7.4 The averages of the standard deviation of time to default, first-to-default- and last-to-default-time
- 7.5 k th-to-default spread versus correlation for a basket with three underlyings
- 7.6 Default spread versus correlation between reference asset and swap counterparty
- 7.7 Credit spread swap
- 7.8 Example of a Credit-linked Note
- 8.1 Classification of CDOs
- 8.2 Example of a cash flow CDO
- 8.3 Example of waterfalls in a cash flow CDO
- 8.4 Example of a synthetic CDO
- 8.5 Equity return distribution of a CDO
- 8.6 CFO modeling scheme
- 8.7 CDO modeling workflow based on default times
- 8.8 Diversification Score as a function of m
- 8.9 Fitting loss distributions by the BET
- 8.10 Tranching a Loss Distribution

Chapter 1

The Basics of Credit Risk Management

Why is credit risk management an important issue in banking? To answer this question let us construct an example which is, although simplified, nevertheless not too unrealistic: Assume a major building company is asking its house bank for a loan in the size of ten billion Euro. Somewhere in the bank's credit department a senior analyst has the difficult job to decide if the loan will be given to the customer or if the credit request will be rejected. Let us further assume that the analyst knows that the bank's chief credit officer has known the chief executive officer of the building company for many years, and to make things even worse, the credit analyst knows from recent default studies that the building industry is under hard pressure and that the *bank-internal rating*¹ of this particular building company is just on the way down to a low *subinvestment grade*.

What should the analyst do? Well, the most natural answer would be that the analyst should reject the deal based on the information she or he has about the company and the current market situation. An alternative would be to grant the loan to the customer but to *insure* the loss potentially arising from the engagement by means of some credit risk management instrument (e.g., a so-called *credit derivative*).

Admittedly, we intentionally exaggerated in our description, but situations like the one just constructed happen from time to time and it is never easy for a credit officer to make a decision under such difficult circumstances. A brief look at any typical banking portfolio will be sufficient to convince people that defaulting obligors belong to the daily business of banking the same way as credit applications or ATM machines. Banks therefore started to think about ways of *loan insurance* many years ago, and the insurance paradigm will now directly lead us to the first central building block credit risk management.

¹A rating is an indication of creditworthiness; see [Section 1.1.1.1](#).

1.1 Expected Loss

Situations as the one described in the introduction suggest the need of a *loss protection* in terms of an *insurance*, as one knows it from car or health insurances. Moreover, history shows that even good customers have a potential to default on their financial obligations, such that an insurance for not only the critical but all loans in the bank's credit portfolio makes much sense.

The basic idea behind insurance is always the same. For example, in health insurance the costs of a few sick customers are covered by the total sum of revenues from the fees paid to the insurance company by all customers. Therefore, the fee that a man at the age of thirty has to pay for health insurance protection somehow reflects the insurance company's experience regarding *expected costs* arising from this particular group of clients.

For bank loans one can argue exactly the same way: Charging an appropriate *risk premium* for every loan and collecting these risk premiums in an internal bank account called *expected loss reserve* will create a capital cushion for covering losses arising from defaulted loans.

In probability theory the attribute *expected* always refers to an *expectation* or *mean value*, and this is also the case in risk management. The basic idea is as follows: The bank assigns to every customer a *default probability* (DP), a loss fraction called the *loss given default* (LGD), describing the fraction of the loan's exposure expected to be lost in case of default, and the *exposure at default* (EAD) subject to be lost in the considered time period. The loss of any obligor is then defined by a *loss variable*

$$\tilde{L} = \text{EAD} \times \text{LGD} \times L \quad \text{with} \quad L = \mathbf{1}_D, \quad \mathbb{P}(D) = \text{DP}, \quad (1.1)$$

where D denotes the *event* that the obligor defaults in a certain period of time (most often one year), and $\mathbb{P}(D)$ denotes the probability of D . Although we will not go too much into technical details, we should mention here that underlying our model is some probability space $(\Omega, \mathcal{F}, \mathbb{P})$, consisting of a *sample space* Ω , a σ -*Algebra* \mathcal{F} , and a probability measure \mathbb{P} . The elements of \mathcal{F} are the *measurable events* of the model, and intuitively it makes sense to claim that the event of default should be measurable. Moreover, it is common to identify \mathcal{F} with

the *information* available, and the information if an obligor defaults or survives should be included in the set of measurable events.

Now, in this setting it is very natural to define the *expected loss* (EL) of any customer as the expectation of its corresponding loss variable \tilde{L} , namely

$$\text{EL} = \mathbb{E}[\tilde{L}] = \text{EAD} \times \text{LGD} \times \mathbb{P}(D) = \text{EAD} \times \text{LGD} \times \text{DP}, \quad (1.2)$$

because the expectation of any *Bernoulli* random variable, like $\mathbf{1}_D$, is its event probability. For obtaining representation (1.2) of the EL, we need some additional assumption on the constituents of Formula (1.1), for example, the assumption that EAD and LGD are constant values. This is not necessarily the case under all circumstances. There are various situations in which, for example, the EAD has to be modeled as a random variable due to uncertainties in amortization, usage, and other drivers of EAD up to the chosen planning horizon. In such cases the EL is still given by Equation (1.2) if one can assume that the exposure, the loss given default, and the default event D are independent and EAD and LGD are the expectations of some underlying random variables. But even the independence assumption is questionable and in general very much simplifying. Altogether one can say that (1.2) is the most simple representation formula for the expected loss, and that the more simplifying assumptions are dropped, the more one moves away from closed and easy formulas like (1.2).

However, for now we should not be bothered about the independence assumption on which (1.2) is based: The basic concept of expected loss is the same, no matter if the constituents of formula (1.1) are independent or not. Equation (1.2) is just a convenient way to write the EL in the first case. Although our focus in the book is on *portfolio risk* rather than on *single obligor risk* we briefly describe the three constituents of Formula (1.2) in the following paragraphs. Our convention from now on is that the EAD always is a deterministic (i.e., nonrandom) quantity, whereas the *severity* (SEV) of loss in case of default will be considered as a random variable with expectation given by the LGD of the respective facility. For reasons of simplicity we assume in this chapter that the severity is independent of the variable L in (1.1).

1.1.1 The Default Probability

The task of assigning a default probability to every customer in the bank's credit portfolio is far from being easy. There are essentially two approaches to default probabilities:

- *Calibration of default probabilities from market data.*

The most famous representative of this type of default probabilities is the concept of *Expected Default Frequencies* (EDF) from KMV² Corporation. We will describe the KMV-Model in [Section 1.2.3](#) and in [Chapter 3](#).

Another method for calibrating default probabilities from market data is based on credit spreads of traded products bearing credit risk, e.g., corporate bonds and credit derivatives (for example, *credit default swaps*; see the chapter on *credit derivatives*).

- *Calibration of default probabilities from ratings.*

In this approach, default probabilities are associated with *ratings*, and ratings are assigned to customers either by external rating agencies like *Moody's Investors Services*, *Standard & Poor's* (S&P), or *Fitch*, or by bank-internal rating methodologies. Because ratings are not subject to be discussed in this book, we will only briefly explain some basics about ratings. An excellent treatment of this topic can be found in a survey paper by Crouhy et al. [22].

The remaining part of this section is intended to give some basic indication about the calibration of default probabilities to ratings.

1.1.1.1 Ratings

Basically ratings describe the *creditworthiness* of customers. Hereby quantitative as well as qualitative information is used to evaluate a client. In practice, the rating procedure is often more based on the judgement and experience of the rating analyst than on pure mathematical procedures with strictly defined outcomes. It turns out that in the US and Canada, most issuers of public debt are rated at least by two of the three main rating agencies Moody's, S&P, and Fitch.

²KMV Corp., founded 13 years ago, headquartered in San Francisco, develops and distributes credit risk management products; see www.kmv.com.

Their reports on *corporate bond defaults* are publicly available, either by asking at their local offices for the respective reports or conveniently per web access; see www.moodys.com, www.standardandpoors.com, www.fitchratings.com.

In Germany and also in Europe there are not as many companies issuing traded debt instruments (e.g., bonds) as in the US. Therefore, many companies in European banking books do not have an external rating. As a consequence, banks need to invest³ more effort in their own bank-internal rating system. The natural candidates for assigning a rating to a customer are the credit analysts of the bank. Hereby they have to consider many different *drivers* of the considered firm's economic future:

- Future *earnings* and *cashflows*,
- *debt, short- and long-term liabilities*, and *financial obligations*,
- *capital structure* (e.g., *leverage*),
- *liquidity* of the firm's assets,
- situation (e.g., political, social, etc.) of the firm's home *country*,
- situation of the *market* (e.g., *industry*), in which the company has its main activities,
- *management quality*, *company structure*, etc.

From this by no means exhaustive list it should be obvious that a rating is an *attribute of creditworthiness* which can not be captured by a pure mathematical formalism. It is a best practice in banking that ratings as an outcome of a statistical tool are always re-evaluated by the rating specialist in charge of the rating process. It is frequently the case that this re-evaluation moves the rating of a firm by one or more notches away from the “mathematically” generated rating. In other words, *statistical tools provide a first indication* regarding the rating of a customer, but due to the various *soft factors* underlying a rating, the

³Without going into details we would like to add that banks always should base the decision about creditworthiness on their bank-internal rating systems. As a main reason one could argue that banks know their customers best. Moreover, it is well known that external ratings do not react quick enough to changes in the economic health of a company. Banks should be able to do it better, at least in the case of their long-term relationship customers.

responsibility to assign a final rating remains the duty of the rating analyst.

Now, it is important to know that the rating agencies have established an ordered *scale* of ratings in terms of a letter system describing the creditworthiness of rated companies. The rating categories of Moody's and S&P are slightly different, but it is not difficult to find a mapping between the two. To give an example, [Table 1.1](#) shows the rating categories of S&P as published⁴ in [118].

As already mentioned, Moody's system is slightly different in meaning as well as in rating letters. Their rating categories are Aaa, Aa, A, Baa, Ba, B, Caa, Ca, C, where the creditworthiness is highest for Aaa and poorest for C. Moreover, both rating agencies additionally provide ratings on a *finer scale*, allowing for a more accurate distinction between different credit qualities.

1.1.1.2 Calibration of Default Probabilities to Ratings

The process of assigning a default probability to a rating is called a *calibration*. In this paragraph we will demonstrate how such a calibration works. The end product of a calibration of default probabilities to ratings is a mapping

$$\text{Rating} \mapsto \text{DP}, \quad \text{e.g.,} \quad \{AAA, AA, \dots, C\} \rightarrow [0, 1], \quad R \mapsto \text{DP}(R),$$

such that to every rating R a certain default probability $\text{DP}(R)$ is assigned.

In the sequel we explain by means of Moody's data how a calibration of default probabilities to external ratings can be done. From Moody's website or from other resources it is easy to get access to their recent study [95] of *historic corporate bond defaults*. There one can find a table like the one shown in [Table 1.2](#) (see [95] Exhibit 40) showing historic default frequencies for the years 1983 up to 2000.

Note that in our illustrative example we chose the *fine ratings scale* of Moody's, making finer differences regarding the creditworthiness of obligors.

Now, an important observation is that for best ratings no defaults at all have been observed. This is not as surprising as it looks at first sight: For example rating class *Aaa* is often calibrated with a default probability of 2 bps ("bp" stands for '*basispoint*' and means 0.01%),

⁴Note that we use shorter formulations instead of the exact wording of S&P.

TABLE 1.1: S&P Rating Categories [118].

AAA	<i>best credit quality extremely reliable with regard to financial obligations</i>
AA	<i>very good credit quality very reliable</i>
A	<i>more susceptible to economic conditions still good credit quality</i>
BBB	<i>lowest rating in investment grade</i>
BB	<i>caution is necessary best sub-investment credit quality</i>
B	<i>vulnerable to changes in economic conditions currently showing the ability to meet its financial obligations</i>
CCC	<i>currently vulnerable to nonpayment dependent on favourable economic conditions</i>
CC	<i>highly vulnerable to a payment default</i>
C	<i>close to or already bankrupt payments on the obligation currently continued</i>
D	<i>payment default on some financial obligation has actually occurred</i>

TABLE 1.2: Moody's Historic Corporate Bond Default Frequencies.

Rating	1983	1984	1985	1986	1987	1988
Aaa	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa2	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa3	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A2	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A3	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Baa1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Baa2	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Baa3	0.00%	1.06%	0.00%	4.82%	0.00%	0.00%
Ba1	0.00%	1.16%	0.00%	0.88%	3.73%	0.00%
Ba2	0.00%	1.61%	1.63%	1.20%	0.95%	0.00%
Ba3	2.61%	0.00%	3.77%	3.44%	2.95%	2.59%
B1	0.00%	5.84%	4.38%	7.61%	4.93%	4.34%
B2	10.00%	18.75%	7.41%	16.67%	4.30%	6.90%
B3	17.91%	2.90%	13.86%	16.07%	10.37%	9.72%

Rating	1989	1990	1991	1992	1993	1994
Aaa	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa2	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa3	1.40%	0.00%	0.00%	0.00%	0.00%	0.00%
A1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A2	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A3	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Baa1	0.00%	0.00%	0.76%	0.00%	0.00%	0.00%
Baa2	0.80%	0.00%	0.00%	0.00%	0.00%	0.00%
Baa3	1.07%	0.00%	0.00%	0.00%	0.00%	0.00%
Ba1	0.79%	2.67%	1.06%	0.00%	0.81%	0.00%
Ba2	1.82%	2.82%	0.00%	0.00%	0.00%	0.00%
Ba3	4.71%	3.92%	9.89%	0.74%	0.75%	0.59%
B1	6.24%	8.59%	6.04%	1.03%	3.32%	1.90%
B2	8.28%	22.09%	12.74%	1.54%	4.96%	3.66%
B3	19.55%	28.93%	28.42%	24.54%	11.48%	8.05%

Rating	1995	1996	1997	1998	1999	2000
Aaa	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa2	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa3	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A2	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A3	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Baa1	0.00%	0.00%	0.00%	0.00%	0.00%	0.29%
Baa2	0.00%	0.00%	0.00%	0.32%	0.00%	0.00%
Baa3	0.00%	0.00%	0.00%	0.00%	0.34%	0.98%
Ba1	0.00%	0.00%	0.00%	0.00%	0.47%	0.91%
Ba2	0.00%	0.00%	0.00%	0.61%	0.00%	0.66%
Ba3	1.72%	0.00%	0.47%	1.09%	2.27%	1.51%
B1	4.35%	1.17%	0.00%	2.13%	3.08%	3.25%
B2	6.36%	0.00%	1.50%	7.57%	6.68%	3.89%
B3	4.10%	3.36%	7.41%	5.61%	9.90%	9.92%

essentially meaning that one expects a *Aaa*-default in average twice in 10,000 years. This is a long time to go; so, one should not be surprised that quite often best ratings are lack of any default history. Nevertheless we believe that it would not be correct to take the historical zero-balance as an indication that these rating classes are risk-free opportunities for credit investment. Therefore, we have to find a way to assign small but positive default probabilities to those ratings.

Figure 1.1 shows our “quick-and-dirty working solution” of the problem, where we use the attribute “quick-and-dirty” because in practice one would try to do the calibration a little more sophisticatedly⁵.

However, for illustrative purposes our solution is sufficient, because it shows the main idea. We do the calibration in three steps:

1. Denote by $h_i(R)$ the historic default frequency of rating class R for year i , where i ranges from 1983 to 2000. For example, $h_{1993}(Ba1) = 0.81\%$. Then compute the mean value and the standard deviation of these frequencies over the years, where the rating is fixed, namely

$$m(R) = \frac{1}{18} \sum_{i=1983}^{2000} h_i(R) \quad \text{and}$$

$$s(R) = \frac{1}{17} \sum_{i=1983}^{2000} (h_i(R) - m(R))^2 .$$

The mean value $m(R)$ for rating R is our first guess of the potential default probability assigned to rating R . The standard deviation $s(R)$ gives us some insight about the volatility and therefore about the error we eventually make when believing that $m(R)$ is a good estimate of the default probability of R -rated obligors. Figure 1.1 shows the values $m(R)$ and $s(R)$ for the considered rating classes. Because even best rated obligors are not free of default risk, we write “not observed” in the cells corresponding to $m(R)$ and $s(R)$ for ratings $R=Aaa, Aa1, Aa2, A1, A2, A3$ (ratings where no defaults have been observed) in Figure 1.1.

2. Next, we plot the mean values $m(R)$ into a coordinate system, where the x -axis refers to the rating classes (here numbered from

⁵For example, one could look at investment and sub-investment grades separately.

1 (*Aaa*) to 16 (*B3*)). One can see in the chart in [Figure 1.1](#) that on a *logarithmic scale* the mean default frequencies $m(R)$ can be fitted by a regression line. Here we should add a comment that there is strong evidence from various empirical default studies that default frequencies grow exponentially with decreasing creditworthiness. For this reason we have chosen an exponential fit (linear on logarithmic scale). Using standard regression theory, see, e.g., [106] [Chapter 4](#), or by simply using any software providing basic statistical functions, one can easily obtain the following exponential function fitting our data:

$$DP(x) = 3 \times 10^{-5} e^{0.5075x} \quad (x = 1, \dots, 16).$$

3. As a last step, we use our regression equation for the estimation of default probabilities $DP(x)$ assigned to rating classes x ranging from 1 to 16. [Figure 1.1](#) shows our result, which we now call *a calibration of default probabilities to Moody's ratings*. Note that based on our regression even the best rating *Aaa* has a small but positive default probability. Moreover, our hope is that our regression analysis has smoothed out sampling errors from the historically observed data.

Although there is much more to say about default probabilities, we stop the discussion here. However, later on we will come back to default probabilities in various contexts.

1.1.2 The Exposure at Default

The EAD is the quantity in Equation (1. 2) specifying the exposure the bank does have to its borrower. In general, the exposure consists of two major parts, the *outstandings* and the *commitments*. The outstandings refer to the portion of the exposure already drawn by the obligor. In case of the borrower's default, the bank is exposed to the total amount of the outstandings. The commitments can be divided in two portions, *undrawn* and *drawn*, in the time before default. The total amount of commitments is the exposure the bank has promised to lend to the obligor at her or his request. Historical default experience shows that obligors tend to draw on committed lines of credit in times of financial distress. Therefore, the commitment is also subject to loss in case of the obligor's default, but only the drawn (prior default) amount

Rating	Mean	Standard-Deviation	Default Probability
Aaa	not observed	not observed	0.005%
Aa1	not observed	not observed	0.008%
Aa2	not observed	not observed	0.014%
Aa3	0.08%	0.33%	0.023%
A1	not observed	not observed	0.038%
A2	not observed	not observed	0.063%
A3	not observed	not observed	0.105%
Baa1	0.06%	0.19%	0.174%
Baa2	0.06%	0.20%	0.289%
Baa3	0.46%	1.16%	0.480%
Ba1	0.69%	1.03%	0.797%
Ba2	0.63%	0.86%	1.324%
Ba3	2.39%	2.35%	2.200%
B1	3.79%	2.49%	3.654%
B2	7.96%	6.08%	6.070%
B3	12.89%	8.14%	10.083%




FIGURE 1.1
Calibration of Moody's Ratings to Default Probabilities

of the commitments will actually contribute to the loss on loan. The fraction describing the decomposition of commitments in drawn and undrawn portions is a random variable due to the optional character commitments have (the obligor has the right but not the obligation to draw on committed lines of credit). Therefore it is natural to define the EAD by

$$\text{EAD} = \text{OUTST} + \gamma \times \text{COMM}, \quad (1.3)$$

where OUTST denotes the outstandings and COMM the commitments of the loan, and γ is the expected portion of the commitments likely to be drawn prior to default. More precisely, γ is the expectation of the random variable capturing the uncertain part of the EAD, namely the utilization of the undrawn part of the commitments. Obviously, γ takes place in the unit interval. Recall that we assume the EAD to be a deterministic (i.e., nonrandom) quantity. This is the reason why we directly deal with the expectation γ , hereby ignoring the underlying random variable.

In practice, banks will calibrate γ w.r.t. the creditworthiness of the borrower and the type of the facility involved.

Note that in many cases, commitments include various so-called *covenants*, which are embedded options either the bank has written to the obligor or reserved to itself. Such covenants may, for example, force an obligor in times of financial distress to provide more collateral⁶ or to renegotiate the terms of the loan. However, often the obligor has some informational advantage in that the bank recognizes financial distress of its borrowers with some delay. In case of covenants allowing the bank to close committed lines triggered by some early default indication, it really is a question of time if the bank picks up such indications early enough to react before the customer has drawn on her or his committed lines. The problem of appropriate and quick action of the lending institute is especially critical for obligors with former good credit quality, because banks tend to focus more on critical than on good customers regarding credit lines (bad customers get much more attention, because the bank is already “alarmed” and will be more sensitive in case of early warnings of financial instability). Any stochastic modeling of EAD should take these aspects into account.

⁶Collateral means assets securing a loan, e.g., mortgages, bonds, guarantees, etc. In case a loan defaults, the value of the collateral reduces the loss on the defaulted loan.

The *Basel Committee on Banking Supervision*⁷ in its recent consultative document [103] defines the EAD for on-balance sheet transactions to be identical to the *nominal* amount of the exposure.

For off-balance sheet transactions there are two approaches: For the *foundation approach* the committee proposes to define the EAD on commitments and revolving credits as 75% of the off-balance sheet amount of the exposure. For example, for a committed line of one billion Euro with current outstandings of 600 million, the EAD would be equal to $600 + 75\% \times 400 = 900$ million Euro.

For the *advanced approach*, the committee proposes that banks eligible for this approach will be permitted to use their own internal estimates of EAD for transactions with uncertain exposure. From this perspective it makes much sense for major banks to carefully think about some rigorous methodology for calibrating EAD to borrower- and facility-specific characteristics. For example, banks that are able to calibrate the parameter γ in (1. 3) on a finer scale will have more accurate estimates of the EAD, better reflecting the underlying credit risk. The more the determination of regulatory capital tends towards risk sensitivity, the more will banks with advanced methodology benefit from a more sophisticated calibration of EAD.

1.1.3 The Loss Given Default

The LGD of a transaction is more or less determined by “1 minus recovery rate”, i.e., the LGD quantifies the portion of loss the bank will really suffer in case of default. The estimation of such loss quotes is far from being straightforward, because recovery rates depend on many driving factors, for example on the *quality of collateral* (securities, mortgages, guarantees, etc.) and on the *seniority* of the bank’s claim on the borrower’s assets. This is the reason behind our convention to consider the loss given default as a random variable describing the *severity* of the loss of a facility type in case of default. The notion LGD then refers to the expectation of the severity.

A bank-external source for recovery data comes from the rating agencies. For example Moody’s [95] provides recovery values of defaulted bonds, hereby distinguishing between different seniorities.

⁷The Basle Committee coordinates the rules and guidelines for banking supervision. Its members are central banks and other national offices or government agencies responsible for banking supervision.

Unfortunately many banks do not have good internal data for estimating recovery rates. In fact, although LGD is a key driver of EL, there is in comparison with other risk drivers like the DP little progress made in moving towards a sophisticated calibration. There are initiatives (for example by the ISDA⁸ and other similar organisations) to bring together many banks for sharing knowledge about their practical LGD experience as well as current techniques for estimating it from historical data.

However, one can expect that in a few years LGD databases will have significantly improved, such that more accurate estimates of the LGD for certain banking products can be made.

1.2 Unexpected Loss

At the beginning of this chapter we introduced the EL of a transaction as an insurance or loss reserve in order to cover losses the bank expects from historical default experience. But holding capital as a cushion against *expected* losses is not enough. In fact, the bank should in addition to the expected loss reserve also save money for covering *unexpected* losses exceeding the average experienced losses from past history. As a measure of the magnitude of the deviation of losses from the EL, the standard deviation of the loss variable \tilde{L} as defined in (1.1) is a natural choice. For obvious reasons, this quantity is called the *Unexpected Loss* (UL), defined by

$$UL = \sqrt{\mathbb{V}[\tilde{L}]} = \sqrt{\mathbb{V}[EAD \times SEV \times L]}.$$

1.2.1 Proposition *Under the assumption that the severity and the default event D are uncorrelated, the unexpected loss of a loan is given by*

$$UL = EAD \times \sqrt{\mathbb{V}[SEV] \times DP + LGD^2 \times DP(1 - DP)}.$$

Proof. Taking $\mathbb{V}[X] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$ and $\mathbb{V}[1_D] = DP(1 - DP)$ into account, the assertion follows from a straightforward calculation. \square

⁸International Swap Dealers Association.

1.2.2 Remark Note that the assumption of zero correlation between severity and default event in Proposition 1.2.1 is not always realistic and often just made to obtain a first approximation to the “real” unexpected loss. In fact, it is not unlikely that on average the recovery rate of loans will drop if bad economic conditions induce an increase of default frequencies in the credit markets. Moreover, some types of collateral bear a significant portion of market risk, such that unfavourable market conditions (which might also be the reason for an increased number of default events) imply a decrease of the collateral’s market value. In Section 2.5 we discuss a case where the severity of losses and the default events are random variables driven by a common underlying factor.

Now, so far we have always looked at the credit risk of a single facility, although banks have to manage large portfolios consisting of many different products with different risk characteristics. We therefore will now indicate how one can model the total loss of a credit portfolio.

For this purpose we consider a portfolio consisting of m loans

$$\tilde{L}_i = \text{EAD}_i \times \text{SEV}_i \times L_i, \quad \text{with} \quad (1.4)$$

$$L_i = \mathbf{1}_{D_i}, \quad \mathbb{P}(D_i) = \text{DP}_i.$$

The *portfolio loss* is then defined as the random variable

$$\tilde{L}_{PF} = \sum_{i=1}^m \tilde{L}_i = \sum_{i=1}^m \text{EAD}_i \times \text{SEV}_i \times L_i. \quad (1.5)$$

Analogously to the “standalone” quantities EL and UL we now obtain portfolio quantities EL_{PF} and UL_{PF} , defined by the expectation respectively standard deviation of the portfolio loss. In case of EL we can use the additivity of expectations to obtain

$$\text{EL}_{PF} = \sum_{i=1}^m \text{EL}_i = \sum_{i=1}^m \text{EAD}_i \times \text{LGD}_i \times \text{DP}_i. \quad (1.6)$$

In case of the UL, additivity holds if the loss variables \tilde{L}_i are pairwise uncorrelated (see Bienaymé’s Theorem in [7] Chapter 8). If the loss variables are correlated, additivity is lost. Unfortunately this is the standard case, because correlations are “part of the game” and a main driver of credit risk. In fact, large parts of this book will essentially

be dealing with correlation modeling. The UL of a portfolio is the first risk quantity we meet where correlations respectively covariances play a fundamental role:

$$\begin{aligned}
 \text{UL}_{PF} &= \sqrt{\mathbb{V}[\tilde{L}_{PF}]} & (1.7) \\
 &= \sqrt{\sum_{i=1}^m \sum_{j=1}^m \text{EAD}_i \times \text{EAD}_j \times \text{Cov}[\text{SEV}_i \times L_i, \text{SEV}_j \times L_j]}.
 \end{aligned}$$

Looking at the special case where severities are constant, we can express the portfolio's UL by means of *default correlations*, namely

1.2.3 Proposition *For a portfolio with constant severities we have*

$$\begin{aligned}
 \text{UL}_{PF}^2 &= \sum_{i,j=1}^m \text{EAD}_i \times \text{EAD}_j \times \text{LGD}_i \times \text{LGD}_j \times \\
 &\quad \times \sqrt{\text{DP}_i(1 - \text{DP}_i)\text{DP}_j(1 - \text{DP}_j)} \rho_{ij}
 \end{aligned}$$

where $\rho_{ij} = \text{Corr}[L_i, L_j] = \text{Corr}[\mathbf{1}_{D_i}, \mathbf{1}_{D_j}]$ denotes the default correlation between counterparties i and j .

Proof. The proposition is obvious. \square

Before continuing we want for a moment to think about the meaning and interpretation of correlation. For simplicity let us consider a portfolio consisting of two loans with $\text{LGD} = 100\%$ and $\text{EAD} = 1$. We then only deal with L_i for $i = 1, 2$, and we set $\rho = \text{Corr}[L_1, L_2]$ and $p_i = \text{DP}_i$. Then, the squared UL of our portfolio is obviously given by

$$\text{UL}_{PF}^2 = p_1(1 - p_1) + p_2(1 - p_2) + 2\rho\sqrt{p_1(1 - p_1)}\sqrt{p_2(1 - p_2)}. \quad (1.8)$$

We consider three possible cases regarding the default correlation ρ :

- $\rho = 0$. In this case, the third term in (1.8) vanishes, such that UL_{PF} attains its minimum. This is called the case of *perfect diversification*. The concept of diversification is easily explained. Investing in many different assets generally reduces the overall portfolio risk, because usually it is very unlikely to see a large number of loans defaulting all at once. The less the loans in the portfolio have in common, the higher the chance that default of one obligor does not mean a lot to the economic future of other

loans in the portfolio. The case $\rho = 0$ is the case, where the loans in the portfolio are completely unrelated. Interpreting the UL as a substitute⁹ for portfolio risk, we see that this case minimizes the overall portfolio risk.

- $\rho > 0$. In this case our two counterparties are interrelated in that default of one counterparty increases the likelihood that the other counterparty will also default. We can make this precise by looking at the conditional default probability of counterparty 2 under the condition that obligor 1 already defaulted:

$$\begin{aligned} \mathbb{P}[L_2 = 1 \mid L_1 = 1] &= \frac{\mathbb{P}[L_1 = 1, L_2 = 1]}{\mathbb{P}[L_1 = 1]} = \frac{\mathbb{E}[L_1 L_2]}{p_1} \quad (1.9) \\ &= \frac{p_1 p_2 + \text{Cov}[L_1, L_2]}{p_1} = p_2 + \frac{\text{Cov}[L_1 L_2]}{p_1}. \end{aligned}$$

So we see that positive correlation respectively covariance leads to a conditional default probability higher (because of $\text{Cov}[L_1, L_2] > 0$) than the unconditional default probability p_2 of obligor 2. In other words, in case of positive correlation any default in the portfolio has an important implication on other facilities in the portfolio, namely that there might be more losses to be encountered. The extreme case in this scenario is the case of *perfect correlation* ($\rho = 1$). In the case of $p = p_1 = p_2$, Equation (1.8) shows that in the case of perfect correlation we have $\text{UL}_{PF} = 2\sqrt{p(1-p)}$, essentially meaning that our portfolio contains the risk of only one obligor but with double intensity (*concentration risk*). In this situation it follows immediately from (1.9) that default of one obligor makes the other obligor defaulting almost surely.

- $\rho < 0$. This is the mirrored situation of the case $\rho > 0$. We therefore only discuss the extreme case of perfect anti-correlation ($\rho = -1$). One then can view an investment in asset 2 as an almost *perfect hedge* against an investment in asset 1, if (additionally to $\rho = -1$) the characteristics (exposure, rating, etc.) of the two loans match. Admittedly, this terminology makes much

⁹Note that in contrast to the EL, the UL is the “true” uncertainty the bank faces when investing in a portfolio because it captures the deviation from the expectation.

more sense when following a *marked-to-market*¹⁰ approach to loan valuation, where an increase in market value of one of the loans immediately (under the assumption $\rho = -1$) would imply a decrease in market value of the other loan. However, from (1. 8) it follows that in the case of a perfect hedge the portfolio's UL completely vanishes ($UL_{PF} = 0$). This means that our perfect hedge (investing in asset 2 with correlation -1 w.r.t. a comparable and already owned asset 1) completely eliminates (neutralizes) the risk of asset 1.

We now turn to the important notion of economic capital.

1.2.1 Economic Capital

We have learned so far that banks should hold some capital cushion against unexpected losses. However, defining the UL of a portfolio as the *risk capital* saved for cases of financial distress is not the best choice, because there might be a significant likelihood that losses will exceed the portfolio's EL by more than one standard deviation of the portfolio loss. Therefore one seeks other ways to quantify risk capital, hereby taking a *target level of statistical confidence* into account.

The most common way to quantify risk capital is the concept of *economic capital*¹¹ (EC). For a prescribed level of confidence α it is defined as the α -quantile of the portfolio loss \tilde{L}_{PF} minus the EL of the portfolio,

$$EC_\alpha = q_\alpha - EL_{PF} , \quad (1. 10)$$

where q_α is the α -quantile of \tilde{L}_{PF} , determined by

$$q_\alpha = \inf\{q > 0 \mid \mathbb{P}[\tilde{L}_{PF} \leq q] \geq \alpha\} . \quad (1. 11)$$

For example, if the level of confidence is set to $\alpha = 99.98\%$, then the risk capital EC_α will (on average) be sufficient to cover unexpected losses

¹⁰In a marked-to-market framework loans do not live in a two-state world (default or survival) but rather are evaluated w.r.t. their *market value*. Because until today loans are only traded "over the counter" in *secondary markets*, a marked-to-market approach is more difficult to calibrate. For example, in Europe the secondary loan market is not as well developed as in the United States. However, due to the strongly increasing market of credit derivatives and securitised credit products, one can expect that there will be a transparent and well-developed market for all types of loans in a few years.

¹¹Synonymously called *Capital at Risk* (CaR) or credit *Value-at-Risk* (VaR) in the literature.

in 9,998 out of 10,000 years, hereby assuming a planning horizon of one year. Unfortunately, under such a calibration one can on the other side expect that in 2 out of 10,000 years the economic capital $EC_{99.98\%}$ will not be sufficient to protect the bank from insolvency. This is the downside when calibrating risk capital by means of quantiles. However, today most major banks use an EC framework for their internal credit risk model.

The reason for reducing the quantile q_α by the EL is due to the “best practice” of decomposing the total risk capital (i.e., the quantile) into a first part covering expected losses and a second part meant as a cushion against unexpected losses. Altogether the pricing of a loan typically takes several cost components into account. First of all, the price of the loan should include the costs of administrating the loan and maybe some kind of upfront fees. Second, expected losses are charged to the customer, hereby taking the creditworthiness captured by the customer’s rating into account. More risky customers have to pay a higher risk premium than customers showing high credit quality. Third, the bank will also ask for some compensation for taking the risk of unexpected losses coming with the new loan into the bank’s credit portfolio. The charge for unexpected losses is often calculated as the *contributory* EC of the loan in reference to the lending bank’s portfolio; see [Chapter 5](#). Note that there is an important difference between the EL and the EC charges: The EL charge is independent from the composition of the reference portfolio, whereas the EC charge strongly depends on the current composition of the portfolio in which the new loan will be included. For example, if the portfolio is already well diversified, then the EC charge as a cushion against unexpected losses does not have to be as high as it would be in the case for a portfolio in which, for example, the new loan would induce some concentration risk. Summarizing one can say the EL charges are *portfolio independent*, but EC charges are *portfolio dependent*. This makes the calculation of the contributory EC in pricing tools more complicated, because one always has to take the complete reference portfolio into account. *Risk contributions* will be discussed in [Chapter 5](#).

An alternative to EC is a risk capital based on *Expected Shortfall* (ESF). A capital definition according to ESF very much reflects an insurance point of view of the credit risk business. We will come back to ESF and its properties in [Chapter 5](#).




FIGURE 1.2
The portfolio loss distribution

1.2.2 The Loss Distribution

All risk quantities on a portfolio level are based on the portfolio loss variable \tilde{L}_{PF} . Therefore it does not come much as a surprise that the distribution of \tilde{L}_{PF} , the so-called *loss distribution* of the portfolio, plays a central role in credit risk management. In Figure 1.2 it is illustrated that all risk quantities of the credit portfolio can be identified by means of the loss distribution of the portfolio. This is an important observation, because it shows that in cases where the distribution of the portfolio loss can only be determined in an empirical way one can use empirical statistical quantities as a proxy for the respective “true” risk quantities.

In practice there are essentially two ways to generate a loss distribution. The first method is based on *Monte Carlo simulation*; the second is based on a so-called *analytical approximation*.

1.2.2.1 Monte Carlo Simulation of Losses

In a Monte Carlo simulation, losses are simulated and tabulated in form of a *histogram* in order to obtain an *empirical loss distribution*

of the underlying portfolio. The *empirical distribution function* can be determined as follows:

Assume we have simulated n potential portfolio losses $\tilde{L}_{PF}^{(1)}, \dots, \tilde{L}_{PF}^{(n)}$, hereby taking the driving distributions of the single loss variables and their correlations¹² into account. Then the empirical loss distribution function is given by

$$F(x) = \frac{1}{n} \sum_{j=1}^n \mathbf{1}_{[0,x]}(\tilde{L}_{PF}^{(j)}) . \tag{1. 12}$$

Figure 1.3 shows the shape of the density (histogram of the randomly generated numbers $(\tilde{L}_{PF}^{(1)}, \dots, \tilde{L}_{PF}^{(n)})$) of the empirical loss distribution of some test portfolio.

From the empirical loss distribution we can derive all the portfolio risk quantities introduced in the previous paragraphs. For example, the α -quantile of the loss distribution can directly be obtained from our simulation results $\tilde{L}_{PF}^{(1)}, \dots, \tilde{L}_{PF}^{(n)}$ as follows:

Starting with order statistics of $\tilde{L}_{PF}^{(1)}, \dots, \tilde{L}_{PF}^{(n)}$, say

$$\tilde{L}_{PF}^{(i_1)} \leq \tilde{L}_{PF}^{(i_2)} \leq \dots \leq \tilde{L}_{PF}^{(i_n)} ,$$

the α -quantile q_α of the empirical loss distribution (for any confidence level α) is given by

$$q_\alpha = \begin{cases} \alpha \tilde{L}_{PF}^{(i_{[n\alpha]})} + (1 - \alpha) \tilde{L}_{PF}^{(i_{[n\alpha]+1})} & \text{if } n\alpha \in \mathbb{N} \\ \tilde{L}_{PF}^{(i_{[n\alpha]})} & \text{if } n\alpha \notin \mathbb{N} \end{cases} \tag{1. 13}$$

where $[n\alpha] = \min \{k \in \{1, \dots, n\} \mid n\alpha \leq k\} .$

The economic capital can then be estimated by

$$EC_\alpha = q_\alpha - \frac{1}{n} \sum_{j=1}^n \tilde{L}_{PF}^{(j)} . \tag{1. 14}$$

In an analogous manner, any other risk quantity can be obtained by calculating the corresponding empirical statistics.

¹²We will later see that correlations are incorporated by means of a *factor model*.




FIGURE 1.3

An empirical portfolio loss distribution obtained by Monte Carlo simulation. The histogram is based on a portfolio of 2,000 middle-size corporate loans.

Approaching the loss distribution of a large portfolio by Monte Carlo simulation always requires a sound *factor model*; see [Section 1.2.3](#). The classical statistical reason for the existence of factor models is the wish to explain the variance of a variable in terms of underlying factors. Despite the fact that in credit risk we also wish to explain the variability of a firm's economic success in terms of global underlying influences, the necessity for factor models comes from two major reasons.

First of all, the correlation between single loss variables should be made interpretable in terms of *economic variables*, such that large losses can be explained in a sound manner. For example, a large portfolio loss might be due to the *downturn* of an industry common to many counterparties in the portfolio. Along this line, a factor model can also be used as a tool for *scenario analysis*. For example, by setting an industry factor to a particular fixed value and then starting the Monte Carlo simulation again, one can study the impact of a down- or upturn of the respective industry.

The second reason for the need of factor models is a reduction of the computational effort. For example, for a portfolio of 100,000 transactions, $\frac{1}{2} \times 100,000 \times 99,000$ correlations have to be calculated. In contrast, modeling the correlations in the portfolio by means of a factor model with 100 indices reduces the number of involved correlations by a factor of 1,000,000. We will come back to factor models in 1.2.3 and also in later chapters.

1.2.2.2 Analytical Approximation

Another approach to the portfolio loss distribution is by analytical approximation. Roughly speaking, the analytical approximation maps an actual portfolio with unknown loss distribution to an equivalent portfolio with known loss distribution. The loss distribution of the equivalent portfolio is then taken as a substitute for the “true” loss distribution of the original portfolio.

In practice this is often done as follows. Choose a family of distributions characterized by its first and second moment, showing the typical shape (i.e., right-skewed with fat tails¹³) of loss distributions as illustrated in [Figure 1.2](#).

¹³In our terminology, a distribution has *fat tails*, if its quantiles at high confidence are higher than those of a normal distribution with matching first and second moments.




FIGURE 1.4
Analytical approximation by some beta distribution

From the known characteristics of the original portfolio (e.g., rating distribution, exposure distribution, maturities, etc.) calculate the first moment (EL) and estimate the second moment (UL).

Note that the EL of the original portfolio usually can be calculated based on the information from the rating, exposure, and LGD distributions of the portfolio.

Unfortunately the second moment can not be calculated without any assumptions regarding the default correlations in the portfolio; see Equation (1.8). Therefore, one now has to make an assumption regarding an *average default correlation* ρ . Note that in case one thinks in terms of *asset value models*, see Section 24.1, one would rather guess an average *asset correlation* instead of a default correlation and then calculate the corresponding default correlation by means of Equation (2.5.1). However, applying Equation (1.8) by setting all default correlations ρ_{ij} equal to ρ will provide an estimated value for the original portfolio's UL.

Now one can choose from the parametrized family of loss distribution the distribution best matching the original portfolio w.r.t. first and second moments. This distribution is then interpreted as the loss distribution of an equivalent portfolio which was selected by a *moment matching* procedure.

Obviously the most critical part of an analytical approximation is the

Obviously the most critical part of an analytical approximation is the determination of the average asset correlation. Here one has to rely on practical experience with portfolios where the average asset correlation is known. For example, one could compare the original portfolio with a set of typical bank portfolios for which the average asset correlations are known. In some cases there is empirical evidence regarding a reasonable range in which one would expect the unknown correlation to be located. For example, if the original portfolio is a retail portfolio, then one would expect the average asset correlation of the portfolio to be a small number, maybe contained in the interval $[1\%, 5\%]$. If the original portfolio would contain loans given to large firms, then one would expect the portfolio to have a high average asset correlation, maybe somewhere between 40% and 60%. Just to give another example, the new *Basel Capital Accord* (see [Section 1.3](#)) assumes an average asset correlation of 20% for corporate loans; see [103]. In [Section 2.7](#) we estimate the average asset correlation in Moody's universe of rated corporate bonds to be around 25%. Summarizing we can say that calibrating¹⁴ an average correlation is on one hand a typical source of *model risk*, but on the other hand nevertheless often supported by some practical experience.

As an illustration of how the moment matching in an analytical approximation works, assume that we are given a portfolio with an EL of 30 bps and an UL of 22.5 bps, estimated from the information we have about some credit portfolio combined with some assumed average correlation.

Now, in [Section 2.5](#) we will introduce a typical family of two-parameter loss distributions used for analytical approximation. Here, we want to approximate the loss distribution of the original portfolio by a beta distribution, matching the first and second moments of the original portfolio. In other words, we are looking for a random variable

$$X \sim \beta(a, b) ,$$

representing the percentage portfolio loss, such that the parameters a and b solve the following equations:

$$0.003 = \mathbb{E}[X] = \frac{a}{a + b} \quad \text{and} \quad (1.15)$$

¹⁴The calibration might be more honestly called a “guestimate”, a mixture of a guess and an estimate.

$$0.00225^2 = \mathbb{V}[X] = \frac{ab}{(a+b)^2(a+b+1)} .$$

Hereby recall that the probability density φ_X of X is given by

$$\varphi_X(x) = \beta_{a,b}(x) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} x^{a-1}(1-x)^{b-1} \quad (1. 16)$$

($x \in [0, 1]$) with first and second moments

$$\mathbb{E}[X] = \frac{a}{a+b} \quad \text{and} \quad \mathbb{V}[X] = \frac{ab}{(a+b)^2(a+b+1)} .$$

Equations (1. 15) represent the moment matching addressing the “correct” beta distribution matching the first and second moments of our original portfolio. It turns out that $a = 1.76944$ and $b = 588.045$ solve equations (1. 15). [Figure 1.4](#) shows the probability density of the so calibrated random variable X .

The analytical approximation takes the random variable X as a proxy for the unknown loss distribution of the portfolio we started with. Following this assumption, the risk quantities of the original portfolio can be approximated by the respective quantities of the random variable X . For example, quantiles of the loss distribution of the portfolio are calculated as quantiles of the beta distribution. Because the “true” loss distribution is substituted by a closed-form, analytical, and well-known distribution, all necessary calculations can be done in fractions of a second. The price we have to pay for such convenience is that all calculations are subject to significant *model risk*. Admittedly, the beta distribution as shown in [Figure 1.4](#) has the shape of a loss distribution, but there are various two-parameter families of probability densities having the typical shape of a loss distribution. For example, some gamma distributions, the F-distribution, and also the distributions introduced in [Section 2.5](#) have such a shape. Unfortunately they all have different tails, such that in case one of them would approximate really well the unknown loss distribution of the portfolio, the others automatically would be the wrong choice. Therefore, the selection of an appropriate family of distributions for an analytical approximation is a remarkable source of model risk. Nevertheless there are some families of distributions that are established as *best practice* choices for particular cases. For example, the distributions in [Section 2.5](#) are a very natural choice for analytical approximations, because they are limit distributions of a well understood model.

In practice, analytical approximation techniques can be applied quite successfully to so-called *homogeneous portfolios*. These are portfolios where all transactions in the portfolio have comparable risk characteristics, for example, no exposure concentrations, default probabilities in a band with moderate bandwidth, only a few (better: one single!) industries and countries, and so on. There are many portfolios satisfying such constraints. For example, many retail banking portfolios and also many portfolios of smaller banks can be evaluated by analytical approximations with sufficient precision.

In contrast, a full Monte Carlo simulation of a large portfolio can last several hours, depending on the number of counterparties and the number of scenarios necessary to obtain sufficiently rich tail statistics for the chosen level of confidence.

The main advantage of a Monte Carlo simulation is that it accurately captures the correlations inherent in the portfolio instead of relying on a whole bunch of assumptions. Moreover, a Monte Carlo simulation takes into account all the different risk characteristics of the loans in the portfolio. Therefore it is clear that Monte Carlo simulation is the “state-of-the-art” in credit risk modeling, and whenever a portfolio contains quite different transactions from the credit risk point of view, one should not trust too much in the results of an analytical approximation.

1.2.3 Modeling Correlations by Means of Factor Models

Factor models are a well established technique from multivariate statistics, applied in credit risk models, for identifying underlying drivers of correlated defaults and for reducing the computational effort regarding the calculation of correlated losses. We start by discussing the basic meaning of a *factor*.

Assume we have two firms A and B which are positively correlated. For example, let A be DaimlerChrysler and B stand for BMW. Then, it is quite natural to explain the positive correlation between A and B by the correlation of A and B with an *underlying factor*; see [Figure 1.5](#). In our example we could think of the *automotive industry* as an underlying factor having significant impact on the economic future of the companies A and B. Of course there are probably some more underlying factors driving the riskiness of A and B. For example, DaimlerChrysler is to a certain extent also influenced by a factor for *Germany*, the *United States*, and eventually by some factors incorporating *Aero Space* and *Financial Companies*. BMW is certainly correlated

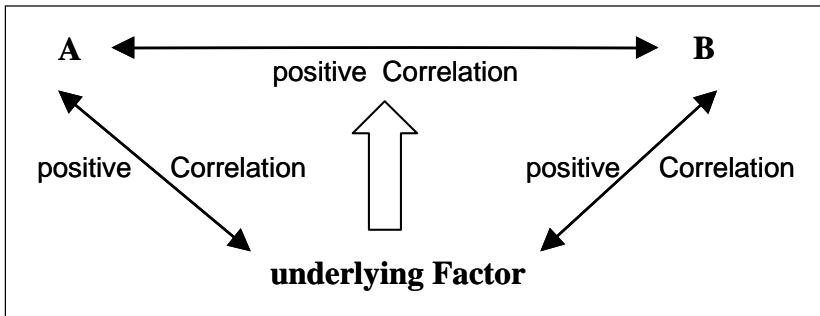


FIGURE 1.5
Correlation induced by an underlying factor

with a country factor for *Germany* and probably also with some other factors. However, the crucial point is that factor models provide a way to express the correlation between A and B exclusively by means of their correlation with common factors. As already mentioned in the previous section, we additionally wish underlying factors to be *interpretable* in order to identify the reasons why two companies experience a down- or upturn at about the same time. For example, assume that the automotive industry gets under pressure. Then we can expect that companies A and B also get under pressure, because their fortune is related to the automotive industry. The part of the *volatility* of a company's financial success (e.g., incorporated by its asset value process) related to systematic factors like industries or countries is called the *systematic risk* of the firm. The part of the firm's asset volatility that can not be explained by systematic influences is called the *specific* or *idiosyncratic risk* of the firm. We will make both notions precise later on in this section.

The *KMV[®]-Model* and *CreditMetricsTM*, two well-known industry models, both rely on a sound modeling of underlying factors. Before continuing let us take the opportunity to say a few words about the firms behind the models.

KMV is a small company, founded about 30 years ago and recently acquired by Moody's, which develops and distributes software for man-

aging credit portfolios. Their tools are based on a modification of Merton's *asset value model*, see [Chapter 3](#), and include a tool for estimating default probabilities (*Credit Monitor*TM) from market information and a tool for managing credit portfolios (*Portfolio Manager*TM). The first tool's main output is the *Expected Default Frequency*TM (EDF), which can nowadays also be obtained online by means of a newly developed web-based KMV-tool called *Credit Edge*TM. The main output of the *Portfolio Manager*TM is the loss distribution of a credit portfolio. Of course, both products have many more interesting features, and to us it seems that most large banks and insurance use at least one of the major KMV products. A reference to the basics of the KMV-Model is the survey paper by Crosbie [19].

*CreditMetrics*TM is a trademark of the *RiskMetrics*TM Group, a company which is a spin-off of the former *JPMorgan* bank, which now belongs to the *Chase Group*. The main product arising from the *CreditMetrics*TM framework is a tool called *CreditManager*TM, which incorporates a similar functionality as KMV's *Portfolio Manager*TM. It is certainly true that the technical documentation [54] of *CreditMetrics*TM was kind of a pioneering work and has influenced many bank-internal developments of credit risk models. The great success of the model underlying *CreditMetrics*TM is in part due to the philosophy of its authors Gupton, Finger, and Bhatia to make credit risk methodology available to a broad audience in a fully transparent manner.

Both companies continue to contribute to the market of credit risk models and tools. For example, the *RiskMetrics*TM Group recently developed a tool for the valuation of *Collateralized Debt Obligations*, and KMV recently introduced a new release of their *Portfolio Manager*TM PM2.0, hereby presenting some significant changes and improvements.

Returning to the subject of this section, we now discuss the factor models used in KMV's *Portfolio Manager*TM and *CreditMetrics*TM *CreditManager*TM. Both models incorporate the idea that every firm admits a process of *asset values*, such that default or survival of the firm depends on the state of the asset values at a certain planning horizon. If the process has fallen below a certain critical threshold, called the *default point* of the firm in KMV terminology, then the company has defaulted. If the asset value process is above the critical threshold, the firm survives. Asset value models have their roots in Merton's seminal paper [86] and will be explained in detail in [Chapter 3](#) and also to some extent in [Section 2.4.1](#).

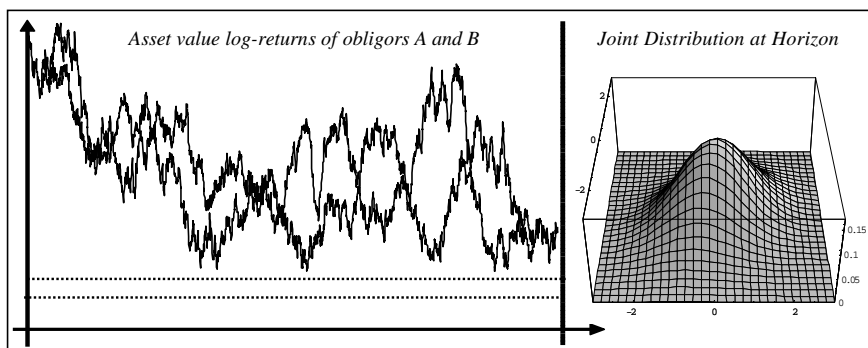


FIGURE 1.6
Correlated processes of obligor's asset value log-returns

Figure 1.6 illustrates the asset value model for two counterparties. Two correlated processes describing two obligor's asset values are shown. The correlation between the processes is called the *asset correlation*. In case the asset values are modeled by *geometric Brownian motions* (see Chapter 3), the asset correlation is just the correlation of the driving Brownian motions. At the planning horizon, the processes induce a bivariate asset value distributions. In the classical Merton model, where asset value processes are correlated geometric Brownian motions, the log-returns of asset values are normally distributed, so that the joint distribution of two asset value log-returns at the considered horizon is bivariate normal with a correlation equal to the asset correlation of the processes, see also Proposition 2.5.1. The dotted lines in Figure 1.6 indicate the critical thresholds or default points for each of the processes. Regarding the calibration of these default points we refer to Crosbie [19] for an introduction.

Now let us start with the KMV-Model, which is called the *Global Correlation ModelTM*. Regarding references we must say that KMV itself does not disclose the details of their factor model. But, nevertheless, a summary of the model can be found in the literature, see, e.g., Crouhy, Galai, and Mark [21]. Our approach to describing KMV's factor model is slightly different than typical presentations, because later on we will write the relevant formulas in a way supporting a convenient algorithm for the calculation of asset correlations.

Following Merton's model, KMV's factor model focuses on the asset value log-returns r_i of counterparties ($i = 1, \dots, m$) at a certain planning horizon (typically 1 year), admitting a representation

$$r_i = \beta_i \Phi_i + \varepsilon_i \quad (i = 1, \dots, m). \quad (1. 17)$$

Here, Φ_i is called the *composite factor* of firm i , because in multi-factor models Φ_i typically is a weighted sum of several factors. Equation (1. 17) is nothing but a standard *linear regression* equation, where the sensitivity coefficient, β_i , captures the linear correlation of r_i and Φ_i . In analogy to the *capital asset pricing model* (CAPM) (see, e.g., [21]) β is called the *beta* of counterparty i . The variable ε_i represents the *residual* part of r_i , essentially meaning that ε_i is the error one makes when substituting r_i by $\beta_i \Phi_i$. Merton's model lives in a log-normal world¹⁵, so that $\mathbf{r} = (r_1, \dots, r_m) \sim N(\boldsymbol{\mu}, \Gamma)$ is multivariate Gaussian with a correlation matrix Γ . The composite factors Φ_i and ε_i are accordingly also normally distributed. Another basic assumption is that ε_i is independent of the Φ_i 's for every i . Additionally the residuals ε_i are assumed to be uncorrelated¹⁶. Therefore, the returns r_i are exclusively correlated by means of their composite factors. This is the reason why Φ_i is thought of as the *systematic* part of r_i , whereas ε_i due to its independence from all other involved variables can be seen as a random effect just relevant for counterparty i . Now, in regression theory one usually decomposes the *variance* of a variable in a systematic and a specific part. Taking variances on both sides of Equation (1. 17) yields

$$\mathbb{V}[r_i] = \underbrace{\beta_i^2 \mathbb{V}[\Phi_i]}_{\text{systematic}} + \underbrace{\mathbb{V}[\varepsilon_i]}_{\text{specific}} \quad (i = 1, \dots, m). \quad (1. 18)$$

Because the variance of r_i captures the risk of unexpected movements of the asset value of counterparty i , the decomposition (1. 18) can be seen as a splitting of total risk of firm i in a *systematic* and a *specific* risk. The former captures the variability of r_i coming from the variability of the composite factor, which is $\beta_i^2 \mathbb{V}[\Phi_i]$; the latter arises from the variability of the residual variable, $\mathbb{V}[\varepsilon_i]$. Note that some people say *idiosyncratic* instead of *specific*.

¹⁵Actually, although the KMV-Model in principal follows Merton's model, it does not really work with Gaussian distributions but rather relies on an empirically calibrated framework; see Crosbie [19] and also [Chapter 3](#).

¹⁶Recall that in the case of Gaussian variables *uncorrelated* is equivalent to *independent*.

Alternatively to the beta of a firm one could also look at the *coefficient of determination* of the regression Equation (1. 17). The coefficient of determination quantifies how much of the variability of r_i can be explained by Φ_i . This quantity is usually called the *R-squared*, R^2 , of counterparty i and constitutes an important input parameter in all credit risk models based on asset values. It is usually defined as the systematic part of the variance of the standardized¹⁷ returns $\tilde{r}_i = (r_i - \mathbb{E}[r_i])/\sqrt{\mathbb{V}[r_i]}$, namely

$$R_i^2 = \frac{\beta_i^2 \mathbb{V}[\Phi_i]}{\mathbb{V}[r_i]} \quad (i = 1, \dots, m). \quad (1. 19)$$

The residual part of the total variance of the standardized returns \tilde{r}_i is then given by $1 - R_i^2$, thereby quantifying the percentage value of the specific risk of counterparty i .

Now we will look more carefully at the composite factors. The decomposition of a firm's variance in a systematic and a specific part is the first level in KMV's three-level factor model; see [Figure 1.7](#). The subsequent level is the decomposition of the firm's composite Φ in industry and country indices.

Before writing down the level-2 decomposition, let us rewrite Equation (1. 17) in vector notation¹⁸, more convenient for further calculations. For this purpose denote by $\beta = (\beta_{ij})_{1 \leq i, j \leq m}$ the diagonal matrix in $\mathbb{R}^{m \times m}$ with $\beta_{ij} = \beta_i$ if $i = j$ and $\beta_{ij} = 0$ if $i \neq j$. Equation (1. 17) then can be rewritten in vector notation as follows:

$$\mathbf{r} = \beta \Phi + \boldsymbol{\varepsilon}, \quad (1. 20)$$

$$\Phi^T = (\Phi_1, \dots, \Phi_m), \quad \boldsymbol{\varepsilon}^T = (\varepsilon_1, \dots, \varepsilon_m).$$

For the second level, KMV decomposes every Φ_i w.r.t. an *industry* and *country* breakdown,

$$\Phi_i = \sum_{k=1}^K w_{i,k} \Psi_k \quad (i = 1, \dots, m), \quad (1. 21)$$

where $\Psi_1, \dots, \Psi_{K_0}$ are industry indices and $\Psi_{K_0+1}, \dots, \Psi_K$ are country indices. The coefficients $w_{i,1}, \dots, w_{i,K_0}$ are called the *industry weights*

¹⁷That is, normalized in order to have mean zero and variance one.

¹⁸Note that in the sequel we write vectors as column vectors.

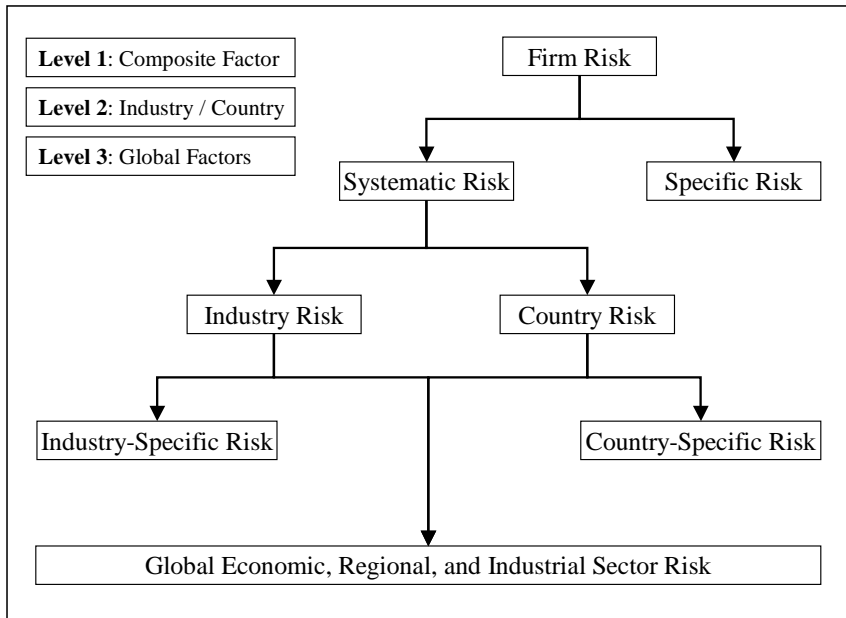


FIGURE 1.7

Three-level factor structure in KMV's Global Correlation Model™, see also comparable presentations in the literature, e.g. Figure 9.9. in [21].

and the coefficients $w_{i,K_0+1}, \dots, w_{i,K}$ are called the *country weights* of counterparty i . It is assumed that $w_{i,k} \geq 0$ for all i and k , and that

$$\sum_{k=1}^{K_0} w_{i,k} = \sum_{k=K_0+1}^K w_{i,k} = 1 \quad (i = 1, \dots, m).$$

In vector notation, (1. 20) combined with (1. 21) can be written as

$$\mathbf{r} = \beta \mathbf{W} \Psi + \varepsilon, \tag{1. 22}$$

where $\mathbf{W} = (w_{i,k})_{i=1, \dots, m; k=1, \dots, K}$ denotes the matrix of industry and country weights for the counterparties in the portfolio, and $\Psi^T = (\Psi_1, \dots, \Psi_K)$ means the vector of industry and country indices. This constitutes the second level of the Global Correlation ModelTM.

At the third and last level, a representation by a weighted sum of *independent global factors* is constructed for representing industry and country indices,

$$\Psi_k = \sum_{n=1}^N b_{k,n} \Gamma_n + \delta_k \quad (k = 1, \dots, K), \tag{1. 23}$$

where δ_k denotes the Ψ_k -specific residual. Such a decomposition is typically done by a *principal components analysis* (PCA) of the industry and country indices. In vector notation, (1. 23) becomes

$$\Psi = \mathbf{B} \Gamma + \delta \tag{1. 24}$$

where $\mathbf{B} = (b_{k,n})_{k=1, \dots, K; n=1, \dots, N}$ denotes the matrix of *industry and country betas*, $\Gamma^T = (\Gamma_1, \dots, \Gamma_N)$ is the global factor vector, and $\delta^T = (\delta_1, \dots, \delta_K)$ is the vector of industry and country residuals. Combining (1. 22) with (1. 24), we finally obtain

$$\mathbf{r} = \beta \mathbf{W} (\mathbf{B} \Gamma + \delta) + \varepsilon. \tag{1. 25}$$

So in the KMV-Model, the vector of the portfolio's returns $\mathbf{r}^T = (r_1, \dots, r_m)$ can conveniently be written by means of underlying factors. Note that for computational purposes Equation (1. 25) is the most convenient one, because the underlying factors are independent. In contrast, for an economic interpretation and for scenario analysis one would rather prefer Equation (1. 22), because the industry and country indices are easier to interpret than the global factors constructed by

PCA. In fact, the industry and country indices have a clear economic meaning, whereas the global factors arising from a PCA are of synthetic type. Although they admit some vague interpretation as shown in [Figure 1.7](#), their meaning is not as clear as is the case for the industry and country indices.

As already promised, the calculation of asset returns in the model as introduced above is straightforward now. First of all, we standardize the asset value log-returns,

$$\tilde{r}_i = \frac{r_i - \mathbb{E}[r_i]}{\sigma_i} \quad (i = 1, \dots, m)$$

where σ_i denotes the volatility of the asset value log-return of counterparty i . From Equation (1. 25) we then obtain a representation of standardized log-returns,

$$\tilde{r}_i = \frac{\beta_i}{\sigma_i} \tilde{\Phi}_i + \frac{\tilde{\varepsilon}_i}{\sigma_i} \quad \text{where} \quad \mathbb{E}[\tilde{\Phi}_i] = \mathbb{E}[\tilde{\varepsilon}_i] = 0 . \quad (1. 26)$$

Now, the asset correlation between two counterparties is given by

$$\text{Corr}[\tilde{r}_i, \tilde{r}_j] = \mathbb{E}[\tilde{r}_i \tilde{r}_j] = \frac{\beta_i \beta_j}{\sigma_i \sigma_j} \mathbb{E}[\tilde{\Phi}_i \tilde{\Phi}_j] \quad (1. 27)$$

because KMV assumes the residuals $\tilde{\varepsilon}_i$ to be uncorrelated and independent of the composite factors. For calculation purposes it is convenient to get rid of the volatilities σ_i and the betas β_i in Equation (1. 27). This can be achieved by replacing the betas by the R-squared parameters of the involved firms. From Equation (1. 19) we know that

$$R_i^2 = \frac{\beta_i^2}{\sigma_i^2} \mathbb{V}[\Phi_i] \quad (i = 1, \dots, m). \quad (1. 28)$$

Therefore, Equation (1. 27) combined with (1. 28) yields

$$\begin{aligned} \text{Corr}[\tilde{r}_i, \tilde{r}_j] &= \frac{R_i}{\sqrt{\mathbb{V}[\Phi_i]}} \frac{R_j}{\sqrt{\mathbb{V}[\Phi_j]}} \mathbb{E}[\tilde{\Phi}_i \tilde{\Phi}_j] \quad (1. 29) \\ &= \frac{R_i}{\sqrt{\mathbb{V}[\tilde{\Phi}_i]}} \frac{R_j}{\sqrt{\mathbb{V}[\tilde{\Phi}_j]}} \mathbb{E}[\tilde{\Phi}_i \tilde{\Phi}_j] \end{aligned}$$

because by construction we have $\mathbb{V}[\Phi_i] = \mathbb{V}[\tilde{\Phi}_i]$.

Based on Equation (1. 25) we can now easily compute asset correlations according to (1. 29). After standardization, (1. 25) changes to

$$\tilde{\mathbf{r}} = \tilde{\boldsymbol{\beta}}\mathbf{W}(\mathbf{B}\tilde{\boldsymbol{\Gamma}} + \tilde{\boldsymbol{\delta}}) + \tilde{\boldsymbol{\varepsilon}}, \quad (1. 30)$$

where $\tilde{\boldsymbol{\beta}} \in \mathbb{R}^{m \times m}$ denotes the matrix obtained by scaling every diagonal element in $\boldsymbol{\beta}$ by $1/\sigma_i$, and

$$\mathbb{E}[\tilde{\boldsymbol{\Gamma}}] = 0, \quad \mathbb{E}[\tilde{\boldsymbol{\varepsilon}}] = 0, \quad \mathbb{E}[\tilde{\boldsymbol{\delta}}] = 0.$$

Additionally, the residuals $\tilde{\boldsymbol{\delta}}$ and $\tilde{\boldsymbol{\varepsilon}}$ are assumed to be uncorrelated and independent of $\tilde{\boldsymbol{\Gamma}}$. We can now calculate asset correlations according to (1. 29) just by computing the matrix

$$\mathbb{E}[\tilde{\boldsymbol{\Phi}}\tilde{\boldsymbol{\Phi}}^T] = \mathbf{W} \left[\mathbf{B}\mathbb{E}[\tilde{\boldsymbol{\Gamma}}\tilde{\boldsymbol{\Gamma}}^T]\mathbf{B}^T + \mathbb{E}[\tilde{\boldsymbol{\delta}}\tilde{\boldsymbol{\delta}}^T] \right] \mathbf{W}^T \quad (1. 31)$$

because the matrix of standardized composite factors is given by $\tilde{\boldsymbol{\Phi}} = \mathbf{W}(\mathbf{B}\tilde{\boldsymbol{\Gamma}} + \tilde{\boldsymbol{\delta}})$. Let us quickly prove that (1. 31) is true. By definition, we have

$$\begin{aligned} \mathbb{E}[\tilde{\boldsymbol{\Phi}}\tilde{\boldsymbol{\Phi}}^T] &= \mathbb{E} \left[\mathbf{W}(\mathbf{B}\tilde{\boldsymbol{\Gamma}} + \tilde{\boldsymbol{\delta}}) \left(\mathbf{W}(\mathbf{B}\tilde{\boldsymbol{\Gamma}} + \tilde{\boldsymbol{\delta}}) \right)^T \right] \\ &= \mathbf{W} \mathbb{E} \left[(\mathbf{B}\tilde{\boldsymbol{\Gamma}} + \tilde{\boldsymbol{\delta}})(\mathbf{B}\tilde{\boldsymbol{\Gamma}} + \tilde{\boldsymbol{\delta}})^T \right] \mathbf{W}^T \\ &= \mathbf{W} \left(\mathbf{B}\mathbb{E}[\tilde{\boldsymbol{\Gamma}}\tilde{\boldsymbol{\Gamma}}^T]\mathbf{B}^T + \underbrace{\mathbf{B}\mathbb{E}[\tilde{\boldsymbol{\Gamma}}\tilde{\boldsymbol{\delta}}^T]}_{=0} + \underbrace{\mathbb{E}[\tilde{\boldsymbol{\delta}}(\mathbf{B}\tilde{\boldsymbol{\Gamma}})^T]}_{=0} + \mathbb{E}[\tilde{\boldsymbol{\delta}}\tilde{\boldsymbol{\delta}}^T] \right) \mathbf{W}^T. \end{aligned}$$

The two expectations above vanish due to our orthogonality assumptions. This proves (1. 31). Note that in equation (1. 31), $\mathbb{E}[\tilde{\boldsymbol{\Gamma}}\tilde{\boldsymbol{\Gamma}}^T]$ is a diagonal matrix (because we are dealing with orthogonal global factors) with diagonal elements $\mathbb{V}[\Gamma_n]$ ($n = 1, \dots, N$), and $\mathbb{E}[\tilde{\boldsymbol{\delta}}\tilde{\boldsymbol{\delta}}^T]$ is a diagonal matrix with diagonal elements $\mathbb{V}[\delta_k]$ ($k = 1, \dots, K$). Therefore, the calculation of asset correlations according to (1. 31) can conveniently be implemented in case one knows the variances of global factors, the variances of industry and country residuals, and the beta of the industry and country indices w.r.t. the global factors. KMV customers have access to this information and can use Equation (1. 31) for calculating asset correlations. In fact, KMV also offers a tool for calculating the asset correlation between any two firms contained in the KMV database, namely a tool called *GCorrTM*. However, Equation (1. 31) nevertheless

is useful to know, because it allows for calculating the asset correlation between firms even if they are not contained in the KMV database. In such cases one has to estimate the industry and country weights and the R-squared of the two firms. Applying Equation (1. 31) for $m=2$ immediately yields the respective asset correlation corresponding to the Global Correlation ModelTM.

The factor model of CreditMetricsTM is quite similar to KMV's factor model just described. So there is no need to start all over again, and we refer to the CreditMetricsTM Technical Document [54] for more information. However, there are two fundamental differences between the models which are worthwhile and important to be mentioned:

First, KMV's Global Correlation ModelTM is calibrated w.r.t. *asset* value processes, whereas the factor model of CreditMetricsTM uses *equity* processes instead of asset value processes, thereby taking *equity correlations as a proxy for asset correlations*; see [54], page 93. We consider this difference to be fundamental, because a very important feature of the KMV-Model is that it really manages the admittedly difficult process of translating equity and market information into asset values; see [Chapter 3](#).

Second, CreditMetricsTM uses indices¹⁹ referring to a *combination* of some industry in some particular country, whereas KMV considers industries and countries separately. For example, a German automotive company in the CreditMetricsTM factor model would get a 100%-weight w.r.t. an index describing the *German automotive* industry, whereas in the Global Correlation ModelTM this company would have industry *and* country weights equal to 100% w.r.t. an *automotive index* and a *country index* representing Germany. Both approaches are quite different and have their own advantages and disadvantages.

1.3 Regulatory Capital and the Basel Initiative

This section needs a disclaimer upfront. Currently, the regulatory capital approach is in the course of revision, and to us it does not make much sense to report in detail on the current state of the discussion. In

¹⁹MSCI indices; see www.msci.com.

recent documentation (from a technical point of view, [103] is a reference to the current approach based on internal ratings) many paragraphs are subject to be changed. In this book we therefore only briefly indicate what regulatory capital means and give some overview of the evolving process of regulatory capital definitions.

In 1983 the banking supervision authorities of the main industrialized countries (G7) agreed on rules for banking regulation, which should be incorporated into national regulation laws. Since the national regulators discussed these issues, hosted and promoted by the *Bank of International Settlement* located in *Basel* in Switzerland, these rules were called *The Basel Capital Accord*.

The best known rule therein is the *8-percent rule*. Under this rule, banks have to prove that the capital they hold is larger than 8% of their so-called *risk-weighted assets* (RWA), calculated for all balance sheet positions. This rule implied that the capital basis for banks was mainly driven by the exposure of the lendings to their customers. The RWA were calculated by a simple *weighting scheme*. Roughly speaking, for loans to any government institution the *risk weight* was set to 0%, reflecting the broad opinion that the governments of the world's industry nations are likely to meet their financial obligations. The risk weight for lendings to OECD banks was fixed at 20%. Regarding corporate loans, the committee agreed on a risk weight of 100%, no matter if the borrowing firm is a more or less risky obligor. The RWA were then calculated by adding up all of the bank's weighted credit exposures, yielding a *regulatory capital* of $8\% \times \text{RWA}$.

The main weakness of this capital accord was that it made no distinction between obligors with different creditworthiness. In 1988 an amendment to this Basel Accord opened the door for the use of internal models to calculate the regulatory capital for off-balance sheet positions in the *trading book*. The trading book was mostly seen as containing deals bearing *market risk*, and therefore the corresponding internal models captured solely the market risk in the trading business. Still, corporate bonds and derivatives contributed to the RWA, since the default risk was not captured by the market risk models.

In 1997 the *Basel Committee on Banking Supervision* allowed the banks to use so-called *specific risk models*, and the eligible instruments did no longer fall under the 8%-rule. Around that time regulators recognized that banks already internally used sophisticated models to handle the credit risk for their balance sheet positions with an emphasis

on default risk. These models were quite different from the standard specific risk models. In particular they produced a loss distribution of the entire portfolio and did not so much focus on the volatility of the spreads as in most of the specific risk models.

At the end of the 20th century, the Basel Committee started to look intensively at the models presented in this book. However, in their recent proposal they decided not to introduce these models into the regulatory framework at this stage. Instead they promote the use of *internal ratings* as main drivers of regulatory capital. Despite the fact that they also use some formulas and insights from the study of portfolio models (see [103]), in particular the notion of asset correlations and the CreditMetricsTM/KMV one-factor model (see Section 2.5), the recently proposed regulatory framework does not take bank-specific portfolio effects into account²⁰. In the documentation of the so-called *Internal Ratings-Based Approach* (IRB) [103], the only quantity reflecting the portfolio as a whole is the *granularity adjustment*, which is still in discussion and likely to be removed from the capital accord. In particular, *industrial or regional diversification effects* are not reflected by regulatory capital if the new Basel Accord in its final form, which will be negotiated in the near future, keeps the approach documented in [103]. So in order to better capture the risk models widely applied in banks all over the world, some further evolution of the Basel process is necessary.

²⁰A loan A requires the same amount of capital, independent of the bank granting the loan, thus ignoring the possibility that loan A increases the concentration risk in the bank's own portfolio but not in another.

Chapter 2

Modeling Correlated Defaults

In this chapter we will look at default models from a more abstract point of view, hereby providing a framework in which today's industry models can be embedded. Let us start with some general remarks.

Regarding random variables and probabilities we repeat our remark from the beginning of the previous chapter by saying that we always assume that an appropriate probability space $(\Omega, \mathcal{F}, \mathbb{P})$ has been chosen, reflecting the “probabilistic environment” necessary to make the respective statement.

Without loss of generality we will always assume a valuation horizon of one year. Let's say we are looking at a credit portfolio with m counterparties. Every counterparty in the portfolio admits a rating R_i as of today, and by means of some rating calibration as explained in [Section 1.1.1.1](#) we know the default probability p_i corresponding to rating R_i . One year from today the rating of the considered counterparty may have changed due to a change in its creditworthiness. Such a rating change is called a *rating migration*. More formally we denote the range of possible ratings by $\{0, \dots, d\}$, where $d \in \mathbb{N}$ means the *default state*,

$$R_i \in \{0, \dots, d\} \quad \text{and} \quad p_i = \mathbb{P}[R_i \rightarrow d],$$

where the notation $R \rightarrow R'$ denotes a rating migration from rating R to rating R' within one year. In this chapter we will focus on a *two-state* approach, essentially meaning that we restrict ourselves to a setting where

$$d = 1, \quad L_i = R_i \in \{0, 1\}, \quad p_i = \mathbb{P}[L_i = 1].$$

Two-state models neglect the possibility of rating changes; only *default* or *survival* is considered. However, generalizing a two-state to a multi-state model is straightforward and will be done frequently in subsequent chapters.

In [Chapter 1](#) we defined loss variables as indicators of default events; see [Section 1.1](#). In the context of two-state models, an approach by means of *Bernoulli* random variables is most natural. When it comes to

the modeling of defaults, CreditMetrics™ and the KMV-Model follow this approach. Another common approach is the modeling of defaults by *Poisson* random variables. CreditRisk⁺ (see Section 2.4.2) from Credit Suisse Financial Products is among the major industry models and a well-known representative of this approach.

There are attempts to bring Bernoulli and Poisson models in a common mathematical framework (see, e.g., Gordy [51] and Hickman and Koyluoglu [74]) and to some extent there are indeed relations and common roots of the two approaches; see Section 2.3. However, in [12] it is shown that the models are not really compatible, because the corresponding *mixture models* (Bernoulli respectively Poisson variables have to be *mixed* in order to introduce correlations into the models) generate loss distributions with significant tail differences. See Section 2.5.3.

Today we can access a rich literature investigating general frameworks for modeling correlated defaults and for embedding the existing industry models in a more abstract framework. See, e.g., Crouhy, Galai and Mark [20], Gordy [51], Frey and McNeil [45], and Hickman and Koyluoglu [74], just to mention a few references.

For the sequel we make a notational convention. Bernoulli random variables will always be denoted by L , whereas Poisson variables will be denoted by L' . In the following section we first look at the Bernoulli¹ model, but then also turn to the case of Poissonian default variables. In Section 2.3 we briefly compare both approaches.

2.1 The Bernoulli Model

A vector of random variables $\mathbf{L} = (L_1, \dots, L_m)$ is called a (Bernoulli) *loss statistics*, if all marginal distributions of \mathbf{L} are Bernoulli:

$$L_i \sim B(1; p_i), \quad \text{i.e.,} \quad L_i = \begin{cases} 1 & \text{with probability } p_i \\ 0 & \text{with probability } 1 - p_i \end{cases} .$$

¹Note that the Bernoulli model benefits from the convenient property that the mixture of Bernoulli variables again yields a Bernoulli-type random variable.

The *loss* resp. *percentage loss* of \mathbf{L} is defined² as

$$L = \sum_{i=1}^m L_i \quad \text{resp.} \quad \frac{L}{m} .$$

The probabilities $p_i = \mathbb{P}[L_i = 1]$ are called *default probabilities* of \mathbf{L} .

The reasoning underlying our terminology is as follows:

A credit portfolio is nothing but a collection of, say m , transactions or deals with certain counterparties. Every counterparty involved creates basically (in a two-state model) two future scenarios: Either the counterparty defaults³, or the counterparty survives⁴. In the case of default of obligor i the indicator variable L_i equals 1; in the case of survival we have $L_i = 0$. In this way, every portfolio generates a natural loss statistics w.r.t. the particular valuation horizon (here, one year). The variable L defined above is then called the *portfolio loss*, no matter if quoted as an absolute or percentage value.

Before we come to more interesting cases we should for the sake of completeness briefly discuss the quite unrealistic case of independent defaults.

The most simple type of a loss statistic can be obtained by assuming a *uniform default probability* p and the lack of dependency between counterparties. More precisely, under these assumptions we have

$$L_i \sim B(1; p) \quad \text{and} \quad (L_i)_{i=1, \dots, m} \text{ independent.}$$

In this case, the absolute portfolio loss L is a *convolution* of i.i.d. Bernoulli variables and therefore follows a *binomial* distribution with parameters m and p , $L \sim B(m; p)$.

If the counterparties are still assumed to be independent, but this time admitting different default probabilities,

$$L_i \sim B(1; p_i) \quad \text{and} \quad (L_i)_{i=1, \dots, m} \text{ independent,}$$

²Note that in the sequel we sometimes write L for denoting the *gross* loss as well as the *percentage* loss of a loss statistics. But from the context the particular meaning of L will always be clear.

³Note that there exist various default definitions in the banking world; as long as nothing different is said, we always mean by default a *payment default on any financial obligation*.

⁴Meets the financial expectations of the bank regarding contractually promised cash flows.

we again obtain the portfolio loss L as a convolution of the single loss variables, but this time with first and second moments

$$\mathbb{E}[L] = \sum_{i=1}^m p_i \quad \text{and} \quad \mathbb{V}[L] = \sum_{i=1}^m p_i(1 - p_i) . \quad (2. 1)$$

This follows from $\mathbb{E}[L_i] = p_i$, $\mathbb{V}[L_i] = p_i(1 - p_i)$, and the additivity of expectations resp. variances⁵.

Now, it is well known that in probability theory *independence* makes things easy. For example the *strong law of large numbers* works well with independent variables and the *central limit theorem* in its most basic version lives from the assumption of independence. If in credit risk management we could assume independence between counterparties in a portfolio, we could – due to the central limit theorem – assume that the portfolio loss (approximable) is a Gaussian variable, at least for large portfolios. In other words, we would never be forced to work with Monte Carlo simulations, because the portfolio loss would conveniently be given in a closed (namely Gaussian) form with well-known properties.

Unfortunately in credit risk modeling we can not expect to find independence of losses. Moreover, it will turn out that *correlation is the central challenge* in credit portfolio risk. Therefore, we turn now to more realistic elaborations of loss statistics.

One basic idea for modeling correlated defaults (by mixing) is the randomization of the involved default probabilities in a correlated manner. We start with a so-called standard *binary mixture model*; see Joe [67] for an introduction to this topic.

2.1.1 A General Bernoulli Mixture Model

Following our basic terminology, we obtain the loss of a portfolio from a loss statistics $\mathbf{L} = (L_1, \dots, L_m)$ with Bernoulli variables $L_i \sim B(1; P_i)$. But now we think of the loss probabilities as random variables $\mathbf{P} = (P_1, \dots, P_m) \sim \mathbf{F}$ with some distribution function \mathbf{F} with support in $[0, 1]^m$. Additionally, we assume that conditional on a realization $\mathbf{p} = (p_1, \dots, p_m)$ of \mathbf{P} the variables L_1, \dots, L_m are independent. In more

⁵For having additivity of variances it would be sufficient that the involved random variables are pairwise uncorrelated and integrable (see [7], [Chapter 8](#)).

mathematical terms we express the *conditional independence* of the losses by writing

$$L_i | P_i = p_i \sim B(1; p_i), \quad (L_i | \mathbf{P} = \mathbf{p})_{i=1, \dots, m} \text{ independent.}$$

The (unconditional) joint distribution of the L_i 's is then determined by the probabilities

$$\begin{aligned} & \mathbb{P}[L_1 = l_1, \dots, L_m = l_m] & (2. 2) \\ & = \int_{[0,1]^m} \prod_{i=1}^m p_i^{l_i} (1 - p_i)^{1-l_i} d\mathbf{F}(p_1, \dots, p_m), \end{aligned}$$

where $l_i \in \{0, 1\}$. The first and second moments of the single losses L_i are given by

$$\mathbb{E}[L_i] = \mathbb{E}[P_i], \quad \mathbb{V}[L_i] = \mathbb{E}[P_i] (1 - \mathbb{E}[P_i]) \quad (i = 1, \dots, m). \quad (2. 3)$$

The first equality is obvious from (2. 2). The second identity can be seen as follows:

$$\begin{aligned} \mathbb{V}[L_i] &= \mathbb{V}[\mathbb{E}[L_i | \mathbf{P}]] + \mathbb{E}[\mathbb{V}[L_i | \mathbf{P}]] & (2. 4) \\ &= \mathbb{V}[P_i] + \mathbb{E}[P_i(1 - P_i)] = \mathbb{E}[P_i] (1 - \mathbb{E}[P_i]) . \end{aligned}$$

The covariance between single losses obviously equals

$$\text{Cov}[L_i, L_j] = \mathbb{E}[L_i L_j] - \mathbb{E}[L_i] \mathbb{E}[L_j] = \text{Cov}[P_i, P_j] . \quad (2. 5)$$

Therefore, the *default correlation* in a Bernoulli mixture model is

$$\text{Corr}[L_i, L_j] = \frac{\text{Cov}[P_i, P_j]}{\sqrt{\mathbb{E}[P_i] (1 - \mathbb{E}[P_i])} \sqrt{\mathbb{E}[P_j] (1 - \mathbb{E}[P_j])}} . \quad (2. 6)$$

Equation (2. 5) respectively Equation (2. 6) show that the dependence between losses in the portfolio is fully captured by the covariance structure of the multivariate distribution \mathbf{F} of \mathbf{P} . [Section 2.4](#) presents some examples for a meaningful specification of \mathbf{F} .

2.1.2 Uniform Default Probability and Uniform Correlation

For portfolios where all exposures are of approximately the same size and type in terms of risk, it makes sense to assume a uniform default probability and a uniform correlation among transactions in the portfolio. As already mentioned in [Section 1.2.2.2](#), retail portfolios and some portfolios of smaller banks are often of a quite homogeneous structure, such that the assumption of a uniform default probability and a simple correlation structure does not harm the outcome of calculations with such a model. In the literature, portfolios with uniform default probability and uniform default correlation are called *uniform portfolios*. Uniform portfolio models generate perfect candidates for analytical approximations. For example, the distributions in [Section 2.5](#) establish a typical family of two-parameter loss distributions used for analytical approximations.

The assumption of *uniformity* yields *exchangeable*⁶ Bernoulli variables $L_i \sim B(1; P)$ with a random default probability $P \sim F$, where F is a distribution function with support in $[0, 1]$. We assume conditional independence of the L_i 's just as in the general case. The joint distribution of the L_i 's is then determined by the probabilities

$$\mathbb{P}[L_1 = l_1, \dots, L_m = l_m] = \int_0^1 p^k (1-p)^{m-k} dF(p), \quad (2.7)$$

$$\text{where } k = \sum_{i=1}^m l_i \quad \text{and} \quad l_i \in \{0, 1\}.$$

The probability that exactly k defaults occur is given by

$$\mathbb{P}[L = k] = \binom{m}{k} \int_0^1 p^k (1-p)^{m-k} dF(p). \quad (2.8)$$

Of course, Equations (2.3) and (2.6) have their counterparts in this special case of Bernoulli mixtures: The uniform default probability of borrowers in the portfolio obviously equals

$$\bar{p} = \mathbb{P}[L_i = 1] = \mathbb{E}[L_i] = \int_0^1 p dF(p) \quad (2.9)$$

⁶That is, $(L_1, \dots, L_m) \sim (L_{\pi(1)}, \dots, L_{\pi(m)})$ for any permutation π .

and the uniform default correlation of two different counterparties is given by

$$\begin{aligned} \rho = \text{Corr}[L_i, L_j] &= \frac{\mathbb{P}[L_i = 1, L_j = 1] - \bar{p}^2}{\bar{p}(1 - \bar{p})} \quad (2. 10) \\ &= \frac{\int_0^1 p^2 dF(p) - \bar{p}^2}{\bar{p}(1 - \bar{p})}. \end{aligned}$$

Note that in the course of this book we typically use “ ρ ” to denote *default* correlations and “ ϱ ” for denoting *asset* correlations.

We now want to briefly discuss some immediate consequences of Equation (2. 10). First of all it implies that

$$\text{Corr}[L_i, L_j] = \frac{\mathbb{V}[P]}{\bar{p}(1 - \bar{p})} \quad (\text{recall: } P \sim F).$$

This shows that the higher the volatility of P , the higher the default correlation inherent in the corresponding Bernoulli loss statistics.

Additionally, it implies that the dependence between the L_i ’s is either positive or zero, because variances are nonnegative. In other words, in this model we can not implement some negative dependencies between the default risks of obligors.

The case $\text{Corr}[L_i, L_j] = 0$ happens if and only if the variance of F vanishes to zero, essentially meaning that there is no randomness at all regarding P . In such a case, F is a *Dirac measure* $\varepsilon_{\bar{p}}$, concentrated in \bar{p} , and the absolute portfolio loss L follows a binomial distribution with default probability \bar{p} .

The other extreme case regarding (2. 10), $\text{Corr}[L_i, L_j] = 1$, implies a “rigid” behaviour of single losses in the portfolio: Either all counterparties default or all counterparties survive simultaneously. The corresponding distribution F of P is then a Bernoulli distribution, such that $P = 1$ with probability \bar{p} and $P = 0$ with probability $1 - \bar{p}$. This means that sometimes (such events occur with probability \bar{p}), all counterparties default and the total portfolio exposure is lost. In other scenarios (occurring with probability $1 - \bar{p}$), all obligors survive and not even one dollar is lost. The rigidity of loss statistics is “perfect” in this situation.

Realistic scenarios live somewhere between the two discussed extreme cases $\text{Corr}[L_i, L_j] = 0$ and $\text{Corr}[L_i, L_j] = 1$.

2.2 The Poisson Model

In the case of the Poisson approach, defaults of counterparties $i = 1, \dots, m$ are modeled by Poisson-distributed random variables

$$L'_i \sim Pois(\lambda_i), \quad L'_i \in \{0, 1, 2, \dots\}, \quad p_i = \mathbb{P}[L'_i \geq 1], \quad (2. 11)$$

where p_i again denotes the default probability of obligor i . Note that (2. 11) allows for multiple defaults of a single obligor. The likelihood of the event that obligor i defaults more than once is given by

$$\mathbb{P}[L'_i \geq 2] = 1 - e^{-\lambda_i}(1 + \lambda_i),$$

which is typically a small number. For example, in the case of $\lambda_i = 0.01$ we would obtain $\mathbb{P}[L'_i \geq 2] = 0.5$ basispoints. In other words, when simulating a Poisson-distributed default variable with $\lambda_i = 0.01$ we can expect that only 1 out of 20,000 scenarios is not applicable because of a multiple default. On the other side, for obligors with good credit quality (for example, a AAA-borrower with a default probability of 2 basispoints), a multiple-default probability of 0.5 basispoints is a relatively high number.

The intensity λ_i is typically quite close to the default probability p_i , due to

$$p_i = \mathbb{P}[L'_i \geq 1] = 1 - e^{-\lambda_i} \approx \lambda_i \quad (2. 12)$$

for small values of λ_i . Equation (2. 12) shows that the one-year default probability equals the probability that an *exponential waiting time* with intensity λ_i takes place in the first year.

In general, the sum of independent variables $L'_1 \sim Pois(\lambda_1)$, $L'_2 \sim Pois(\lambda_2)$ has distribution⁷ $Pois(\lambda_1 + \lambda_2)$. Assuming independence, the portfolio's total number of losses would be given by

$$L' = \sum_{i=1}^m L'_i \sim Pois\left(\sum_{i=1}^m \lambda_i\right). \quad (2. 13)$$

Correlation is introduced into the model by again following a mixture approach, in this case with Poisson variables (see also Joe [67], [Section 7.2](#)).

⁷More generally, $(Pois(\lambda))_{\lambda \geq 0}$ is a *convolution semigroup*; see, e.g., [7].

2.2.1 A General Poisson Mixture Model

Now the loss statistics is a random vector $\mathbf{L}' = (L'_1, \dots, L'_m)$ of Poisson random variables $L'_i \sim Pois(\Lambda_i)$, where $\mathbf{\Lambda} = (\Lambda_1, \dots, \Lambda_m)$ is a random vector with some distribution function \mathbf{F} with support in $[0, \infty)^m$. Additionally, we assume that conditional on a realization $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_m)$ of $\mathbf{\Lambda}$ the variables L'_1, \dots, L'_m are independent:

$$L'_i |_{\Lambda_i = \lambda_i} \sim Pois(\lambda_i), \quad (L'_i |_{\mathbf{\Lambda} = \boldsymbol{\lambda}})_{i=1, \dots, m} \text{ independent.}$$

The (unconditional) joint distribution of the variables L'_i is given by

$$\begin{aligned} & \mathbb{P}[L'_1 = l'_1, \dots, L'_m = l'_m] \quad (2.14) \\ &= \int_{[0, \infty)^m} e^{-(\lambda_1 + \dots + \lambda_m)} \prod_{i=1}^m \frac{\lambda_i^{l'_i}}{l'_i!} d\mathbf{F}(\lambda_1, \dots, \lambda_m), \end{aligned}$$

where $l'_i \in \{0, 1, 2, \dots\}$. Analogously to the Bernoulli case we obtain

$$\mathbb{E}[L'_i] = \mathbb{E}[\Lambda_i] \quad (i = 1, \dots, m) \quad (2.15)$$

$$\mathbb{V}[L'_i] = \mathbb{V}[\mathbb{E}[L'_i | \mathbf{\Lambda}]] + \mathbb{E}[\mathbb{V}[L'_i | \mathbf{\Lambda}]] = \mathbb{V}[\Lambda_i] + \mathbb{E}[\Lambda_i].$$

Again we have $\text{Cov}[L'_i, L'_j] = \text{Cov}[\Lambda_i, \Lambda_j]$, and the correlation between defaults is given by

$$\text{Corr}[L'_i, L'_j] = \frac{\text{Cov}[\Lambda_i, \Lambda_j]}{\sqrt{\mathbb{V}[\Lambda_i] + \mathbb{E}[\Lambda_i]} \sqrt{\mathbb{V}[\Lambda_j] + \mathbb{E}[\Lambda_j]}}. \quad (2.16)$$

In the same manner as in the Bernoulli model this shows that correlation is exclusively induced by means of the distribution function \mathbf{F} of the *random intensity vector* $\mathbf{\Lambda}$.

2.2.2 Uniform Default Intensity and Uniform Correlation

Analogously to the Bernoulli model, one can introduce a *Poisson uniform portfolio model* by restriction to one uniform intensity and one uniform correlation among transactions in the portfolio. More explicitly, the uniform portfolio model in the Poisson case is given by Poisson variables $L'_i \sim Pois(\Lambda)$ with a random intensity $\Lambda \sim F$, where F is a distribution function with support in $[0, \infty)$, and the L'_i 's are

assumed to be conditionally independent. The joint distribution of the L_i 's is given by

$$\mathbb{P}[L'_1 = l'_1, \dots, L'_m = l'_m] = \int_0^\infty e^{-m\lambda} \frac{\lambda^{(l'_1 + \dots + l'_m)}}{l'_1! \dots l'_m!} dF(\lambda). \quad (2.17)$$

Because (see the beginning of [Section 2.2](#)) conditional on $\Lambda = \lambda$ the portfolio loss is again a Poisson distribution with intensity $m\lambda$, the probability of exactly k defaults equals

$$\begin{aligned} \mathbb{P}[L' = k] &= \int_0^\infty \mathbb{P}[L' = k \mid \Lambda = \lambda] dF(\lambda) \quad (2.18) \\ &= \int_0^\infty e^{-m\lambda} \frac{m^k \lambda^k}{k!} dF(\lambda). \end{aligned}$$

Again, note that due to the unbounded support of the Poisson distribution the absolute loss L' can exceed the number of “physically” possible defaults. As already mentioned at the beginning of this section, the probability of a multiple-defaults event is small for typical parametrizations. In the Poisson framework, the uniform default probability of borrowers in the portfolio is defined by

$$\begin{aligned} \bar{p} = \mathbb{P}[L'_i \geq 1] &= \int_0^\infty \mathbb{P}[L'_i \geq 1 \mid \Lambda = \lambda] dF(\lambda) \quad (2.19) \\ &= \int_0^\infty (1 - e^{-\lambda}) dF(\lambda). \end{aligned}$$

The counterpart of Formula (2.16) is

$$\text{Corr}[L'_i, L'_j] = \frac{\mathbb{V}[\Lambda]}{\mathbb{V}[\Lambda] + \mathbb{E}[\Lambda]} \quad (i \neq j). \quad (2.20)$$

Formula (2.20) is especially intuitive if seen in the context of *dispersion*, where the dispersion of a distribution is its *variance to mean ratio*

$$D_X = \frac{\mathbb{V}[X]}{\mathbb{E}[X]} \quad \text{for any random variable } X. \quad (2.21)$$

The dispersion of the Poisson distribution is equal to 1. Therefore, the Poisson distribution is kind of a benchmark when deciding about *overdispersion* ($D_X > 1$) respectively *underdispersion* ($D_X < 1$). In

general, nondegenerate⁸ *Poisson mixtures are overdispersed* due to (2. 15). This is a very important property of Poisson mixtures, because before using such a model for credit risk measurement one has to make sure that overdispersion can be observed in the data underlying the calibration of the model. Formula (2. 20) can be interpreted by saying that the correlation between the number of defaults of different counterparties increases with the dispersion of the random intensity Λ . For proving this statement we write Formula (2. 20) in the form

$$\text{Corr}[L'_i, L'_j] = \frac{D_\Lambda}{D_\Lambda + 1} \quad (i \neq j). \quad (2. 22)$$

From (2. 22) it follows that *an increase in dispersion increases the mixture effect, which, in turn, strengthens the dependence between obligor's defaults.*

2.3 Bernoulli Versus Poisson Mixture

The *law of small numbers*⁹ implies that for large m and small p

$$B(m; p) \approx \text{Pois}(pm).$$

Setting $\lambda = pm$, this shows that under the assumption of independent defaults the portfolio absolute gross loss $L = \sum L_i$ of a Bernoulli loss statistics (L_1, \dots, L_m) with a uniform default probability p can be approximated by a Poisson variable $L' \sim \text{Pois}(\lambda)$. But the law of small numbers is by no means an argument strong enough to support the unfortunately widespread opinion that Bernoulli and Poisson approaches are more or less compatible. In order to show that both approaches have significant systematic differences, we turn back to the default correlations induced by the models; see (2. 6), combined with (2. 4), and (2. 16). In the Bernoulli case we have

$$\begin{aligned} \text{Corr}[L_i, L_j] &= \quad (2. 23) \\ &= \frac{\text{Cov}[P_i, P_j]}{\sqrt{\mathbb{V}[P_i] + \mathbb{E}[P_i(1 - P_i)]} \sqrt{\mathbb{V}[P_j] + \mathbb{E}[P_j(1 - P_j)]}}, \end{aligned}$$

⁸The random intensity Λ is not concentrated in a single point, $P_\Lambda \neq \varepsilon_\lambda$.

⁹That is, approximation of binomial distributions by means of Poisson distributions.

whereas in the Poisson case we obtain

$$\text{Corr}[L'_i, L'_j] = \frac{\text{Cov}[\Lambda_i, \Lambda_j]}{\sqrt{\mathbb{V}[\Lambda_i] + \mathbb{E}[\Lambda_i]} \sqrt{\mathbb{V}[\Lambda_j] + \mathbb{E}[\Lambda_j]}} . \quad (2. 24)$$

Looking only at the driving random variables P_i, P_j respectively Λ_i, Λ_j , we see that in the denominators of (2. 23) and (2. 24) we compare

$$\mathbb{V}[P_i] + \mathbb{E}[P_i(1 - P_i)] = \mathbb{V}[P_i] + \mathbb{E}[P_i] - \mathbb{E}[P_i^2] \quad (2. 25)$$

$$\text{with} \quad \mathbb{V}[\Lambda_i] + \mathbb{E}[\Lambda_i] .$$

Now, analogous to the deterministic case (2. 12), we can – even in the random case – expect P_i and Λ_i to be of the same order of magnitude. To keep things simple, let us for a moment assume that P_i and Λ_i have the same first and second moments. In this case Equation (2. 25) combined with (2. 23) and (2. 24) shows that the Bernoulli model always induces a higher default correlation than the Poisson model. But higher default correlations result in fatter tails of the corresponding loss distributions. In other words one could say that given equal first and second moments of P_i and Λ_i , the expectations of L_i and L'_i will match, but the variance of L'_i will always exceed the variance of L_i , thereby inducing lower default correlations.

So there is a *systematic* difference between the Bernoulli and Poisson mixture models. In general one can expect that for a given portfolio the Bernoulli model yields a loss distribution with a fatter tail than a comparably (e.g., by a first and second moment matching) calibrated Poisson model. This difference is also reflected by the industry models from CreditMetrics™ / KMV Corporation (Portfolio Manager) and Credit Suisse Financial Products (CreditRisk⁺). In [Section 2.5.3](#) we come back to this issue.

2.4 An Overview of Today's Industry Models

In the last five years, several industry models for measuring credit portfolio risk have been developed. Besides the main commercial models we find in large international banks various so-called *internal models*, which in most cases are more or less inspired by the well-known commercial products. For most of the industry models it is easy to

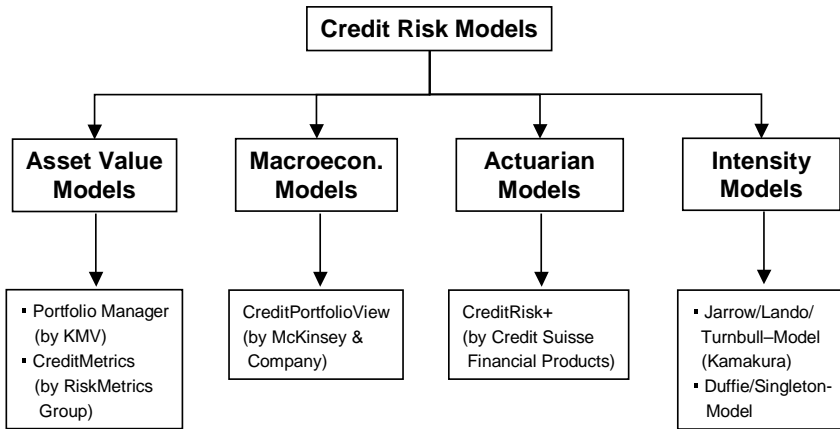


FIGURE 2.1
Today's Best-Practice Industry Models.

find some kind of technical documentation describing the mathematical framework of the model and giving some idea about the underlying data and the calibration of the model to the data. An exception is KMV's Portfolio ManagerTM, where most of the documentation is proprietary or confidential. However, even for the KMV-Model the basic idea behind the model can be explained without reference to nonpublic sources. In [Section 1.2.3](#) we already briefly introduced CreditMetricsTM and the KMV-Model in the context of asset value factor models. In [Chapter 3](#) we present a mathematically more detailed but nontechnical introduction to the type of asset value models KMV is incorporating.

Before looking at the main models, we want to provide the reader with a brief overview. [Figure 2.1](#) shows the four main types of industry models and indicates the companies behind them. [Table 2.1](#) summarizes the main differences between the models.

CreditRisk⁺ could alternatively be placed in the group of intensity models, because it is based on a Poisson mixture model incorporating random intensities. Nevertheless in [Figure 2.1](#) we prefer to stress the difference between CreditRisk⁺ and the *dynamic* intensity models based on intensity *processes* instead of on a *static* intensity.

Dynamic intensity models will be briefly discussed in [Section 2.4.4](#)

TABLE 2.1: Overview: Main Differences between Industry Models.

	KMV-Model	CreditMetrics	Credit* Portfolio View	CreditRisk⁺	Intensity Models
Risk Driver	Asset Value Process	Asset Value Process	Macro-economic Factors	Default Intensity	Intensity Process
Definition of Risk	Distance to Default (DtD)	Mark-to-Model of Loan Value	Mark-to-Model of Loan Value	Default Risk only	Default Risk only
Risk Scale	DtD on contin. Scale	Down/Upgrade and Default	Down/Upgrade and Default	Default	Default
Transition Probabilities	EDF-Concept, high migration probabilities	Historic Rating Changes, e.g. from S&P	Stochastic, via Macrofactors	Not Implemented	Not Implemented
Correlations	Asset Value Factor Model	Equity Value Factor Model	Implicit by Macroeconomy	Implicit by Sectors	Correlated Intensity Proc.
Severity	Stochastic (Beta-Distr.) and Fixed	Stochastic (Beta-Distr.) and Fixed	Stochastic, Empirically Calibrated	Deterministic LGD, Stoch. Modifications	Deterministic LGD

* Credit Portfolio View in the CPV-Macro mode. In the CPV-Direct mode, segment-specific default probabilities are drawn from a gamma distribution instead of simulating macroeconomic factors as input into a logit function representing a segments conditional default probability.

and to some extent in the context of *securitizations*. From a mathematician’s point of view they provide a “mathematically beautiful” approach to credit risk modeling, but from the introductory point of view we adopted for writing this book we must say that an appropriate presentation of dynamic intensity models is beyond the scope of the book. We therefore decided to provide the reader only with some references to the literature combined with introductory remarks about the approach; see [Section 2.4.4](#)

In this book, our discussion of CreditPortfolioView¹⁰ (CPV) is kept shorter than our presentation of CreditMetricsTM, the KMV-Model, and the actuarial model CreditRisk⁺. The reason for not going too much into details is that CPV can be considered as a *general framework* for credit risk modeling, which is then tailor-made for client’s needs. In our presentation we mainly focus on the *systematic risk model* of CPV.

2.4.1 CreditMetricsTM and the KMV-Model

For some background on CreditMetricsTM and KMV we refer to [Section 1.2.3](#). Note that for both models we focus on their “default-only mode”, hereby ignoring the fact that both models incorporate a *mark-to-model* approach. In the default-only mode, both models are of Bernoulli type, deciding about default or survival of a firm by comparing the firm’s asset value at a certain horizon with some critical threshold. If the firm value at the horizon is below this threshold, then the firm is considered to be in default. If the firm value is above the threshold, the firm survived the considered time period. In more mathematical terms, for m counterparties denote their asset value at the considered valuation horizon $t = T$ by $A_T^{(i)}$. It is assumed that for every company i there is a critical threshold C_i such that the firm defaults in the period $[0, T]$ if and only if $A_T^{(i)} < C_i$. In the framework of Bernoulli loss statistics A_T can be viewed as a latent variable driving the default event. This is realized by defining

$$L_i = \mathbf{1}_{\{A_T^{(i)} < C_i\}} \sim B\left(1; \mathbb{P}[A_T^{(i)} < C_i]\right) \quad (i = 1, \dots, m). \quad (2. 26)$$

In both models it is assumed that the asset value process is dependent on underlying factors reflecting industrial and regional influences, thereby driving the economic future of the firm. For the convenience

¹⁰By McKinsey & Company.

of the reader we now recall some formulas from [Section 1.2.3](#). The parametrization w.r.t. underlying factors typically is implemented at the standardized¹¹ log-return level, i.e., the asset value log-returns $\log(A_T^{(i)}/A_0^{(i)})$ after standardization admit a representation¹²

$$r_i = R_i\Phi_i + \varepsilon_i \quad (i = 1, \dots, m). \quad (2. 27)$$

Here R_i is defined as in (1. 28), Φ_i denotes the firm’s *composite factor*, and ε_i is the *firm-specific effect* or (as it is also often called) the *idiosyncratic* part of the firm’s asset value log-return. In both models, the factor Φ_i is a *superposition* of many different *industry* and *country* indices. *Asset correlations* between counterparties are exclusively captured by the correlation between the respective composite factors. The specific effects are assumed to be independent among different firms and independent of the composite factors. The quantity R_i^2 reflects how much of the volatility of r_i can be explained by the volatility of the composite factor Φ_i . Because the composite factor is a superposition of *systematic* influences, namely industry and country indices, R_i^2 quantifies the *systematic risk* of counterparty i .

In CreditMetricsTM as well as in the (parametric) KMV world, asset value log-returns are assumed to be normally distributed, such that due to standardization we have

$$r_i \sim N(0, 1), \quad \Phi_i \sim N(0, 1), \quad \text{and} \quad \varepsilon_i \sim N(0, 1 - R_i^2).$$

We are now in a position to rewrite (2. 26) in the following form:

$$L_i = \mathbf{1}_{\{r_i < c_i\}} \sim B(1; \mathbb{P}[r_i < c_i]) \quad (i = 1, \dots, m), \quad (2. 28)$$

where c_i is the threshold corresponding to C_i after exchanging $A_T^{(i)}$ by r_i . Applying (2. 27), the condition $r_i < c_i$ can be written as

$$\varepsilon_i < c_i - R_i\Phi_i \quad (i = 1, \dots, m). \quad (2. 29)$$

Now, in both models, the standard valuation horizon is $T = 1$ year. Denoting the one-year default probability of obligor i by p_i , we naturally have $p_i = \mathbb{P}[r_i < c_i]$. Because $r_i \sim N(0, 1)$ we immediately obtain

$$c_i = N^{-1}[p_i] \quad (i = 1, \dots, m), \quad (2. 30)$$

¹¹Shifted and scaled in order to obtain a random variable with mean zero and standard deviation one.

¹²Note that for reasons of a simpler notation we here write r_i for the standardized log-returns, in contrast to the notation in [Section 1.2.3](#), where we wrote \tilde{r}_i .

where $N[\cdot]$ denotes the cumulative standard normal distribution function. Scaling the idiosyncratic component towards a standard deviation of one, this changes (2. 29) into

$$\tilde{\varepsilon}_i < \frac{N^{-1}[p_i] - R_i\Phi_i}{\sqrt{1 - R_i^2}}, \quad \tilde{\varepsilon}_i \sim N(0, 1). \quad (2. 31)$$

Taking into account that $\tilde{\varepsilon}_i \sim N(0, 1)$, we altogether obtain the following representation for the one-year default probability of obligor i conditional on the factor Φ_i :

$$p_i(\Phi_i) = N \left[\frac{N^{-1}[p_i] - R_i\Phi_i}{\sqrt{1 - R_i^2}} \right] \quad (i = 1, \dots, m). \quad (2. 32)$$

The only random part of (2. 32) is the composite factor Φ_i . Conditional on $\Phi_i = z$, we obtain the conditional one-year default probability by

$$p_i(z) = N \left[\frac{N^{-1}[p_i] - R_i z}{\sqrt{1 - R_i^2}} \right]. \quad (2. 33)$$

Combined with (2. 28) this shows that we are in a Bernoulli mixture setting exactly the same way as elaborated in [Section 2.1.1](#). More formally we can – analogously to (2. 2) – specify the portfolio loss distribution by the probabilities (here we assume again $l_i \in \{0, 1\}$)

$$\begin{aligned} & \mathbb{P}[L_1 = l_1, \dots, L_m = l_m] \quad (2. 34) \\ &= \int_{[0,1]^m} \prod_{i=1}^m q_i^{l_i} (1 - q_i)^{1-l_i} d\mathbf{F}(q_1, \dots, q_m), \end{aligned}$$

where the distribution function \mathbf{F} is now explicitly given by

$$\mathbf{F}(q_1, \dots, q_m) = N_m [p_1^{-1}(q_1), \dots, p_m^{-1}(q_m); \Gamma], \quad (2. 35)$$

where $N_m[\cdot; \Gamma]$ denotes the cumulative multivariate centered Gaussian distribution with correlation matrix Γ , and $\Gamma = (\varrho_{ij})_{1 \leq i, j \leq m}$ means the asset correlation matrix of the log-returns r_i according to (2. 27).

In case that the composite factors are represented by a weighted sum of industry and country indices $(\Psi_j)_{j=1, \dots, J}$ of the form

$$\Phi_i = \sum_{j=1}^J w_{ij} \Psi_j \quad (2. 36)$$

(see [Section 1.2.3](#)), the conditional default probabilities (2. 33) appear as

$$p_i(z) = N \left[\frac{N^{-1}[p_i] - R_i(w_{i1}\psi_1 + \dots + w_{iJ}\psi_J)}{\sqrt{1 - R_i^2}} \right], \quad (2. 37)$$

with industry and country index realizations $(\psi_j)_{j=1,\dots,J}$. By varying these realizations and then recalculating the conditional probabilities (2. 37) one can perform a simple *scenario stress testing*, in order to study the impact of certain changes of industry or country indices on the default probability of some obligor.

2.4.2 CreditRisk⁺

CreditRisk⁺ is a credit risk model developed by Credit Suisse Financial Products (CSFP). It is more or less based on a typical insurance mathematics approach, which is the reason for its classification as an *actuarial model*. Regarding its mathematical background, the main reference is the CreditRisk⁺ Technical Document [18]. In light of this chapter one could say that CreditRisk⁺ is a typical representative of the group of Poisson mixture models. In this paragraph we only summarize the model, focussing on *defaults only* and not on *losses* in terms of money, but in [Chapter 4](#) a more comprehensive introduction (taking exposure distributions into account) is presented.

As mixture distribution CreditRisk⁺ incorporates the *gamma distribution*. Recall that the gamma distribution is defined by the probability density

$$\gamma_{\alpha,\beta}(x) = \frac{1}{\beta^\alpha \Gamma(\alpha)} e^{-x/\beta} x^{\alpha-1} \quad (x \geq 0),$$

where $\Gamma(\cdot)$ denotes¹³ the gamma function. The first and second moments of a gamma-distributed random variable Λ are

$$\mathbb{E}[\Lambda] = \alpha\beta, \quad \mathbb{V}[\Lambda] = \alpha\beta^2; \quad (2. 38)$$

see [Figure 2.2](#) for an illustration of gamma densities.

¹³We will also write $X \sim \Gamma(\alpha, \beta)$ for any gamma-distributed random variable X with parameters α and β . Additionally, we use Γ to denote the correlation matrix of a multivariate normal distribution. However, it should be clear from the context which current meaning the symbol Γ has.




Figure 2.2: Shape of Gamma Distributions for parameters $(\alpha, \beta) \in \{(2, 1/2), (5, 1/5)\}$.

Instead of incorporating a factor model (as we have seen it in the case of CreditMetricsTM and KMV's Portfolio ManagerTM in [Section 1.2.3](#)), CreditRisk⁺ implements a so-called *sector model*. However, somehow one can think of a sector as a “factor-inducing” entity, or – as the CreditRisk⁺ Technical Document [18] says it – every sector could be thought of as generated by a single underlying factor. In this way, sectors and factors are somehow comparable objects. From an interpretational point of view, sectors can be identified with industries, countries, or regions, or any other systematic influence on the economic performance of counterparties with a positive weight in this sector. Each sector $s \in \{1, \dots, m_S\}$ has its own *gamma-distributed random intensity* $\Lambda^{(s)} \sim \Gamma(\alpha_s, \beta_s)$, where the variables $\Lambda^{(1)}, \dots, \Lambda^{(m_S)}$ are assumed to be *independent*.

Now let us assume that a credit portfolio of m loans to m different obligors is given. In the *sector model* of CreditRisk⁺, every obligor i admits a breakdown into sector weights $w_{is} \geq 0$ with $\sum_{s=1}^{m_S} w_{is} = 1$, such that w_{is} reflects the sensitivity of the default intensity of obligor i to the systematic default risk arising from sector s . The risk of sector s is captured by two parameters: The first driver is the *mean default*

intensity of the sector,

$$\lambda_{(s)} = \mathbb{E}[\Lambda^{(s)}] = \alpha_s \beta_s ;$$

see also (4. 18) in Chapter 4. The second driver is the *default intensity's volatility*

$$\sigma_{(s)} = \mathbb{V}[\Lambda^{(s)}] = \alpha_s \beta_s^2 .$$

In Section 4.3.2 we indicate some possible approaches for calibrating the sector parameters $\lambda_{(s)}$ and $\sigma_{(s)}$. Every obligor i admits a random default intensity Λ_i with mean value $\mathbb{E}[\Lambda_i] = \lambda_i$, which could be calibrated to the obligor's one-year default probability by means of Formula (2. 12). The sector parametrization of Λ_i is as follows:

$$\Lambda_i = \sum_{s=1}^{m_S} w_{is} \lambda_i \frac{\Lambda^{(s)}}{\lambda_{(s)}} \quad (i = 1, \dots, m); \quad (2. 39)$$

see also Formula (4. 29). This shows that two obligors are correlated if and only if there is at least one sector such that both obligors have a positive sector weight with respect to this sector. Only in such cases two obligors admit a common source of *systematic default risk*. Note that (2. 39) is consistent with the assumption that λ_i equals the expected default intensity of obligor i . The default risk of obligor i is then modeled by a mixed Poisson random variable L'_i with random intensity Λ_i .

Note that in accordance with (2. 12) any conditional default intensity of obligor i arising from realizations $\theta_1, \dots, \theta_{m_S}$ of the sector's default intensities $\Lambda^{(1)}, \dots, \Lambda^{(m_S)}$ generates a conditional one-year default probability $p_i(\theta_1, \dots, \theta_{m_S})$ of obligor i by setting

$$\begin{aligned} p_i(\theta_1, \dots, \theta_{m_S}) &= \mathbb{P}[L'_i \geq 1 \mid \Lambda_1 = \theta_1, \dots, \Lambda_{m_S} = \theta_{m_S}] \quad (2. 40) \\ &= 1 - e^{-\lambda_i \sum_{s=1}^{m_S} w_{is} \theta_s / \lambda_{(s)}} . \end{aligned}$$

Let L' denote the random variable representing the number of defaults in the portfolio. We already mentioned that CreditRisk⁺ is a Poisson mixture model. More explicitly, it is assumed that L' is a Poisson variable with random intensity $\Lambda^{(1)} + \dots + \Lambda^{(m_S)}$. Additionally, it is naturally required to obtain the portfolio's defaults as the sum of single obligor defaults, and indeed (2. 39) obviously is consistent with $L' = L'_1 + \dots + L'_m$ when defining the sector's mean intensity by

$$\lambda_{(s)} = \sum_{i=1}^m w_{is} \lambda_i ;$$

see also Formula (4. 30) in [Section 4.3.2](#).

Now, on the portfolio level, the “trick” CreditRisk⁺ uses in order to obtain a nice closed-form distribution of portfolio defaults is *sector analysis*. Given that we know distribution of defaults in every single sector, the portfolio’s default distribution then just turns out to be the *convolution* of the sector distributions due to the independence of the sector variables $\Lambda^{(1)}, \dots, \Lambda^{(m_s)}$. So we only have to find the sector’s default distributions.

When focussing on single sectors, it is a standard result from elementary statistics (see, e.g., [109] Section 8.6.1) that any gamma-mixed Poisson distribution follows a *negative binomial distribution*, . See [Figure 2.7](#) in [Section 2.5.2](#). Therefore every sector has its own individually parametrized negative binomial distribution of sector defaults, such that the portfolio’s default distribution indeed can be obtained as a convolution of negative binomial distributions. As a consequence, the *generating function* of the portfolio loss can be explicitly written in a closed form; see Formula (4. 35) in [Chapter 4](#).

So far we have only discussed the distribution of defaults. The corresponding loss distributions for a single sector are given as the *compound* distribution arising from two independent random effects, where the first random effect is due to the uncertainty regarding the number of defaults (negative binomially distributed) in the sector and the second random effect arises from the uncertainty regarding the exposures affected by the sector defaults; see [Section 4.3.2](#). On the portfolio level the loss distribution again is the convolution of sector loss distributions. The final formula for the generating function of the portfolio loss is presented in (4. 36).

2.4.3 CreditPortfolioView

As mentioned before there are good reasons for keeping the presentation of CPV in a summarizing style. However, for readers interested in more details we refer to the papers by Wilson [127,128], and the technical documentation [85] of CPV in its recent version¹⁴. Both sources and the overview in Crouhy et al. [21], Section 8.10, have been valuable references for writing this paragraph.

¹⁴We are grateful to McKinsey & Company for sending us the technical documentation [85] of CPV as a source for writing this section.

So far we restricted the discussion in this chapter to default modeling. For our summary of CreditPortfolioView (CPV) we will now also include rating migrations as already indicated in the introduction of this chapter.

CPV has its roots in two papers by Wilson [127,128]. Based on these two papers, McKinsey & Company¹⁵ developed CreditPortfolioView during the years since then as a tool for supporting consulting projects in credit risk management. Summarizing one could say that CPV is a ratings-based portfolio model incorporating the dependence of default and migration probabilities on the *economic cycle*. Consequently default probabilities and migration matrices are subject to *random fluctuations*.

Let us start with some general remarks regarding *migration matrices*. Mathematically speaking a migration matrix is a *stochastic*¹⁶ *matrix* in $\mathbb{R}^{n \times n}$, where n depends on the number of rating classes incorporated. For example, the rating agencies (Moody's or S&P) typically publish migration matrices w.r.t. two different dimensions, namely $n = 8$ (standard case) and n substantially larger than 8, reflecting the finer rating scale as shown in [Figure 1.2](#). Migration matrices will be extensively studied later on when discussing the *term structure of default probabilities*; see [Section 6.3.3](#).

The basic observation underlying CPV is that migration probabilities show random fluctuations due to the volatility of the economic cycle. Very much reflecting the terminology in this chapter, CPV calls any migration matrix observed in a particular year a *conditional* migration matrix, because it is sampled conditional on the economic conditions of the considered year. Calculating the average of conditional migration matrices sampled over a series of years will give us an *unconditional* migration matrix reflecting *expected migration paths*. Such average migration matrices can be found in the rating agency reports, or can be calculated from bank-internal data.

Now let us assume that an unconditional migration matrix has been chosen. We denote this matrix by $\overline{M} = (\overline{m}_{ij})$ where i, j range from 1 to 8. Compatible to the notation at the beginning of this chapter we denote rating classes by R_i . Rating class R_1 stands for the best possible credit quality, whereas R_8 is the *default state*, such that $\overline{m}_{i8} =$

¹⁵McKinsey & Company is an international management consulting firm.

¹⁶A matrix (m_{ij}) is called *stochastic* if $\sum_j m_{ij} = 1$ for every row i .

$\mathbb{P}[R_i \rightarrow R_8]$ denotes the probability that obligors with rating R_i at the beginning of a year go into default until the end of that year. In general it is assumed that R_i is more creditworthy than R_j if and only if $i < j$; compare also to [Figure 1.2](#). Because the default state is *absorbing*¹⁷, we additionally have $\bar{m}_{8j} = 0$ for $j = 1, \dots, 7$ and $\bar{m}_{88} = 1$. Note that in this notation \bar{m}_{i8} takes over the role of p_i in previous paragraphs, where p_i denoted the one-year default probability of some customer i . Also recall that default probabilities are typically *rating-driven* so that there is no need to distinguish between two obligors with the same rating when interested in their default probabilities; see also [Section 1.1.1](#).

CPV assumes that there are several *risk segments* differently reacting to the overall economic conditions. For example, typical risk segments refer to industry groups. In our presentation we will not be bothered about the interpretation of risk segments; so, we just assume that there are m_S such segments. Moreover, to keep our presentation free from index-overloading we restrict ourselves to a one-year view. For each segment CPV simulates a conditional migration matrix based on the average migration matrix \bar{M} and a so-called *shift algorithm*. The shift algorithm works in three steps:

1. A segment-specific conditional default probability p_s is simulated for every segment $s = 1, \dots, m_S$. The probability p_s is the same for all rating classes, and we will later explain the simulation method CPV uses for generating those probabilities. Any simulated vector (p_1, \dots, p_{m_S}) can be considered as an aggregated (second-level) *scenario* in a Monte Carlo simulation of CPV. Underlying the generation of such a scenario is the simulation of macroeconomic factors driving (p_1, \dots, p_{m_S}) .
2. A so-called *risk index* r_s representing the state of the economy seen in light of segment s is calculated by means of the ratio

$$r_s = \frac{p_s}{\bar{p}_s} \quad (s = 1, \dots, m_S), \quad (2.41)$$

where \bar{p}_s denotes the unconditional default probability of segment s , incorporating the average default potential of segment s .

¹⁷Absorbing means that the default state is a trap with no escape.

3. Last, a conditional migration matrix $M^{(s)} = (m_{ij}^{(s)})$ for segment s w.r.t. a scenario (p_1, \dots, p_{m_s}) is defined by

$$m_{ij}^{(s)} = \alpha_{ij}(r_s - 1) + \bar{m}_{ij} \quad (s = 1, \dots, m_s). \quad (2. 42)$$

The *shift coefficients* α_{ij} have to be calibrated by the user of CPV, although CPV contains some standard values that can be chosen in case a user does not want to specify the shift factors individually. The shift factors depend on the considered migration path $R_i \rightarrow R_j$, hereby expressing the sensitivity of $\mathbb{P}[R_i \rightarrow R_j]$ w.r.t. a change in the segment's risk index r_s . Because they are intended to reflect rating class behaviour rather than a segment's reaction to macroeconomic conditions, they are uniform w.r.t. different segments. Calculating the row sums of the shifted migration matrices we obtain

$$\sum_{j=1}^8 m_{ij}^{(s)} = (r_s - 1) \sum_{j=1}^8 \alpha_{ij} + \sum_{j=1}^8 \bar{m}_{ij}.$$

Therefore, in order to guarantee that the shifted matrix $M^{(s)}$ is stochastic, CPV assumes $\sum_{j=1}^8 \alpha_{ij} = 0$. If a concrete realization of (2. 42) results in a migration probability $m_{ij}^{(s)}$ shifted to the negative, CPV performs a correction by setting such negative values to zero. In such cases a renormalization of the rows of $M^{(s)}$ is necessary in order to obtain a stochastic matrix.

For some reasons to be explained later the *shift matrix* (α_{ij}) is supposed to satisfy some more conditions, namely

$$\alpha_{ij} \geq 0 \quad \text{for } i < j \quad \text{and} \quad \alpha_{ij} \leq 0 \quad \text{for } i > j. \quad (2. 43)$$

This assumption is compatible to the condition that the shift matrix has row sums equal to zero. Because the upper (lower) triangle matrix of a migrating matrix contains the probabilities for a rating downgrade (upgrade), the conditions on the shift matrix are not as arbitrary as it seems at first glance. Just below we come back to this issue.

Any conditional migration matrix $M^{(s)}$ is relevant to all obligors in segment s . Thinking about Formula (2. 42) and applying the conditions in (2. 43) we see that one can distinguish between three different types of scenarios:

- $r_s < 1$:

In such situations the simulation suggests an *expansion* of the economy, admitting a potential for a lower number of downgrades and a higher number of upgrades, reflecting favourable economic conditions.

- $r_s = 1$:

This is the *average* macroeconomic scenario. Formula (2. 42) shows that in such cases the impact of the shift coefficients vanishes to zero such that the shifted migration probability $m_{ij}^{(s)}$ agrees with the unconditional migration probability \bar{m}_{ij} for all combinations of i and j .

- $r_s > 1$:

This scenario refers to a *recession*. Downgrades are more likely and the potential for upgrades is reduced when compared to average conditions.

Note that because CPV treats different segments differently the concept of segment-specific risk indices allows for a great flexibility. In CPV terminology the process of generating macro-scenarios in a Monte-Carlo simulation is called the *systematic risk model*.

Based on any outcome of the systematic risk model, CPV constructs a conditional loss distribution for the considered portfolio. In a last step all conditional loss distributions are aggregated to an unconditional portfolio loss distribution. The details of how CPV tabulates losses for obtaining scenario-specific and unconditional distributions are rather technical and can be found in the technical documentation [85] of CPV. There one also finds information about other special aspects of the model, for example the implementation of country risk or the method for discounting cash flows to a present value. A nice and CPV-unique feature is the ability to incorporate *stressed tails* in the systematic risk model (to be used in the *CPV Direct* mode, see our discussion on CPV-Macro and CPV-Direct later in this section) in order to study the impact of heavy recessions.

Remaining to be done in this section is a brief description of how CPV manages to simulate the segment-specific conditional default probabilities p_s . Here CPV supports two distinct modes of calibration:

- *CPV Macro*

If CPV is run in the macro mode, default and rating migration

shifts are explained by a *macroeconomic regression model*. The macroeconomic model underlying systematic influences on the economic future of obligors is calibrated by means of time series of empirical data; see the original papers by Wilson [127, 128]. The calibration of CPV Macro is more complicated than the alternative CPV Direct. The difficulties in calibrating CPV Macro are mainly due to the many parameters that have to be estimated; see Formula (2. 44) and (2. 45).

- *CPV Direct*

In this mode of CPV, the segment-specific conditional default probabilities p_s are directly drawn from a *gamma distribution*. In other words, the conditional probability determining a segment's risk index is not implied by some underlying macroeconomic factor model. Working with CPV Direct, the user can avoid all the difficulties some macroeconomic regression model incorporates. The effort in sector calibration is reduced to the calibration of two parameters of some gamma distribution for each risk segment.

Originally CPV contained only the macro approach. CPV Direct was developed later on in order to make the calibration of the model easier.

2.4.3.1 CPV Macro

In CPV Macro, macroeconomic variables drive the distribution of default probabilities and migration matrices for each risk segment. Typical candidates for macroeconomic factors are the *unemployment rate*, the growth rate of the *Gross Domestic Product (GDP)*, *interest* or *currency exchange rates*, and other variables reflecting the macroeconomy of a country. The regression model underlying CPV Macro can be described as follows.

Let us again assume we work with m_S risk segments. Every risk segment s is represented by a *macroeconomic index* $Y_{s,t}$ where t refers to the particular time the index is considered. The index $Y_{s,t}$ itself is represented by a weighted sum of *macroeconomic variables*,

$$Y_{s,t} = w_{s,0} + \sum_{k=1}^K w_{s,k} X_{s,k,t} + \varepsilon_{s,t} , \quad (2. 44)$$

where $X_{s,k,t}$ are macroeconomic variables at time t , relevant to the economic performance of segment s , $(w_{s,k})_{k=0,\dots,K}$ are coefficients that

have to be calibrated w.r.t. segment s , and $\varepsilon_{s,t}$ describes the *residual* random fluctuation of $Y_{s,t}$ not explainable by the fluctuation of the $X_{s,k,t}$'s. For every segment s , a calibration analogous to (2. 44) has to be done. Typically for such regression models the residual variables $\varepsilon_{s,t}$, $s = 1, \dots, m_S$, are assumed to be i.i.d. normally distributed and independent of the variables $X_{s,k,t}$.

The macroeconomic variables $X_{s,k,t}$ are parametrized by an autoregressive model with time lag t_0 , where t_0 has to be specified in the model. More explicitly, it is assumed that the macroeconomic variables can be written as

$$X_{s,k,t} = \theta_{k,0} + \sum_{j=1}^{t_0} \theta_{k,j} X_{s,k,t-j} + \gamma_{s,k,t} . \quad (2. 45)$$

In his original papers Wilson used an *AR(2)-process* ($t_0 = 2$). The conditional default probability for obligors in segment s is specified in terms of the segment-specific macroeconomic index conditioned on its respective realization $Y_{s,t} = y_{s,t}$ and is given by the *logit function*

$$p_{s,t} = \frac{1}{1 + \exp(y_{s,t})} . \quad (2. 46)$$

Once the logit function has been calculated for a scenario, the systematic risk model will calculate the shifted migration matrix as a preparation for the tabulation of the conditional loss distribution.

2.4.3.2 CPV Direct

CPV Direct models the overall macroeconomy (represented by the probabilities p_s) by a multivariate gamma distribution

$$\mathbf{\Gamma} = \left(\Gamma(\gamma_{1,1}, \gamma_{1,2}), \dots, \Gamma(\gamma_{m_S,1}, \gamma_{m_S,2}) \right),$$

where the parameter pairs $(\gamma_{s,1}, \gamma_{s,2})$ have to be calibrated to each segment. A main issue is the calibration of the *correlation* matrix of $\mathbf{\Gamma}$. In general these challenges are much easier to master than calibrating the macroeconomic indices by means of an autoregression as it is suggested by CPV Macro.

The parameters of the gamma distribution of a segment are calibrated by specifying the mean and the volatility of the random variable generating the segment's default probability p_s . The parameters

$(\gamma_{s,1}, \gamma_{s,2})$ are then determined by a moment matching based on (2.38).

Note that the support of the gamma distribution is \mathbb{R}_+ , so that it theoretically can happen to draw a number $p_s > 1$, which obviously can not be interpreted as a probability. This is an unpleasant side effect when drawing a random number that is supposed to be a probability from a gamma distribution. However, to some extent this very much reminds us of the comparably unpleasant possibility of obtaining a multiple default of a single obligor in the CreditRisk⁺ framework; see Section 2.4.2 and Chapter 4. In practice such scenarios are not very likely and will be “thrown away” by the Monte Carlo engine of CPV.

2.4.4 Dynamic Intensity Models

We already mentioned that this section is intended to be not more than a brief “remark” with some references. Dynamic intensity models have been extensively studied by Duffie and Singleton [31,32]. In Duffie and Gârleanu [29] intensity models are applied to the valuation of *collateralized debt obligations*. The theory underlying intensity models has much in common with interest rate term structure models, which are mathematically complex and beyond the scope of this book. For readers interested in the theory we refer to the already mentioned papers by Duffie et al. and also to Jarrow, Lando, and Turnbull [64] and Lando [77].

In the sequel we briefly summarize some basics of one representative of the group of intensity models. First of all, the basic assumption is that every obligor admits a *default time* such that default happens in a time interval $[0, T]$ if and only if the default time of the considered obligor appears to be smaller than the planning horizon T . The default times are driven by an *intensity process*, a so-called *basic affine process*, whose evolution is described by the *stochastic differential equation*

$$d\lambda(t) = \kappa(\theta - \lambda(t))dt + \sigma\sqrt{\lambda(t)}dB(t) + \Delta J(t) , \quad (2.47)$$

where $(B(t))_{t \geq 0}$ is a standard Brownian motion and $\Delta J(t)$ denotes the jump that occurs – if it occurs – at time t . Hereby J is a pure jump process, independent of B , whose jump sizes are independent, positive and exponentially distributed with mean μ and admitting jump times according to an independent Poisson process with mean jump arrival rate l . The parameter set $(\kappa, \theta, \sigma, \mu, l)$ can be adjusted to control the

manner in which default risk changes over time, e.g., one can vary the mean reversion rate κ , the long-run mean $\bar{m} = \theta + l\mu/\kappa$, or the relative contributions to the total variance of $\lambda(t)$ that are attributed to the jump risk and diffusion volatility.

Conditional on a realization $(\lambda(t))_{t \geq 0}$ of the stochastic process solving (2. 47), the default times of obligors are independent *Poisson arrivals* with intensities $\lambda(t)$. From this point of view the dynamic intensity model can be considered as a time-continuous extension of the CreditRisk⁺ framework.

The unconditional survival probability $q(t)$ is given by

$$q(t) = \mathbb{E} \left[e^{-\int_0^t du \lambda(u)} \right].$$

The conditional survival probability for a time interval of length $s \geq 0$ given survival up to time t can be calculated by

$$q(s + t | t) = \mathbb{E}_t \left[e^{-\int_t^{t+s} du \lambda(u)} \right] = e^{\alpha(s) + \beta(s)\lambda(s)},$$

where the explicit solutions for the coefficients $\alpha(s)$ and $\beta(s)$ can be found in the above-mentioned papers by Duffie et al.

Because the sum of independent basic affine processes with common parameters κ , σ , and μ , governing respectively the mean-reversion rate, diffusive volatility, and mean jump size, again yields a basic affine process, one can introduce dependencies between the default times of the counterparties in a considered portfolio. Each obligor's default intensity can thus be represented by means of a one-factor *Markov model*

$$\lambda_i = X_c + X_i ,$$

where X_c and X_1, \dots, X_N are independent basic affine processes with respective parameters $(\kappa, \theta_c, \sigma, \mu, l_c)$ and $(\kappa, \theta_i, \sigma, \mu, l_i)$. The so constructed process λ_i then again is a basic affine process with parameters $(\kappa, \theta, \sigma, \mu, l)$, where $\theta = \theta_c + \theta_i$ and $l = l_c + l_i$. One could interpret X_c as a *state variable* governing the *common aspects of performance*, whereas X_i , seen as a state variable, contributes the *obligor-specific* or *idiosyncratic risk*. Obviously, this can be extended to handle multi-factor models by introducing additional basic affine processes for each of a collection of *sectors*. Each obligor then admits an intensity process

$$\lambda_i = X_c + X_i + X_{s(i)} ,$$

where the *sector factor* $X_{s(i)}$ is common to all obligors in that sector. Here $s(i)$ denotes the sector in which obligor i takes place

A possible simulation algorithm to generate default times τ_1, \dots, τ_n up to some time horizon T with given intensities $\lambda_1, \dots, \lambda_n$ is the *multi-compensator* method. In this method it is assumed that the *compensator* $\Lambda_i(t) = \int_0^t \lambda_i(u)du$ can be simulated for all i and t . Then n independent unit-mean exponentially distributed random variables Z_1, \dots, Z_n are drawn. For each i one has $\tau_i > T$ if $\Lambda_i(T) < Z_i$. In this case the obligor survived the time interval $[0, T]$. Otherwise, the default time of the obligor is given by $\tau_i = \min\{t : \Lambda_i(t) = Z_i\}$ (see the references for more details).

Here we stop our discussion and hope that, although we kept our introduction to the basics of dynamic intensity models rather short, the reader nevertheless got some flavour of what these models are all about.

2.5 One-Factor/Sector Models

In Sections 2.1.2 respectively 2.2.2 we discussed portfolios admitting a uniform default probability respectively intensity and a uniform default correlation for Bernoulli respectively Poisson mixture models. In this paragraph we look in more detail at portfolios with uniform dependency structure, namely *one-factor* respectively *one-sector* models.

2.5.1 The CreditMetricsTM/KMV One-Factor Model

The one-factor model in the context of CreditMetricsTM and KMV is completely described by specializing equations (2. 27) and (2. 32) to the case of only one single factor common to all counterparties, hereby assuming that the asset correlation between obligors is uniform. More explicitly, this means that the composite factors Φ_i of all obligors are equal to one single factor, usually denoted by $Y \sim N(0, 1)$. Moreover, instead of (2. 27) one can write¹⁸

$$r_i = \sqrt{\varrho} Y + \sqrt{1 - \varrho} Z_i \quad (i = 1, \dots, m), \quad (2. 48)$$

¹⁸Note that here one could more generally work with ϱ_i instead of ϱ . Note also that the term $\sqrt{\varrho}$ takes over the role of R_i in Equation (2. 27).

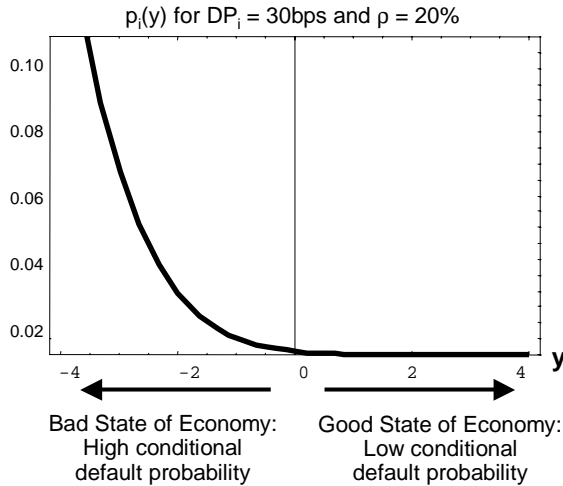


FIGURE 2.3
CreditMetrics™/KMV One-Factor Model: Conditional default probability as a function of the factor realizations $Y = y$.

where $\sqrt{1 - \rho} Z_i$, with $Z_i \sim N(0, 1)$, takes over the role of the residual ε_i and ρ is the uniform asset correlation between the asset value log-returns $r_i \sim N(0, 1)$. In one-factor models, the uniform asset correlation ρ equals the *R-squared* as described in (1. 19), uniform to all obligors. As before, it is assumed that the residuals Z_i constitute an independent family, also independent of the factor Y .


Under the assumption of a single factor and a uniform ρ , Equation (2. 32) turns into

$$p_i(Y) = N \left[\frac{N^{-1}[p_i] - \sqrt{\rho} Y}{\sqrt{1 - \rho}} \right] \quad (i = 1, \dots, m). \quad (2. 49)$$

Figure 2.3 illustrates the dependence of the conditional default probability $p_i(y)$ on realizations y of the single factor Y .

Figure 2.4 shows for three fixed *states of economy* $Y = -3, 0, 3$ the conditional default probability $p_i(y)$ as a function of the average one-year default probability p_i arising in formula (2. 49), which we denoted by DP_i in the introductory Chapter 1. Figures 2.3 and 2.4 also give

$p_i(y)$ in dependence on DP_i for fixed $\rho = 20\%$ and three fixed realizations $Y=y$



Y = - 3

Bad State of Economy:

Conditional default probab.
increases more than linearly
with increasing DP

Y = 0

Typical State of Economy:

Conditional default probab.
increases proportionally
with increasing DP

Y = 3

Good State of Economy:

Conditional default probab.
increases less than linearly
with increasing DP

FIGURE 2.4

CreditMetricsTM/KMV One-Factor Model: Conditional default probability as a function of the average one-year default probability $p_i=DP_i$.

an interpretation of the behaviour of conditional default probabilities in terms of the *economic cycle* captured by the single factor Y .

Before proceeding, we calculate the *joint default probability* (JDP) of two obligors.

2.5.1 Proposition *In a one-factor portfolio model with uniform asset correlation ϱ and loss statistics (L_1, \dots, L_m) with $L_i \sim B(1; p_i(Y))$, where $p_i(Y)$ is defined as in (2. 49), the joint default probability (JDP) of two obligors is given by the bivariate normal integral*

$$JDP_{ij} = \mathbb{P}[L_i = 1, L_j = 1] = N_2[N^{-1}[p_i], N^{-1}[p_j]; \varrho] ,$$

where $N_2[\cdot, \cdot; \varrho]$ denotes the cumulative bivariate normal distribution function with correlation ϱ .

Proof. The joint default probability can be calculated as

$$\mathbb{P}[L_i = 1, L_j = 1] = \mathbb{P}(r_i < N^{-1}[p_i], r_j < N^{-1}[p_j]) .$$

By construction, the correlation between the asset value log-returns $r_i, r_j \sim N(0, 1)$ is ϱ . This proves the proposition. \square

We now want to prove that with increasing portfolio size m (“size” in terms of the number of loans in the portfolio) the portfolio loss distribution converges to a closed-form *limit distribution*. References for the sequel are Finger [36], Gordy [52], Schönbucher [111], Vasicek [124], and Ong [123], Example 9.2. In the following we denote by

$$E_i = \text{EAD}_i \times \text{LGD}_i$$

the exposure that is lost in case obligor i defaults; see [Chapter 1](#) for the meaning of EAD_i and LGD_i . Here we allow for *random* LGDs but deterministic (i.e., fixed) EADs. Moreover, we will not exclude that the LGDs also depend on the state of economy Y in some (not necessarily more detailed specified) way. Such a dependence of the default *and* the recovery rates on the same underlying factor is certainly reasonable, because historic observations show that recovery rates tend to decrease in times where default rates rise up sharply; see, e.g., Altman et al. [2] and Frye [48, 49] for a more detailed discussion of recoveries and their relation to default rates.

Altogether we are looking at a Bernoulli mixture model¹⁹, such that the counterparties are modeled by random variables

$$E_i L_i, \quad L_i \sim B(1; p_i(Y)), \quad Y \sim N(0, 1), \quad (2.50)$$

$$\left((\text{LGD}_i \times L_i) |_{Y=y} \right)_{i=1, \dots, m} \text{ independent,}$$

where we assume that all involved random variables are defined on a common probability space; see Remark 2.5.6. The last condition in (2.50) means that we assume conditional independence of *losses* rather than independence of *default indicators*. For reasons of a shorter and more convenient notation we write in the sequel η_i for LGD_i ,

$$\eta_i = \text{LGD}_i.$$

For a portfolio of m obligors²⁰, the portfolio loss relative to the portfolio's total exposure is given by

$$L = L^{(m)} = \sum_{i=1}^m w_i \eta_i L_i \quad \text{where} \quad w_i = \frac{\text{EAD}_i}{\sum_{j=1}^m \text{EAD}_j}. \quad (2.51)$$

We now want to prove that with increasing number of obligors m some limit behaviour of the portfolio loss $L^{(m)}$ can be established. For this we first of all need some technical assumption, essentially taking care that in the limit the portfolio is free of any dominating single exposures.

2.5.2 Assumption *In the following we consider an infinite number of loans with exposures EAD_i . We assume that the following holds:*

$$\sum_{i=1}^m \text{EAD}_i \uparrow \infty \quad (m \rightarrow \infty),$$

$$\sum_{m=1}^{\infty} \left(\frac{\text{EAD}_m}{\sum_{i=1}^m \text{EAD}_i} \right)^2 < \infty.$$

¹⁹Note that the following notation, although intuitive, is not mathematically rigorous. Later on in the proof of Proposition 2.5.4 we will follow a mathematically more precise notation.

²⁰Here we make the simplifying assumption that the number of loans in the portfolio equals the number of obligors involved. This can be achieved by aggregating different loans of a single obligor into one loan. Usually the DP, EAD, and LGD of such an aggregated loan are exposure-weighted average numbers.

The first condition says that the total exposure of the portfolio strictly increases to infinity with increasing number of obligors. The second condition implies that the exposure weights shrink very rapidly with increasing number of obligors. Altogether this makes sure that the exposure share of each loan in the portfolio tends to zero.

Condition 2.5.2 is by no means a strict assumption. As an example consider the following situation:

2.5.3 Example Assuming $a \leq EAD_i \leq b$ for some $0 < a \leq b$ and all i , we obtain

$$\sum_{i=1}^m EAD_i \geq m a \uparrow \infty \quad (m \rightarrow \infty),$$

$$\sum_{m=1}^{\infty} \left(\frac{EAD_m}{\sum_{i=1}^m EAD_i} \right)^2 \leq \sum_{m=1}^{\infty} \frac{b^2}{m^2 a^2} = \frac{b^2}{a^2} \sum_{m=1}^{\infty} \frac{1}{m^2} < \infty,$$

such that Assumption 2.5.2 is fulfilled in this case.

Now we are in a position to prove the following statement.

2.5.4 Proposition Assumption 2.5.2 is sufficient to guarantee that in the limit the percentage portfolio loss $L^{(m)}$ defined in (2. 51) and the conditional expectation $\mathbb{E}[L^{(m)}|Y]$ are equal almost surely, such that

$$\mathbb{P} \left[\lim_{m \rightarrow \infty} \left(L^{(m)} - \mathbb{E}[L^{(m)}|Y] \right) = 0 \right] = 1.$$

Proof. Fix $y \in \mathbb{R}$. Define the conditional probability measure \mathbb{P}_y by

$$\mathbb{P}_y(\cdot) = \mathbb{P}[\cdot | Y = y].$$

Consider the random variable

$$X_k = EAD_k(\eta_k L_k - \mathbb{E}[\eta_k L_k | Y]).$$

With respect to \mathbb{P}_y , the random sequence $(X_k)_{k \geq 1}$ is independent due to (2. 50) and centered by definition. We now define $\tau_m = \sum_{i=1}^m EAD_i$, such that $(\tau_m)_{m \geq 1}$ is a positive sequence strictly increasing to infinity due to Assumption 2.5.2. If we could prove that

$$\sum_{k=1}^{\infty} \frac{1}{\tau_k^2} \mathbb{E} [X_k^2] < \infty, \tag{2. 52}$$

then a version²¹ of the *strong law of large numbers* (see [7]) would yield

$$\lim_{m \rightarrow \infty} \frac{1}{\tau_m} \sum_{k=1}^m X_k = 0 \quad \mathbb{P}_y - \text{almost surely.} \quad (2. 53)$$

We therefore prove (2. 52) next. From Assumption 2.5.2 we get

$$\sum_{k=1}^{\infty} \frac{1}{\tau_k^2} \mathbb{E} [X_k^2] \leq \sum_{k=1}^{\infty} \frac{4 \times \text{EAD}_k^2}{\tau_k^2} < \infty$$

due to the uniform boundedness of $(\eta_k L_k - \mathbb{E}[\eta_k L_k | Y])$. So we have established (2. 53) for every $y \in \mathbb{R}$. We can now write

$$\mathbb{P} \left[\lim_{m \rightarrow \infty} (L^{(m)} - \mathbb{E}[L^{(m)} | Y]) = 0 \mid Y = y \right] = 1 \quad \text{for every } y \in \mathbb{R}.$$

But then almost sure convergence also holds unconditionally,

$$\begin{aligned} & \mathbb{P} \left[\lim_{m \rightarrow \infty} (L^{(m)} - \mathbb{E}[L^{(m)} | Y]) = 0 \right] = \\ & = \int \mathbb{P} \left[\lim_{m \rightarrow \infty} (L^{(m)} - \mathbb{E}[L^{(m)} | Y]) = 0 \mid Y = y \right] dP_Y(y) = 1. \end{aligned}$$

Therefore the proposition is proved. \square

2.5.5 Corollary *In the case that $(\eta_i L_i)_{i \geq 1}$ are not only conditionally independent but also identically distributed, Proposition 2.5.4 can be reformulated as follows:*

There exists some measurable function $p : \mathbb{R} \rightarrow \mathbb{R}$ such that for $m \rightarrow \infty$ the portfolio loss $L^{(m)}$ converges to $p \circ Y$ almost surely. Moreover, $p \circ Y$ equals $\mathbb{E}[\eta_1 L_1 | Y]$ almost surely.

Proof. Because the conditional expectation $\mathbb{E}[L^{(m)} | Y]$ is by definition $\sigma(Y)$ -measurable, where $\sigma(Y)$ denotes the σ -Algebra generated by Y , there exists some measurable function $p : \mathbb{R} \rightarrow \mathbb{R}$ with $\mathbb{E}[L^{(m)} | Y] = p \circ Y$; see [70], Lemma 1.13. Combined with Proposition 2.5.4 and the assumption that all losses are identically distributed this concludes the proof of Corollary 2.5.5. \square

²¹This version of the LLNs is based on Kronecker's Lemma (see [7]), saying that whenever $(x_k)_{k \geq 1}$ and $(\tau_k)_{k \geq 1}$ are sequences with the latter being positive and strictly increasing to infinity, such that $\sum_{k=1}^{\infty} x_k / \tau_k$ converges, we obtain $\lim_{m \rightarrow \infty} \tau_m^{-1} \sum_{k=1}^m x_k = 0$.

The important conclusion from the convergence results above is that in the limit the randomness of the portfolio loss $L^{(m)}$ solely depends on the randomness of the factor Y . By increasing the number of obligors in the portfolio, the specific risk is completely removed, such that in the limit only *systematic risk* arising from the volatility of the factor Y remains in the portfolio.

2.5.6 Remark The proof of Proposition 2.5.4 does not rely on the particular distribution we use for the factor Y . To make this more precise let us look at a probability space suitable for uniform portfolios. We need a factor Y and residual variables Z_1, Z_2, \dots , which are random variables in \mathbb{R} defined on some not necessarily more specified probability spaces $(\Omega_Y, \mathcal{F}_Y, P_Y)$, $(\Omega_1, \mathcal{F}_1, P_1)$, $(\Omega_2, \mathcal{F}_2, P_2)$, and so on. A suitable probability space for Proposition 2.5.4 is the product space

$$(\Omega, \mathcal{F}, P) = (\Omega_Y, \mathcal{F}_Y, P_Y) \otimes (\Omega_1, \mathcal{F}_1, P_1) \otimes (\Omega_2, \mathcal{F}_2, P_2) \otimes \dots ,$$

because we always assume the variables Y, Z_1, Z_2, \dots to be independent. For every $\omega = (y, z_1, z_2, \dots) \in \Omega$ the loss variables $L_i(\omega)$ are given by *latent variable indicators* evaluated w.r.t. the realization ω ,

$$L_i(\omega) = \mathbf{1}_{\{\sqrt{\varrho} y + \sqrt{1-\varrho} z_i < c_i\}} .$$

It is not difficult to argue that the proof of Proposition 2.5.4 only relies on the conditional independence of the variables Z_i w.r.t. Y and the asymptotics of the portfolio weights according to Assumption 2.5.2. In the case that the factor Y and the residuals Z_i are normally distributed, (Ω, \mathcal{F}, P) turns out to be an *infinite dimensional Gaussian space*, but due to the more generally applicable proof we can use the same convergence argument for distributions other than normal. For example, the t -distribution is a natural candidate to replace the normal distribution; see [Section 2.6.1](#).

Now let us apply our findings to *uniform portfolios* by assuming that $p_i = p$ for all obligors i , such that the assumptions of Proposition 2.5.5 are fulfilled. In the CreditMetricsTM respectively KMV framework, the factor Y and the residual variables Z_1, Z_2, \dots follow a standard normal distribution. For reasons of simplicity we assume constant LGDs ($\eta_i = 100\%$). In this framework, the function p from Proposition 2.5.5 can be explicitly derived by applying Equation (2. 49), and taking into

account that we are in a Bernoulli framework,

$$\mathbb{E}[L^{(m)}|Y] = \sum_{i=1}^m w_i \mathbb{E}[L_i|Y] = N \left[\frac{N^{-1}[p] - \sqrt{\varrho} Y}{\sqrt{1 - \varrho}} \right] =: p(Y),$$

such that Proposition 2.5.4 guarantees that

$$L^{(m)} \xrightarrow{m \rightarrow \infty} p(Y) = N \left[\frac{N^{-1}[p] - \sqrt{\varrho} Y}{\sqrt{1 - \varrho}} \right] \text{ almost surely.} \quad (2.54)$$

So for portfolios with a sufficiently large portfolio size m satisfying Assumption 2.5.2, the percentage quote of defaulted loans for a given state of economy $Y = y$ is approximately equal to the conditional default probability $p(y)$. In the limit we obtain a portfolio loss variable $p(Y)$ describing the fraction of defaulted obligors in an *infinitely fine-grained* credit portfolio.

We now want to derive the cumulative distribution function and the probability density of the limit loss variable $p(Y)$, $Y \sim N(0, 1)$, with $p(\cdot)$ as in (2.54). Denote the portfolio's percentage number of defaults in an infinitely fine-grained portfolio (again assuming constant LGDs of 100%) by L . We then have for every $0 \leq x \leq 1$

$$\begin{aligned} \mathbb{P}[L \leq x] &= \mathbb{P}[p(Y) \leq x] & (2.55) \\ &= \mathbb{P} \left[-Y \leq \frac{1}{\sqrt{\varrho}} \left(N^{-1}[x] \sqrt{1 - \varrho} - N^{-1}[p] \right) \right] \\ &= N \left[\frac{1}{\sqrt{\varrho}} \left(N^{-1}[x] \sqrt{1 - \varrho} - N^{-1}[p] \right) \right]. \end{aligned}$$

In the sequel we will denote this distribution function by

$$F_{p,\varrho}(x) = \mathbb{P}[L \leq x] \quad (x \in [0, 1]). \quad (2.56)$$

The corresponding probability density can be derived by calculating the derivative of $F_{p,\varrho}(x)$ w.r.t. x , which is

$$\begin{aligned} f_{p,\varrho}(x) &= \frac{\partial F_{p,\varrho}(x)}{\partial x} = \sqrt{\frac{1 - \varrho}{\varrho}} \times & (2.57) \\ &\times \exp \left(-\frac{1}{2\varrho} \left((1 - 2\varrho) (N^{-1}[x])^2 - 2\sqrt{1 - \varrho} N^{-1}[x] N^{-1}[p] + (N^{-1}[p])^2 \right) \right) \end{aligned}$$

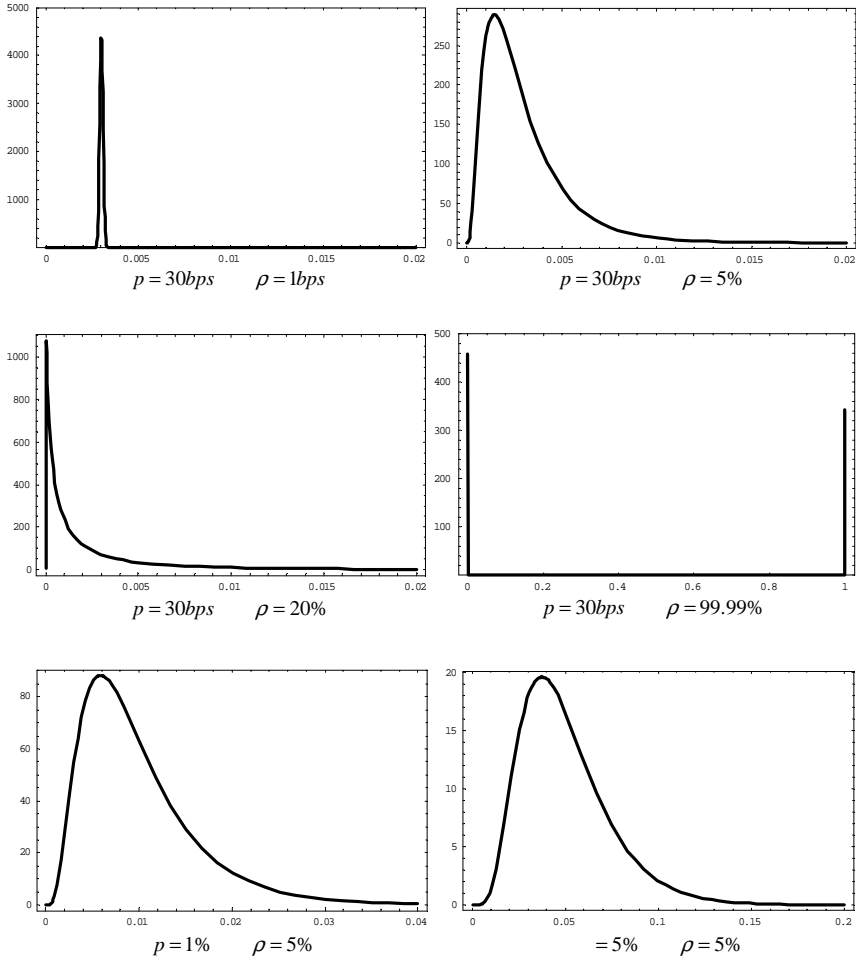


FIGURE 2.5

The probability density $f_{p,\rho}$ for different combinations of p and ρ (note that the x-axes of the plots are differently scaled).

$$= \sqrt{\frac{1-\varrho}{\varrho}} \exp\left(\frac{1}{2} (N^{-1}[x])^2 - \frac{1}{2\varrho} \left(N^{-1}[p] - \sqrt{1-\varrho} N^{-1}[x]\right)^2\right).$$

Figure 2.5 shows the loss densities $f_{p,\varrho}$ for different values of p and ϱ .

It could be guessed from Figure 2.5 that regarding the extreme cases w.r.t. p and ϱ some reasonable limit of $f_{p,\varrho}$ should exist. Indeed, one can easily prove the following statement:

2.5.7 Proposition *The density $f_{p,\varrho}$ admits four extreme cases induced by the extreme values of the parameters p and ϱ , namely*

1. $\varrho = 0$:

This is the correlation-free case with loss variables

$$L_i = \mathbf{1}_{\{r_i = Z_i < N^{-1}[p]\}} \sim B(1; p),$$

taking (2. 48) into account. In this case, the absolute (size- m) portfolio loss $\sum L_i$ follows a binomial distribution, $\sum_{i=1}^m L_i \sim B(m; mp)$, and the percentage portfolio loss L_m converges by arguments analogous to Proposition 2.5.4 (or just by an application of the Law of Large Numbers) to p almost surely. Therefore, $f_{p,0}$ is the density²² of a degenerate distribution²³ concentrated in p . This is illustrated by the first plot in Figure 2.5, where an almost vanishing correlation ($\varrho=1$ bps) yields an $f_{p,\varrho}$, which is almost just a peak in $p=30$ bps.

2. $\varrho = 1$:

In this case one has perfect correlation between all loss variables in the portfolio (see also Section 1.2, where the term “perfect correlation” was mentioned the first time). In this case we can replace the percentage portfolio loss L_m by $L_1 \sim B(1; p)$, which is no longer dependent on m . Therefore, the limit ($m \rightarrow \infty$) percentage portfolio loss L is also Bernoulli $B(1; p)$, such that $\mathbb{P}[L = 1] = p$ and $\mathbb{P}[L = 0] = 1 - p$. The case of (almost) perfect correlation is illustrated in the fourth plot ($p=30$ bps, $\varrho=99.99\%$) of Figure 2.5, clearly showing the shape of a distribution concentrated in only two points, yielding an “all or nothing” loss.

3. $p = 0$:

All obligors survive almost surely, such that $\mathbb{P}[L = 0] = 1$.

²²More precisely, it is a *delta distribution*.

²³More explicitly, we are talking about a *Dirac measure*.

4. $p = 1$:

All obligors default almost surely, such that $\mathbb{P}[L = 1] = 1$.

Proof. A proof is straightforward. \square

For the infinitely fine-grained limit portfolio (encoded by the portfolio's percentage loss variable L) it is very easy to calculate quantiles at any given level of confidence.

2.5.8 Proposition *For any given level of confidence α , the α -quantile $q_\alpha(L)$ of a random variable $L \sim F_{p,\varrho}$ is given by*

$$q_\alpha(L) = p\left(-q_\alpha(Y)\right) = N\left[\frac{N^{-1}[p] + \sqrt{\varrho} q_\alpha(Y)}{\sqrt{1-\varrho}}\right]$$

where $Y \sim N(0, 1)$ and $q_\alpha(Y)$ denotes the α -quantile of the standard normal distribution.

Proof. The function $p(\cdot)$ is strictly decreasing, as illustrated by [Figure 2.3](#). Therefore it follows that

$$\begin{aligned} \mathbb{P}[L \leq p(-q_\alpha(Y))] &= \mathbb{P}[p(Y) \leq p(-q_\alpha(Y))] \\ &= \mathbb{P}[Y \geq -q_\alpha(Y)] = \mathbb{P}[-Y \leq q_\alpha(Y)] , \end{aligned}$$

taking (2. 55) into account. This proves the proposition. \square .

By definition (see [Section 1.2](#)) the *Unexpected Loss* (UL) is the standard deviation of the portfolio loss distribution. In the following proposition the UL of an infinitely fine-grained uniform portfolio is calculated.

2.5.9 Proposition *The first and second moments of a random variable $L \sim F_{p,\varrho}$ are given by*

$$\mathbb{E}[L] = p \quad \text{and} \quad \mathbb{V}[L] = N_2[N^{-1}[p], N^{-1}[p]; \varrho] - p^2 ,$$

where N_2 is defined as in [Proposition 2.5.1](#).

Proof. That the first moment equals p follows just by construction of $F_{p,\varrho}$. Regarding the second moment, we write $\mathbb{V}[L] = \mathbb{E}[L^2] - \mathbb{E}[L]^2$. We already know $\mathbb{E}[L]^2 = p^2$. So it only remains to show that $\mathbb{E}[L^2] = N_2[N^{-1}[p], N^{-1}[p]; \varrho]$. For proving this, we use a typical ‘‘conditioning

trick.” For this purpose, let $X_1, X_2 \sim N(0, 1)$ denote two independent standard normal random variables, independent from the random variable

$$X = \frac{N^{-1}[p] - \sqrt{\varrho} Y}{\sqrt{1 - \varrho}} \sim N(\mu, \sigma^2)$$

with $\mu = \frac{N^{-1}[p]}{\sqrt{1 - \varrho}}, \quad \sigma^2 = \frac{\varrho}{1 - \varrho}.$

We write g_{μ, σ^2} for the density of X . Then, we can write $\mathbb{E}[L^2]$ as

$$\begin{aligned} \mathbb{E}[L^2] &= \mathbb{E}[p(Y)^2] = \mathbb{E}[N(X)^2] \\ &= \int_{\mathbb{R}} \mathbb{P}[X_1 \leq X \mid X = x] \mathbb{P}[X_2 \leq X \mid X = x] dg_{\mu, \sigma^2}(x) \\ &= \int_{\mathbb{R}} \mathbb{P}[X_1 \leq X, X_2 \leq X \mid X = x] dg_{\mu, \sigma^2}(x) \\ &= \mathbb{P}[X_1 - X \leq 0, X_2 - X \leq 0]. \end{aligned}$$

The variables $X_i - X$ are normally distributed with expectation and variance

$$\mathbb{E}[X_i - X] = -\frac{N^{-1}[p]}{\sqrt{1 - \varrho}} \quad \text{and} \quad \mathbb{V}[X_i - X] = 1 + \frac{\varrho}{1 - \varrho}.$$

The correlation between $X_1 - X$ and $X_2 - X$ equals ϱ . Standardizing²⁴ $X_1 - X$ and $X_2 - X$, we conclude $\mathbb{E}[L^2] = N_2[N^{-1}[p], N^{-1}[p]; \varrho]$. \square

The next proposition reports on higher moments of $F_{p, \varrho}$.

2.5.10 Proposition *The higher moments of $L \sim F_{p, \varrho}$ are given by*

$$\mathbb{E}[L^m] = N_m[(N^{-1}[p], \dots, N^{-1}[p]), C_\varrho]$$

where $N_m[\dots]$ denotes the m -dimensional normal distribution function and $C_\varrho \in \mathbb{R}^{m \times m}$ is a matrix with 1 on the diagonal and ϱ off-diagonal.

Proof. The proof relies on the same argument as the proof of Proposition 2.5.9. A generalization to $m \geq 2$ is straightforward. \square

²⁴Shifting and scaling a random variable in order to achieve mean zero and standard deviation one.

TABLE 2.2: Economic Capital EC_α for an infinitely fine-grained portfolio (portfolio loss

$L \sim F_{p,\varrho}$) w.r.t. p and ϱ , for $\alpha = 99.5\%$.

99.50%	p in %							
p in bps	1%	5%	10%	15%	20%	30%	40%	50%
10	0.12%	0.39%	0.72%	1.06%	1.41%	2.14%	2.86%	3.54%
20	0.22%	0.71%	1.28%	1.87%	2.48%	3.77%	5.14%	6.55%
30	0.32%	0.99%	1.78%	2.58%	3.42%	5.20%	7.13%	9.21%
40	0.41%	1.26%	2.24%	3.24%	4.28%	6.50%	8.93%	11.60%
50	0.49%	1.51%	2.67%	3.85%	5.07%	7.69%	10.58%	13.80%
60	0.57%	1.75%	3.08%	4.42%	5.82%	8.80%	12.11%	15.83%
70	0.65%	1.98%	3.47%	4.97%	6.52%	9.85%	13.55%	17.73%
80	0.73%	2.20%	3.84%	5.49%	7.20%	10.84%	14.90%	19.50%
90	0.81%	2.42%	4.20%	5.99%	7.84%	11.79%	16.18%	21.18%
100	0.88%	2.63%	4.55%	6.48%	8.46%	12.69%	17.40%	22.76%
150	1.23%	3.60%	6.15%	8.67%	11.25%	16.71%	22.75%	29.60%
200	1.55%	4.47%	7.57%	10.60%	13.67%	20.11%	27.18%	35.12%
250	1.85%	5.28%	8.86%	12.33%	15.82%	23.08%	30.96%	39.73%
300	2.14%	6.03%	10.05%	13.90%	17.76%	25.72%	34.26%	43.65%
350	2.41%	6.74%	11.16%	15.36%	19.53%	28.08%	37.17%	47.03%
400	2.67%	7.41%	12.19%	16.70%	21.16%	30.23%	39.76%	49.98%
450	2.92%	8.04%	13.16%	17.96%	22.67%	32.19%	42.09%	52.57%
500	3.16%	8.64%	14.07%	19.13%	24.08%	33.99%	44.19%	54.86%
600	3.62%	9.76%	15.76%	21.28%	26.62%	37.17%	47.82%	58.69%
700	4.04%	10.80%	17.29%	23.20%	28.87%	39.91%	50.84%	61.75%
800	4.44%	11.75%	18.68%	24.93%	30.86%	42.28%	53.38%	64.20%

TABLE 2.3: Economic capital EC_α for an infinitely fine-grained portfolio (portfolio loss $L \sim F_{p,\varrho}$) w.r.t. p and ϱ , for $\alpha = 99.98\%$.

99.98%	ρ in %							
p in bps	1%	5%	10%	15%	20%	30%	40%	50%
10	0.20%	0.82%	1.79%	3.01%	4.50%	8.34%	13.49%	20.22%
20	0.36%	1.41%	2.99%	4.91%	7.18%	12.88%	20.27%	29.60%
30	0.51%	1.94%	4.01%	6.47%	9.35%	16.39%	25.27%	36.18%
40	0.65%	2.41%	4.91%	7.84%	11.20%	19.30%	29.29%	41.27%
50	0.78%	2.86%	5.74%	9.07%	12.85%	21.83%	32.68%	45.41%
60	0.90%	3.28%	6.51%	10.19%	14.35%	24.07%	35.62%	48.90%
70	1.03%	3.67%	7.23%	11.24%	15.72%	26.08%	38.21%	51.89%
80	1.14%	4.05%	7.90%	12.22%	17.00%	27.92%	40.52%	54.51%
90	1.26%	4.42%	8.55%	13.14%	18.19%	29.61%	42.61%	56.82%
100	1.37%	4.77%	9.17%	14.01%	19.30%	31.17%	44.51%	58.88%
150	1.90%	6.36%	11.91%	17.81%	24.09%	37.62%	52.05%	66.63%
200	2.38%	7.77%	14.24%	20.95%	27.94%	42.55%	57.45%	71.75%
250	2.83%	9.03%	16.28%	23.65%	31.18%	46.50%	61.57%	75.39%
300	3.25%	10.19%	18.11%	26.02%	33.97%	49.78%	64.81%	78.07%
350	3.64%	11.26%	19.77%	28.13%	36.41%	52.54%	67.43%	80.09%
400	4.02%	12.26%	21.29%	30.03%	38.57%	54.91%	69.58%	81.64%
450	4.38%	13.19%	22.69%	31.75%	40.51%	56.96%	71.36%	82.84%
500	4.73%	14.07%	23.99%	33.33%	42.25%	58.74%	72.85%	83.76%
600	5.38%	15.68%	26.32%	36.10%	45.27%	61.70%	75.15%	85.01%
700	5.98%	17.14%	28.36%	38.48%	47.78%	64.01%	76.78%	85.69%
800	6.54%	18.45%	30.17%	40.53%	49.89%	66.83%	77.92%	85.98%

TABLE 2.4: Unexpected loss UL for an infinitely fine-grained portfolio (portfolio loss $L \sim F_{p,\rho}$) w.r.t. p and ρ .

p in bps	p in %							
	1%	5%	10%	15%	20%	30%	40%	50%
10	0.03%	0.08%	0.14%	0.19%	0.24%	0.37%	0.53%	0.73%
20	0.06%	0.16%	0.25%	0.33%	0.43%	0.63%	0.88%	1.17%
30	0.09%	0.22%	0.35%	0.47%	0.59%	0.86%	1.18%	1.54%
40	0.12%	0.29%	0.45%	0.59%	0.75%	1.07%	1.44%	1.87%
50	0.15%	0.35%	0.54%	0.71%	0.89%	1.27%	1.69%	2.17%
60	0.17%	0.41%	0.63%	0.83%	1.03%	1.46%	1.93%	2.45%
70	0.20%	0.47%	0.72%	0.94%	1.17%	1.64%	2.15%	2.72%
80	0.22%	0.53%	0.80%	1.05%	1.30%	1.81%	2.36%	2.98%
90	0.25%	0.58%	0.88%	1.15%	1.42%	1.98%	2.57%	3.22%
100	0.27%	0.64%	0.96%	1.26%	1.55%	2.14%	2.77%	3.46%
150	0.38%	0.90%	1.34%	1.74%	2.12%	2.88%	3.68%	4.52%
200	0.49%	1.14%	1.70%	2.18%	2.65%	3.56%	4.49%	5.47%
250	0.59%	1.37%	2.03%	2.60%	3.14%	4.18%	5.23%	6.32%
300	0.69%	1.59%	2.34%	2.99%	3.60%	4.76%	5.92%	7.11%
350	0.78%	1.80%	2.65%	3.37%	4.04%	5.31%	6.57%	7.85%
400	0.87%	2.00%	2.94%	3.73%	4.45%	5.83%	7.18%	8.55%
450	0.95%	2.20%	3.21%	4.07%	4.85%	6.33%	7.77%	9.21%
500	1.04%	2.38%	3.48%	4.40%	5.24%	6.81%	8.32%	9.84%
600	1.20%	2.74%	3.99%	5.03%	5.97%	7.71%	9.37%	11.02%
700	1.35%	3.08%	4.48%	5.62%	6.65%	8.55%	10.34%	12.11%
800	1.49%	3.41%	4.93%	6.18%	7.30%	9.34%	11.25%	13.13%




FIGURE 2.6

Dependence of economic capital EC_α on the chosen level of confidence α .

Given a uniform one-year average default probability p and a uniform asset correlation ϱ , [Tables 2.2](#) and [2.3](#) report on the *Economic Capital* (EC) w.r.t. confidence levels of $\alpha = 99,5\%$ and $\alpha = 99,98\%$ for an infinitely fine-grained portfolio (described by the distribution $F_{p,\varrho}$), hereby assuming an LGD of 100% (see [Section 1.2.1](#) for the definition of EC). Analogously, [Table 2.4](#) shows the *Unexpected Loss* for a given pair (p, ϱ) .

[Figure 2.6](#) illustrates the sensitivity of the EC w.r.t. the chosen confidence level. It can be seen that at high levels of confidence (e.g., from 99,9% on) the impact of every basispoint increase of α on the portfolio EC is enormous.

Another common portfolio-dependent quantity is the so-called *capital multiplier* (CM_α); see also [Chapter 5](#) on *capital allocation*. It is defined as the *EC w.r.t. confidence α in units of UL* (i.e., in units of the portfolio standard deviation). In pricing tools the CM is sometimes assumed to be constant for a portfolio, even when adding new deals to it. The contribution of the new deal to the total EC of the enlarged portfolio is then given by a multiple of the CM. In general, the CM heavily depends on the chosen level of confidence underlying the EC definition. Because for given p and ϱ the CM is just the EC scaled by the inverse of the UL, [Figure 2.6](#) additionally illustrates the shape of the curve describing the dependency of the CM from the assumed level of confidence.

For example, for $p=30$ bps (about a BBB-rating) and $\varrho=20\%$ (the Basel II suggestion for the asset correlation of the benchmark risk weights for corporate loans) the (rounded!) CM of a portfolio with loss variable $L \sim F_{p,\varrho}$ is given by $CM_{99\%} \approx 4$, $CM_{99,5\%} \approx 6$, $CM_{99,9\%} \approx 10$, and $CM_{99,98\%} \approx 16$ (in this particular situation we have an UL of 59 bps, as can be read from the [Figure 2.4](#)).

Now, as a last remark in this section we want to refer back to [Section 1.2.2.2](#), where the *analytical approximation* of portfolio loss distributions is outlined. The distribution $L_{p,\varrho}$, eventually combined with some modifications (e.g., random or deterministic LGDs), is extremely well suited for analytical approximation techniques in the context of asset value (or more generally latent variable) models.

2.5.2 The CreditRisk⁺ One-Sector Model

We already discussed CreditRisk⁺ in [Section 2.4.2](#) and will come back to it in [Chapter 4](#). Therefore this paragraph is just a brief “warming-




FIGURE 2.7
Negative binomial distribution with parameters $(\alpha, \beta) = (1, 30)$.

up” for the next paragraph where we compare the uniform portfolio loss distributions of CreditMetrics™ respectively KMV with the corresponding distribution in the CreditRisk⁺ world.

Assuming infinitely many obligors and only one sector, we obtain a situation comparable to the uniform portfolio model of CreditMetrics™ and KMV.

Under these assumptions, the portfolio loss is distributed according to a *negative binomial distribution* $NB(\alpha, \beta)$ due to a gamma-distributed random intensity. The derivation of the negative binomial distribution in the CreditRisk⁺ framework is extensively discussed in [Chapter 4](#). Denoting the portfolio loss by $L' \sim NB(\alpha, \beta)$, the loss distribution is determined by

$$\mathbb{P}[L' = n] = \binom{n + \alpha - 1}{n} \left(1 - \frac{\beta}{1 + \beta}\right)^\alpha \left(\frac{\beta}{1 + \beta}\right)^n, \quad (2. 58)$$

where α and β are called the *sector parameters* of the sector; see Formula (4. 26). The expectation and the variance of L' are given by

$$\mathbb{E}[L'] = \alpha\beta \quad \text{and} \quad \mathbb{V}[L'] = \alpha\beta(1 + \beta), \quad (2. 59)$$

as derived in Formula (4. 27). [Figure 2.7](#) illustrates the shape of the probability mass function of a negative binomial distribution, here

with parameters $\alpha = 1$ and $\beta = 30$. The expected loss in a portfolio admitting such a loss distribution is

$$EL = \mathbb{E}[L'] = 1 \times 30 = 30,$$

and the unexpected loss (volatility of the portfolio loss) is

$$UL = \sqrt{\mathbb{V}[L']} = \sqrt{1 \times 30 \times (1 + 30)} = 30.5 .$$

We are now prepared for the next section.

2.5.3 Comparison of One-Factor and One-Sector Models

Recalling the discussion about general mixture models at the beginning of this chapter one could say that in this section we compare Bernoulli and Poisson mixture models by means of a typical example.

As a representative for the Bernoulli mixture models we choose the random variable $L \sim F_{p,\varrho}$ describing the percentage loss of an infinitely fine-grained portfolio with uniform default probability p and uniform asset correlation ϱ ; see (2. 55). Such portfolios typically arise in analytical approximations in the CreditMetricsTM respectively KMV framework.

The one-sector model of CreditRisk⁺ as described in the previous paragraph will serve as a representative for Poisson mixture models.

A very natural way to calibrate the two models on a common basis is by *moment matching*. One problem we face here is that L takes place in the unit interval and L' generates random integers. We overcome this problem by fixing some large m , say 20,000, such that the tail probability $\mathbb{P}[L' > m]$ is negligibly small, and transforming L' into a variable

$$\tilde{L}' = \frac{L'}{m} .$$

So we take \tilde{L}' as a proxy for the percentage portfolio loss in the one-sector model in CreditRisk⁺. The moment matching procedure is based on the conditions

$$\mathbb{E}[L] = \mathbb{E}[\tilde{L}'] \quad \text{and} \quad \mathbb{V}[L] = \mathbb{V}[\tilde{L}'] .$$

Hereby we always start with some p and ϱ specifying the distribution of L . We then set

$$\mathbb{E}[\tilde{L}'] = p , \quad \mathbb{V}[\tilde{L}'] = N_2[N^{-1}[p], N^{-1}[p]; \varrho] - p^2 ,$$

TABLE 2.5: Comparison of Bernoulli and Poisson mixture models by means of one-factor respectively one-sector models.

portfolio	p	rho	sigma
1	0.01%	10%	0.02%
2	0.01%	20%	0.04%
3	0.01%	30%	0.06%
4	0.30%	10%	0.35%
5	0.30%	20%	0.59%
6	0.30%	30%	0.86%
7	1.00%	10%	0.96%
8	1.00%	20%	1.55%
9	1.00%	30%	2.14%

portfolio	alpha	beta	Q [99.98%] (KMV)	Q [99.98%] (CR+)
1	0.37	5.38	0.31%	0.19%
2	0.08	25.35	0.85%	0.59%
3	0.03	78.17	1.67%	1.42%
4	0.75	80.25	4.30%	3.14%
5	0.26	232.99	9.65%	6.84%
6	0.12	496.04	16.69%	12.20%
7	1.09	184.32	10.17%	8.11%
8	0.42	476.85	20.30%	15.76%
9	0.22	911.68	32.17%	25.69%

applying Proposition 2.5.9. As a last step we solve (2. 59) for α and β . One always has

$$\alpha = \frac{m \times \mathbb{E}[\tilde{L}']^2}{m \times \mathbb{V}[\tilde{L}'] - \mathbb{E}[\tilde{L}']}, \quad \beta = \frac{m \times \mathbb{V}[\tilde{L}'] - \mathbb{E}[\tilde{L}']}{\mathbb{E}[\tilde{L}']}, \quad (2. 60)$$

e.g., for $p=30$ bps, $\varrho=20\%$, and $m=20,000$ we apply 2.5.9 for

$$\mathbb{V}[L] = N_2[N^{-1}[0.003], N^{-1}[0.003]; 0.2] - 0.003^2 = 0.000035095 .$$

The unexpected loss of L therefore turns out to be $UL=59$ bps. Applying Formulas (2. 60), we get

$$\alpha = 0.26 \quad \text{and} \quad \beta = 232.99 ,$$

so that the distribution of \tilde{L}' is finally determined.

In Table 2.5 high-confidence quantiles of one-factor respectively one-sector models with different parameter settings are compared. It turns out that the Bernoulli mixture model always yields fatter tails than

the Poisson mixture model, hereby confirming our theoretical results from Section 2.3. A more detailed comparison of the KMV-Model and CreditRisk⁺ can be found in [12].

2.6 Loss Distributions by Means of Copula Functions

Copula functions have been used as a statistical tool for constructing multivariate distributions long before they were re-discovered as a valuable technique in risk management. Currently, the literature on the application of copulas to credit risk is growing every month, so that tracking every single paper on this issue starts being difficult if not impossible. A small and by no means exhaustive selection of papers providing the reader with a good introduction as well as with a valuable source of ideas how to apply the copula concept to standard problems in credit risk is Li [78,79], Frey and McNeil [45], Frey, McNeil, and Nyfeler [47], Frees and Valdez [44], and Wang [125]. However, the basic idea of copulas is so simple that it can be easily introduced:

2.6.1 Definition *A copula (function) is a multivariate distribution (function) such that its marginal distributions are standard uniform. A common notation for copulas we will adopt is*

$$C(u_1, \dots, u_m) : [0, 1]^m \rightarrow [0, 1]$$

if considered in R^m (e.g., m obligors, m assets, m latent variables, etc.).

The most commonly applied copula function (e.g., in CreditMetricsTM and the KMV-Model) is the *normal copula*, defined by

$$C(u_1, \dots, u_m) = N_m[N^{-1}[u_1], \dots, N^{-1}[u_m]; \Gamma] , \quad (2. 61)$$

with $N_m(\cdot; \Gamma)$ as in Section 2.4.1. In this section we also elaborate that CreditMetricsTM and the KMV-Model implicitly incorporate copula functions based on the multivariate Gaussian distribution of asset value processes. For example, Proposition 2.5.9 says that the *bivariate normal copula* determines the second moment of the loss distribution of an infinitely fine-grained portfolio. So we implicitly already met copulas in previous paragraphs.

However, in the next section we are going to show how to use copulas in order to construct portfolio loss variables admitting a stronger *tail dependency* than induced by the normal copulas.

But before continuing, we want to quote a Theorem by Sklar [113, 114], saying that copulas are a universal tool for studying multivariate distributions.

2.6.2 Theorem (Sklar [113]) *Let \mathbf{F} be a multivariate m -dimensional distribution function with marginals F_1, \dots, F_m . Then there exists a copula C such that*

$$\mathbf{F}(x_1, \dots, x_m) = C\left(F_1(x_1), \dots, F_m(x_m)\right) \quad (x_1, \dots, x_m \in \mathbb{R}).$$

Moreover, if the marginal distributions F_1, \dots, F_m are continuous, then C is unique.

Proof. For a comprehensive proof see Sklar [113], or, alternatively, the textbook [99] by Nelsen. However, the basic idea (which is already the heart of the proof) of deriving a copula from a given multivariate distribution \mathbf{F} with marginals F_1, \dots, F_m is by imitating what we previously have seen in case of the normal copulas, namely

$$C(u_1, \dots, u_m) = \mathbf{F}\left(F_1^{-1}(u_1), \dots, F_m^{-1}(u_m)\right). \quad (2.62)$$

Now one only has to confirm that C defined by (2.62) does the job. \square

The converse of Proposition 2.6.2 is also true:

2.6.3 Proposition *For any copula C and (marginal) distribution functions F_1, \dots, F_m , the function*

$$\mathbf{F}(x_1, \dots, x_m) = C\left(F_1(x_1), \dots, F_m(x_m)\right) \quad (x_1, \dots, x_m \in \mathbb{R})$$

defines a multivariate distribution function with marginals F_1, \dots, F_m .

Proof. The proof is straightforward. One just has to apply the defining properties of copulas. \square

Summarizing Theorem 2.6.2 and Proposition 2.6.3, one can say that every multivariate distribution with continuous marginals admits a unique *copula representation*. Moreover, copulas and distribution functions are the building blocks to derive new multivariate distributions with prescribed correlation structure and marginal distributions. This immediately brings us to the next section.

2.6.1 Copulas: Variations of a Scheme

Here, we are mainly interested in giving some examples of how the copula approach can be used for constructing loss distributions with fatter tails than it would be for normally distributed asset value log-returns. In this book we restrict ourselves to normal and t -copulas, because they are most common in the credit risk context. For other copulas we refer to Nelsen [99].

For our example we look at a Bernoulli mixture model but replace the multivariate normal asset value log-return vector as used in the models of CreditMetricsTM and KMV by a multivariate t -distributed log-return vector. For the convenience of the reader we first recall some basic test distributions from statistics (see, e.g., [106]):

The Chi-Square Distribution:

The χ^2 -distribution can be constructed as follows: Start with an i.i.d. sample $X_1, \dots, X_n \sim N(0, 1)$. Then, $X_1^2 + \dots + X_n^2$ is said to be χ^2 -distributed with n degrees of freedom. The first and second moments of a random variable $X \sim \chi^2(n)$ are

$$\mathbb{E}[X] = n \quad \text{and} \quad \mathbb{V}[X] = 2n .$$

In some sense the χ^2 distribution is a “derivate” of the gamma-distribution (see 2.4.2), because the $\chi^2(n)$ -distribution equals the gamma-distribution with parameters $\alpha = n/2$ and $\beta = 2$. Therefore we already know the shape of χ^2 -densities from Figure 2.2.

The (Student’s) t -distribution:

The building blocks of the t -distribution are a standard normal variable $Y \sim N(0, 1)$ and a χ^2 -distributed variable $X \sim \chi^2(n)$, such that Y and X are independent. Then the variable Z defined by $Z = Y/\sqrt{X/n}$ is said to be t -distributed with n degrees of freedom. The density of Z is given by

$$t_n(x) = \frac{\Gamma((n+1)/2)}{\sqrt{\pi n} \Gamma(n/2)} \left(1 + \frac{x^2}{n}\right)^{-(n+1)/2} \quad (x \in \mathbb{R}).$$

The first and second moments of a variable $Z \sim t(n)$ are given by

$$\mathbb{E}[Z] = 0 \quad (n \geq 2) \quad \text{and} \quad \mathbb{V}[Z] = \frac{n}{n-2} \quad (n \geq 3).$$




FIGURE 2.8
 $t(3)$ -density versus $N(0, 1)$ -density.

For large n , the t -distribution is close to the normal distribution. More precisely, if F_n denotes the distribution function of a random variable $Z_n \sim t(n)$, then one can show that F_n converges in distribution to the distribution function of a standard normal random variable $Z \sim N(0, 1)$; see [106].

This convergence property is a nice result, because it enables us to start in the following modification of the CreditMetrics™/KMV model close to the normal case by looking at a large n . By systematically decreasing the degrees of freedom we can transform the model step-by-step towards a model with fatter and fatter tails.

In general the t -distribution has more mass in the tails than a normal distribution. Figure 2.8 illustrates this by comparing a standard normal density with the density of a t -distribution with 3 degrees of freedom.

The multivariate t -distribution:

Given a multivariate Gaussian vector $\mathbf{Y} = (Y_1, \dots, Y_m) \sim N(0, \Gamma)$ with correlation matrix Γ , the scaled vector $\Theta\mathbf{Y}$ is said to be multivariate t -distributed with n degrees of freedom if $\Theta = \sqrt{n/X}$ with $X \sim \chi^2(n)$ and Θ is independent of \mathbf{Y} . We denote the distribution of such a variable $\Theta\mathbf{Y}$ by $t(n, \Gamma)$. The matrix Γ is explicitly addressed as the second parameter, because $\Theta\mathbf{Y}$ inherits the correlation structure

from \mathbf{Y} :

$$\text{Corr}[\Theta Y_i, \Theta Y_j] = \text{Corr}[Y_i, Y_j] .$$

This can be easily verified by a conditioning argument (w.r.t. Θ).

Now let $n \geq 3$ and F_n be a t -distribution function with n degrees of freedom. Denote the inverse or quantile function of F_n by F_n^{-1} . For the multivariate t -distribution function with n degrees of freedom and correlation matrix Γ we write $F_{n,\Gamma} \sim t(n, \Gamma)$. Then we can define a t -copula function $C_{n,\Gamma}$ by applying formula (2. 62),

$$C_{n,\Gamma}(u_1, \dots, u_m) = F_{n,\Gamma}\left(F_n^{-1}(u_1), \dots, F_n^{-1}(u_m)\right) \quad (2. 63)$$

$$\text{where } u_1, \dots, u_m \in [0, 1].$$

The copula $C_{n,\Gamma}$ incorporates a *multivariate t-dependency* that we can now combine with any marginal distributions we like. For example, a multivariate distribution function with t -dependency and Gaussian marginals can be defined (for $x_1, \dots, x_m \in \mathbb{R}$) by

$$\Phi_{n,\Gamma}(x_1, \dots, x_m) = C_{n,\Gamma}(N[x_1], \dots, N[x_m]) \quad (2. 64)$$

where $N[\cdot]$ denotes the standard normal distribution function. That indeed $\Phi_{n,\Gamma}$ defines a multivariate distribution function with standard normal marginals is a direct consequence of Proposition 2.6.3. Replacing a normal by a t -dependency will – in accordance with the fact that t -tails are fatter than Gaussian tails – significantly shift mass into the tails of the loss distribution arising from a corresponding asset value model. The fatness of tails strongly depends on the chosen degrees of freedom, so that the calibration of an appropriate n in Formula (2. 64) is an essential challenge when dealing with t -copulas. Although there is much literature about the calibration of non-normal distributions to financial time series in general (see, e.g., Eberlein [33]), so far we do not know about an established standard calibration methodology for fitting t -copulas to a credit portfolio. Here we believe that some further research is necessary.

The impact of different dependency structures can be best illustrated by means of a scatterplot. In [Figure 2.9](#) we look at four different variations:

- **Bivariate Gaussian copula with normal marginals:**

We randomly generated points (X_1, X_2) with

$$X_i = \sqrt{\varrho} Y + \sqrt{1 - \varrho} Z_i \quad (i = 1, 2)$$

with $Y, Z_1, Z_2 \sim N(0, 1)$ i.i.d. and $\varrho = 40\%$.

- **Bivariate t-copula with t-distributed marginals:**

Here we plotted randomly generated pairs (X_1, X_2) with

$$X_i = \sqrt{3}(\sqrt{\varrho} Y + \sqrt{1 - \varrho} Z_i) / \sqrt{W} \quad (i = 1, 2)$$

with $Y, Z_1, Z_2 \sim N(0, 1)$ i.i.d., $W \sim \chi^2(3)$, and $\varrho = 40\%$.

- **Bivariate t-copula with normal marginal distributions:**

The points (X_1, X_2) are generated according to

$$X_i = N^{-1}[F_3(\sqrt{3}(\sqrt{\varrho} Y + \sqrt{1 - \varrho} Z_i) / \sqrt{W})] \quad (i = 1, 2)$$

with $Y, Z_1, Z_2 \sim N(0, 1)$ i.i.d., $W \sim \chi^2(3)$, W independent of Y, Z_1, Z_2 , $\varrho = 40\%$, and F_3 denoting the t-distribution function with 3 degrees of freedom. Generalizing for m instead of 2 dimensions we obtain a multivariate distribution function F with

$$\begin{aligned} F(x_1, \dots, x_m) &= \mathbb{P}[X_1 \leq x_1, \dots, X_m \leq x_m] \\ &= P [T_1 \leq F_3^{-1}(N[x_1]), \dots, T_m \leq F_3^{-1}(N[x_m])] \end{aligned}$$

with $(T_1, \dots, T_m) \sim t(3, \Gamma_\varrho)$ and Γ_ϱ denoting the ϱ -uniform correlation matrix in $\mathbb{R}^{m \times m}$. Therefore, we finally see that

$$F(x_1, \dots, x_m) = C_{3, \Gamma_\varrho}(N[x_1], \dots, N[x_m]) = F_{3, \Gamma_\varrho}(x_1, \dots, x_m).$$

This shows that indeed we simulated copula (2. 64) for $n = 3$.

- **Independence copula with normal marginal distributions:**

We randomly generated points (X_1, X_2) with

$$X_1, X_2 \sim N(0, 1) \quad \text{i.i.d.}$$

The *independence copula* is defined by $C(u_1, \dots, u_m) = u_1 \cdots u_m$.

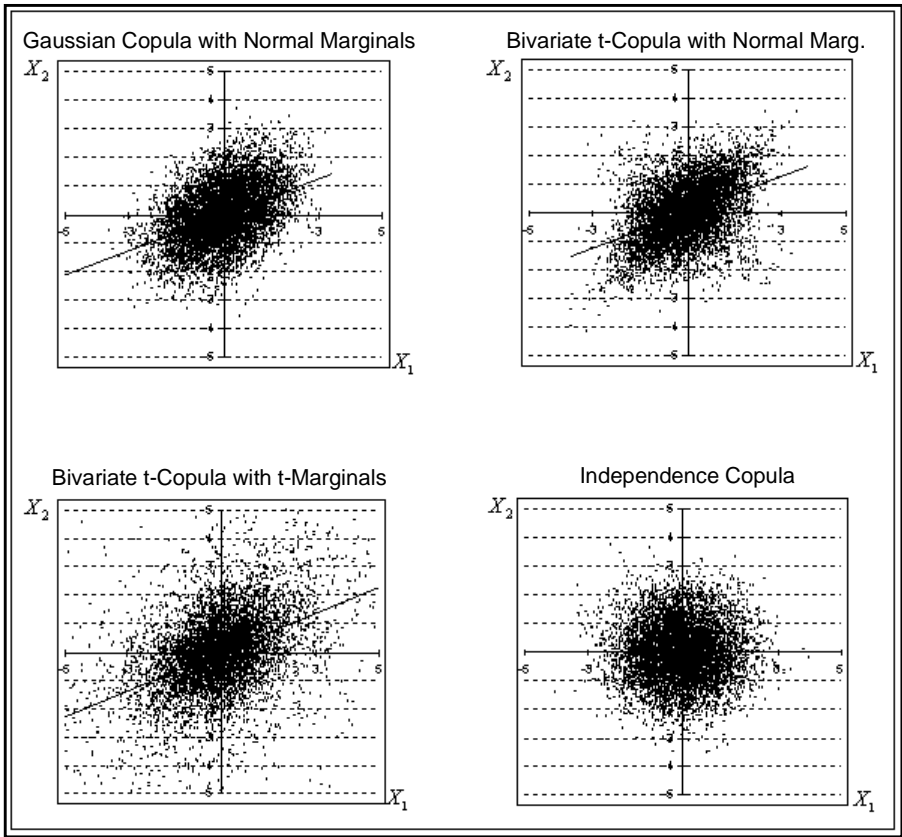


FIGURE 2.9
Normal versus t-dependency with same linear correlation.

TABLE 2.6: Uniform portfolio calculations with t-copulas w.r.t. default probabilities of 50, 80, and 150 basispoints, correlations of 5% and 20%, and degrees of freedom of 10,000, 40, and 10. Quantiles are calculated w.r.t. a confidence of 99%.

Gaussian Copula (not simul.)		
not simulated		
Mean	5%	20%
0.5%	0.5000%	0.5000%
0.8%	0.8000%	0.8000%
1.5%	1.5000%	1.5000%
Quantile	5%	20%
0.5%	1.7470%	4.3017%
0.8%	2.6323%	6.2997%
1.5%	4.5250%	10.3283%
Std.Dev.	5%	20%
0.5%	0.3512%	0.8926%
0.8%	0.5267%	1.2966%
1.5%	0.8976%	2.1205%

T-Copula with df = 10,000		
simulated, 100,000 scenarios		
Mean	5%	20%
0.5%	0.5002%	0.4983%
0.8%	0.8028%	0.8037%
1.5%	1.5034%	1.4944%
Quantile	5%	20%
0.5%	1.7605%	4.3060%
0.8%	2.6351%	6.3342%
1.5%	4.5287%	10.1864%
Std.Dev.	5%	20%
0.5%	0.3522%	0.8946%
0.8%	0.5283%	1.3045%
1.5%	0.8964%	2.0988%

T-Copula with df = 40		
simulated, 100,000 scenarios		
Mean	5%	20%
0.5%	0.4959%	0.4992%
0.8%	0.8006%	0.8009%
1.5%	1.5030%	1.4970%
Quantile	5%	20%
0.5%	2.9674%	5.3814%
0.8%	4.2611%	7.6405%
1.5%	6.6636%	11.9095%
Std.Dev.	5%	20%
0.5%	0.6145%	1.1201%
0.8%	0.8802%	1.5726%
1.5%	1.3723%	2.4653%

T-Copula with df = 10		
simulated, 100,000 scenarios		
Mean	5%	20%
0.5%	0.4990%	0.4973%
0.8%	0.7999%	0.8051%
1.5%	1.5023%	1.5003%
Quantile	5%	20%
0.5%	6.0377%	7.9295%
0.8%	8.0921%	11.2434%
1.5%	11.7042%	16.5620%
Std.Dev.	5%	20%
0.5%	1.2535%	1.7135%
0.8%	1.6574%	2.3104%
1.5%	2.3889%	3.3475%

Figure 2.9 clearly illustrates the impact of different copula variations. For example, the t-copula with normal marginals more or less keeps the Gaussian “tightness” of the point cloud but tends to have greater tail dependency.

We now want to investigate the impact of a change from a normal to a t -copula on the loss distribution of a Bernoulli mixture model based on a uniform asset value model. For this purpose we choose a default probability p , an asset correlation ϱ , and fix n degrees of freedom. Our starting point is Formula (2. 48),

$$r_i = \sqrt{\varrho} Y + \sqrt{1 - \varrho} Z_i \quad (i = 1, \dots, m).$$

By scaling this equation with $\sqrt{n/W}$, $W \sim \chi^2(n)$, W independent of Y, Z_1, \dots, Z_m , we transform the normal copula into a t -copula yielding t -distributed asset value log-returns $(\tilde{r}_1, \dots, \tilde{r}_m) \sim t(n, \Gamma_\varrho)$,

$$\tilde{r}_i = \sqrt{n/W} r_i = \sqrt{n/W} \sqrt{\varrho} Y + \sqrt{n/W} \sqrt{1 - \varrho} Z_i \sim t(n)$$

for $i = 1, \dots, m$. Again denoting the t -distribution function of $t(n)$ by F_n we can write the default point of the model as $F_n^{-1}(p)$. The Bernoulli loss variables are given by $L_i = \mathbf{1}_{\{\tilde{r}_i \leq F_n^{-1}(p)\}}$. The uniform default probability conditional on $Y = y$ and $W = w$ has the following representation:

$$\begin{aligned} p(y, w) &= \mathbb{P}[L_i = 1 \mid Y = y, W = w] \\ &= \mathbb{P}[\tilde{r}_i \leq F_n^{-1}(p) \mid Y = y, W = w] \\ &= P \left[\sqrt{n/W} \sqrt{\varrho} Y + \sqrt{n/W} \sqrt{1 - \varrho} Z_i \leq F_n^{-1}(p) \mid Y = y, W = w \right] \\ &= N \left[\frac{\sqrt{W/n} F_n^{-1}(p) - \sqrt{\varrho} Y}{\sqrt{1 - \varrho}} \mid Y = y, W = w \right] \\ &= N \left[\frac{\sqrt{w/n} F_n^{-1}(p) - \sqrt{\varrho} y}{\sqrt{1 - \varrho}} \right]. \end{aligned}$$

Analogous to the conclusion in (2. 55), not only the single obligor’s conditional default probability but in the limit also the portfolio’s percentage loss is described by $p(y, w)$ given $Y = y$ and $W = w$. We therefore can simulate the portfolio loss in a t -copula model by looking at the distribution of, say, 100,000 samples

$$N \left[\frac{F_n^{-1}[p] \sqrt{W_i/n} - \sqrt{\varrho} Y_i}{\sqrt{1 - \varrho}} \right], \tag{2. 65}$$

$$W_1, \dots, W_{100,000} \sim \chi^2(n), \quad Y_1, \dots, Y_{100,000} \sim N(0, 1) \quad \text{independent.}$$

We have done this exercise for different p 's, ϱ 's, and n 's, and the result is shown in [Table 2.6](#). In the table one can see that for 10,000 degrees of freedom the difference of the portfolio statistics compared to a normal copula is very small and just due to stochastic fluctuations in the simulation. But with decreasing n the portfolio statistics significantly changes. For example, there is a multiplicative difference of almost a factor of 2 between the 99%-quantiles w.r.t. $(p, \varrho) = (0.8\%, 5\%)$ and degrees of freedom of 40 and 10. If we would calculate the quantiles in [Table 2.6](#) w.r.t. higher levels of confidence, the differences would be even higher. Therefore one can easily increase the potential for extreme losses in a uniform t-copula portfolio model by just decreasing the degrees of freedom of the underlying multivariate t-distribution. Unfortunately, a decision of how fat the tails really should be is never easy and sometimes purely subjective. Maybe this is the reason why people very often rely on asset value models based on the Gaussian copula. Gaussian distributions are uniquely determined by their expectation vector and their covariance matrix, such that more complicated calibrations are not necessary. Moreover, as we already indicated, often even the estimation of linear correlations is a great challenge and far from being obvious. We believe that more research combined with empirical evidence is necessary before other than normal copulas will become "best practice" in credit risk management.

Our last point in this section is the following proposition.

2.6.4 Proposition *Given a Bernoulli loss statistics (L_1, \dots, L_m) based on an asset value respectively (more general) latent variables model in the form $L_i = \mathbf{1}_{\{r_i \leq c_i\}}$, the gross loss distribution of (L_1, \dots, L_m) , defined as the distribution of the variable $L = \sum L_i$, is uniquely determined by the set of one-year default probabilities $p_i = \mathbb{P}[r_i \leq c_i]$ and the respective copula function C of (r_1, \dots, r_m) .*

Proof. The distribution of gross losses arising from the loss statistics (L_1, \dots, L_m) is determined by the *joint default probabilities*

$$\begin{aligned} \mathbb{P}[L_{i_1} = 1, \dots, L_{i_k} = 1] &= \mathbb{P}[r_{i_1} \leq c_{i_1}, \dots, r_{i_k} \leq c_{i_k}] = \\ &= C_{i_1, \dots, i_k}(p_{i_1}, \dots, p_{i_k}), \quad \text{with } \{i_1, \dots, i_k\} \subseteq \{1, \dots, m\}, \end{aligned}$$

where C_{i_1, \dots, i_k} denotes the respective k -dimensional marginal distribution of the copula C of (r_1, \dots, r_m) . \square

Given the standard case of a one-year time horizon asset value model, Proposition 2.6.4 says that besides the one-year default probabilities the used copula function completely determines the portfolio loss distribution. In a Gaussian world, the asset correlation as a second parameter in uniform asset value models is the main driver of fat tails. For people allowing for other than normal copulas, changing the copula may have even more impact than just increasing the asset correlation in a Gaussian model.

2.7 Working Example: Estimation of Asset Correlations

We conclude this chapter with a working example regarding the estimation of asset correlations from historic default frequencies. We already saw in Table 1.2 rating agency data from Moody's reporting on historic default frequencies of corporate bond defaults. In the same report [95], Exhibit 39, we also find Table 2.7, showing one-year default rates by year and letter rating from 1970-2000.

What one can clearly see is that observed default frequencies are quite volatile, and a natural interpretation of such volatility is the existence of an *economic cycle*. Although it is the most simple approach, the uniform portfolio model as introduced in Section 2.5.1 already provides us with a useful *parametric framework* in order to estimate the *systematic risk* inherent in Moody's corporate bond portfolio. As reference for the sequel we mention [11], where several approaches for estimating asset correlations are elaborated.

As already indicated above, we use the uniform portfolio model of CreditMetricsTM and KMV as a parametric framework. Table 2.7 includes $R_1=Aaa$, $R_2=Aa$, ..., and $R_6=B$, altogether six rating grades. For every rating class R_i we can calculate the mean \bar{p}_i and the corresponding volatility from the historic default frequencies of class R_i over the years from 1970 to 2000. The result is shown in Tables 2.8 and 2.9 in the mean and standard deviation column.

With Table 2.7 we have the same problem we already faced in Section 1.1.1: There is no default history for upper investment grade bonds. We therefore again fit the historic data by a *linear regression on logarithmic scale*. Hereby we distinguish two regression methods:

TABLE 2.7: Moody's Historic Corporate Bond Default Frequencies from 1970 to 2000.

Rating	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	
Aaa	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Baa	0.27%	0.00%	0.00%	0.45%	0.00%	0.00%	0.00%	0.28%	0.00%	0.00%	0.00%
Ba	4.12%	0.42%	0.00%	0.00%	0.00%	1.02%	1.01%	0.52%	1.08%	0.49%	
B	22.78%	3.85%	7.14%	3.77%	10.00%	5.97%	0.00%	3.28%	5.41%	0.00%	
	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	
Aaa	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.61%	
A	0.00%	0.00%	0.26%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Baa	0.00%	0.00%	0.31%	0.00%	0.36%	0.00%	1.33%	0.00%	0.00%	0.60%	
Ba	0.00%	0.00%	2.72%	0.91%	0.83%	1.75%	2.04%	2.71%	1.24%	2.98%	
B	4.94%	4.49%	2.41%	6.31%	6.72%	8.22%	11.80%	6.25%	6.04%	9.21%	
	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000
Aaa	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
A	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Baa	0.00%	0.28%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.12%	0.10%	0.38%
Ba	3.34%	5.30%	0.30%	0.55%	0.24%	0.67%	0.00%	0.19%	0.61%	1.14%	1.05%
B	16.16%	14.71%	9.03%	5.79%	3.82%	4.80%	1.44%	2.11%	4.26%	5.88%	5.14%

TABLE 2.8: Calibration Results due to Regression I.

Rating	Mean	Stand.Dev.	μ	σ	ρ
Aaa	0.000%	0.000%	0.0001%	0.0023%	34%
Aa	0.020%	0.110%	0.0012%	0.0110%	28%
A	0.008%	0.047%	0.0113%	0.0514%	24%
Baa	0.145%	0.277%	0.1027%	0.2406%	19%
Ba	1.201%	1.330%	0.9348%	1.1270%	14%
B	6.507%	4.762%	8.5040%	5.2788%	10%
Mean	1.31%	1.09%	1.59%	1.12%	22%







TABLE 2.9: Calibration Results due to Regression II.

Rating	Mean	Stand.Dev.	μ	σ	ρ
Aaa	0.000%	0.000%	0.0001%	0.0023%	34%
Aa	0.020%	0.110%	0.0012%	0.0110%	28%
A	0.008%	0.047%	0.0113%	0.0514%	24%
Baa	0.145%	0.277%	0.1027%	0.2406%	19%
Ba	1.201%	1.330%	0.9348%	1.1270%	14%
B	6.507%	4.762%	8.5040%	5.2788%	10%
Mean	1.31%	1.09%	1.59%	1.12%	22%



- *Regression I:*

Here we just set R_1 =Aaa to “not observed” and fit the mean default frequencies $\bar{p}_2, \dots, \bar{p}_6$ by an exponential function yielding fitted default probabilities μ_1, \dots, μ_6 for all rating classes (class Aaa is extrapolated). After that we repeat the same procedure with the volatilities of the default frequency time series of rating classes R_2, \dots, R_6 , this time yielding volatilities $\sigma_1, \dots, \sigma_6$ (class Aaa again extrapolated). The results are shown in [Table 2.8](#).

- *Regression II:*

Regression method II is motivated by the observation that class Aa possibly constitutes an outlier, due to the spike arising from just one observed default frequency in the year 1989. So here we decide to exclude not only Aaa but also Aa from the regression. At the end the default probabilities for Aaa and Aa are extrapolated. [Table 2.9](#) shows the result.

We could continue in this way for rating class A, because in this class we also have only one observation different from zero, namely in year 1982. However, our example is purely illustrative, such that two different regression approaches should be enough to demonstrate the effect. The reason for presenting two approaches is that it shows very clearly that subjective opinions very often play a crucial role in bank-internal calibrations. In fact, there are various ways in which a regression could have been done in order to obtain reasonable default probabilities for every rating class. So people have to make a decision as to which method best reflects their “expert opinion” and their “analytical honesty.” The ρ -columns in [Tables 2.8](#) and [2.9](#) contain estimated average asset correlations for the considered rating classes, and one can see that the different regression approaches are reflected by differences in estimated asset correlations. For example, people relying on *Regression I* would believe in an overall average asset correlation of 25%, whereas people relying on *Regression II* would believe that the overall average asset correlation in Moody’s corporate bond universe is at the lower level of 22%.

Now, it remains to explain how we came up with the asset correlation columns in [Tables 2.8](#) and [2.9](#). For this purpose let us fix a rating class, such that we can drop the index i referring to rating class i . For the chosen rating class, we know that in year j some default frequency

p_j has been observed. The time series p_1, \dots, p_{31} , addressing the historically observed default frequencies for the chosen rating class in the years 1970 up to 2000, is given by the respective row in Table 2.7. In the parametric framework of the CreditMetricsTM/KMV uniform portfolio model, it is assumed that for every year j some realization y_j of a global factor Y drives the realized conditional default probability observed in year j . According to Equation (2. 49) we can write

$$p_j = p(y_j) = N \left[\frac{N^{-1}[p] - \sqrt{\varrho} y_j}{\sqrt{1 - \varrho_i}} \right] \quad (i = 1, \dots, m)$$

where p denotes the “true” default probability of the chosen rating class, and ϱ means the unknown asset correlation of the considered rating class, which will be estimated in the following. The parameter p we do not know exactly, but after a moment’s reflection it will be clear that the observed historic mean default frequency \bar{p} provides us with a good proxy of the “true” mean default rate. Just note that if Y_1, \dots, Y_n are i.i.d.²⁵ copies of the factor Y , then the law of large numbers guarantees that

$$\frac{1}{n} \sum_{j=1}^n p(Y_j) \xrightarrow{n \rightarrow \infty} \mathbb{E}[p(Y)] = p \quad \text{a.s.}$$

Replacing the term on the left side by

$$\bar{p} = \frac{1}{n} \sum_{j=1}^n p_j ,$$

we see that \bar{p} should be reasonably close to the “true” default probability p . Now, a similar argument applies to the sample variances, because we naturally have

$$\frac{1}{n-1} \sum_{j=1}^n \left(p(Y_j) - \overline{p(Y)} \right)^2 \xrightarrow{n \rightarrow \infty} \mathbb{V}[p(Y)] \quad \text{a.s.}$$

where $\overline{p(Y)} = \sum p(Y_j)/n$. This shows that the sample variance

$$s^2 = \frac{1}{n-1} \sum_{j=1}^n (p_j - \bar{p})^2$$

²⁵Here we make the simplifying assumption that the economic cycle, represented by Y_1, \dots, Y_n , is free of *autocorrelation*. In practice one would rather prefer to work with a process incorporating some *intertemporal dependency*, e.g., an AR(1)-process.

should be a reasonable proxy for the “true” variance $\mathbb{V}[p(Y)]$. Recalling Proposition 2.5.9, we obtain

$$\mathbb{V}[p(Y)] = N_2[N^{-1}[p], N^{-1}[p]; \varrho] - p^2, \quad (2.66)$$

and this is all we need for estimating ϱ . Due to our discussion above we can replace the “true” variance $\mathbb{V}[p(Y)]$ by the sample variance σ^2 and the “true” default probability p by the sample mean \bar{p} . After replacing the unknown parameters p and $\mathbb{V}[p(Y)]$ by their corresponding estimated values \bar{p} and s^2 , the asset correlation ϱ is the only “free parameter” in (2.66). It only remains to solve (2.66) for ϱ . The ϱ -values in Tables 2.8 and 2.9 have been calculated by exactly this procedure, hereby relying on the regression-based estimated values μ_i and σ_i^2 . Summarizing one could say that we estimated asset correlations based on the volatility of historic default frequencies.

As a last calculation we want to infer the economic cycle y_1, \dots, y_n for *Regression I*. For this purpose we used an L^2 -solver for calculating y_1, \dots, y_n with

$$\sqrt{\sum_{j=1}^n \sum_{i=1}^6 |p_{ij} - p_i(y_j)|^2} = \min_{(v_1, \dots, v_n)} \sqrt{\sum_{j=1}^n \sum_{i=1}^6 |p_{ij} - p_i(v_j)|^2},$$

where p_{ij} refers to the observed historic loss in rating class R_i in year j , and $p_i(v_j)$ is defined by

$$p_i(v_j) = N \left[\frac{N^{-1}[\bar{p}_i] - \sqrt{\varrho_i} v_j}{\sqrt{1 - \varrho_i}} \right] \quad (i = 1, \dots, 6; j = 1, \dots, 31).$$

Here, ϱ_i refers to the just estimated asset correlations for the respective rating classes. Figure 2.10 shows the result of our estimation of y_1, \dots, y_n . In fact, the result is very intuitive: Comparing the economic cycle y_1, \dots, y_n with the historic mean default path, one can see that any economic downturn corresponds to an increase of default frequencies.

We conclude our example by a brief remark. Looking at Tables 2.8 and 2.9, we find that estimated asset correlations decrease with decreasing credit quality. At first sight this result looks very intuitive, because one could argue that asset correlations increase with firm size, because larger firms could be assumed to carry more systematic risk, and that





FIGURE 2.10
Estimated economic cycle (top) compared to Moody's average historic default frequencies (bottom).

larger firms (so-called “global players”) on average receive better ratings than middle-market corporates. However, although if we possibly see such an effect in the data and our estimations, the uniform portfolio model as we introduced it in this chapter truly is a *two-parameter* model without dependencies between p and ϱ . All possible combinations of p and ϱ can be applied in order to obtain a corresponding loss distribution. From the modeling point of view, there is no rule saying that in case of an increasing p some lower ϱ should be used.

Chapter 3

Asset Value Models

The *asset value model* (AVM) is an important contribution to modern finance. In the literature one can find a tremendous amount of books and papers treating the classical AVM or one of its various modifications. See, e.g., Crouhy, Galai, and Mark [21] (Chapter 9), Sobehart and Keenan [115], and Bohn [13], just to mention a very small selection of especially nicely written contributions.

As already discussed in [Section 1.2.3](#) and also in [Chapter 2](#), two of the most widely used credit risk models are based on the AVM, namely the KMV-Model and CreditMetrics™.

The roots of the AVM are the seminal papers by Merton [86] and Black and Scholes [10], where the *contingent claims* approach to risky debt valuation by *option pricing* theory is elaborated.

3.1 Introduction and a Small Guide to the Literature

The AVM in its original form goes back to Merton [86] and Black and Scholes [10]. Their approach is based on option pricing theory, and we will frequently use this theory in the sequel. For readers not familiar with options we will try to keep our course as self-contained as possible, but refer to the book by Hull [57] for a practitioner's approach and to the book by Baxter and Rennie [8] for a highly readable introduction to the mathematical theory of financial derivatives. Another excellent book more focussing on the underlying stochastic calculus is the one by Lamberton and Lapeyre [76]. For readers without any knowledge of stochastic calculus we recommend the book by Mikosch [87], which gives an introduction to the basic concepts of stochastic calculus with finance in view. To readers with a strong background in probability we recommend the books by Karatzas and Shreve [71,72]. Besides these, the literature on derivative pricing is so voluminous that one can be

sure that there is the optimal book for any reader's taste. All results presented later on can be found in the literature listed above. We therefore will – for the sake of a more fluent presentation – avoid the quotation of particular references but instead implicitly assume that the reader already made her or his particular choice of reference including proofs and further readings.

3.2 A Few Words about Calls and Puts

Before our discussion of Merton's model we want to briefly prepare the reader by explaining some basics on options. The basic assumption underlying option pricing theory is the nonexistence of *arbitrage*, where the word “arbitrage” essentially addresses the opportunity to make a *risk-free profit*. In other words, the common saying that “*there is no free lunch*” is the fundamental principle underlying the theory of financial derivatives.

In the following we will always and without prior notice assume that we are living in a so-called standard¹ Black-Scholes world. In such a world several conditions are assumed to be fulfilled, for example

- stock prices follow *geometric Brownian motions* with constant *drift* μ and constant *volatility* σ ;
- *short selling* (i.e., selling a security without owning it) with full use of proceeds is permitted;
- when buying and selling, *no transaction costs* or *taxes* have to be deducted from proceeds;
- there are *no dividend payments*² during the lifetime of a financial instrument;
- the *no-arbitrage principle* holds;
- security trading is *continuous*;

¹In mathematical finance, various generalizations and improvements of the classical Black-Scholes theory have been investigated.

²This assumption will be kept during the introductory part of this chapter but dropped later on.

- some *riskless* instrument, a so-called *risk-free bond*, can be bought and sold in arbitrary amounts at the riskless rate r , such that, e.g., investing x_0 units of money in a bond today (at time $t = 0$) yields $x_t = x_0 e^{rt}$ units of money at time t ;
- the *risk-free interest rate* $r > 0$ is constant and independent of the maturity of a financial instrument.

As an illustration of how the *no-arbitrage principle* can be used to derive statements about asset values we want to prove the following proposition.

3.2.1 Proposition *Let $(A_t)_{t \geq 0}$ and $(B_t)_{t \geq 0}$ denote the value of two different assets with $A_T = B_T$ at time $T > 0$. Then, if the no-arbitrage principle holds, the values of the assets today (at time 0) also agree, such that $A_0 = B_0$.*

Proof. Assume without loss of generality $A_0 > B_0$. We will show that this assumption contradicts the no-arbitrage principle. As a consequence we must have $A_0 = B_0$. We will derive the contradiction by a simple *investment strategy*, consisting of three steps:

1. short selling of A today, giving us A_0 units of money today;
2. buying asset B today, hereby spending B_0 units of money;
3. investing the residual $A_0 - B_0 > 0$ in the riskless bond today.

At time T , we first of all receive back the money invested in the bond, so that we collect $(A_0 - B_0)e^{rT}$ units of money. Additionally we have to return asset A , which we sold at time $t = 0$, without possessing it. Returning some asset we do not have means that we have to fund the purchase of A . Fortunately we bought B at time $t = 0$, such that selling B for a price of B_T just creates enough income to purchase A at a price of $A_T = B_T$. So for clearing our accounts we were not forced to use the positive payout from the bond, such that at the end we have made some risk-free profit. \square

The investment strategy in the proof of Proposition 3.2.1 is “risk-free” in the sense that the strategy yields some positive profit no matter what the value of the underlying assets at time T might be. The information that the assets A and B will agree at time T is sufficient

for locking-in a guaranteed positive net gain if the asset values at time 0 differ.

Although Proposition 3.2.1 and its proof are almost trivial from the content point of view, they already reflect the typical proof scheme in option pricing theory: For proving some result, the opposite is assumed to hold and an appropriate investment strategy is constructed in order to derive a contradiction to the no-arbitrage principle.

3.2.1 Geometric Brownian Motion

In addition to our bond we now introduce some risky asset A whose values are given by a stochastic process $A = (A_t)_{t \geq 0}$. We call A a *stock* and assume that it evolves like a *geometric Brownian motion* (gBm). This means that the process of asset values is the solution of the *stochastic differential equation*

$$A_t - A_0 = \mu_A \int_0^t A_s ds + \sigma_A \int_0^t A_s dB_s \quad (t \geq 0), \quad (3. 1)$$

where $\mu_A > 0$ denotes the *drift* of A , $\sigma_A > 0$ addresses the *volatility* of A , and $(B_s)_{s \geq 0}$ is a standard Brownian motion; see also (3. 14) where (3. 1) is presented in a slightly more general form incorporating *dividend payments*. Readers with some background in stochastic calculus can easily solve Equation (3. 1) by an application of *Itô’s formula* yielding

$$A_t = A_0 \exp\left(\left(\mu_A - \frac{1}{2} \sigma_A^2\right)t + \sigma_A B_t\right) \quad (t \geq 0). \quad (3. 2)$$

This formula shows that gBm is a really intuitive process in the context of stock prices respectively asset values. Just recall from elementary calculus that the exponential function $f(t) = f_0 e^{ct}$ is the unique solution of the differential equation

$$df(t) = cf(t)dt, \quad f(0) = f_0.$$

Writing (3. 1) formally in the following way,

$$dA_t = \mu_A A_t dt + \sigma_A A_t dB_t, \quad (3. 3)$$

shows that the first part of the stochastic differential equation describing the evolution of gBm is just the “classical” way of describing exponential growth. The difference turning the exponential growth function

into a stochastic process arises from the *stochastic differential* w.r.t. Brownian motion captured by the second term in (3. 3). This differential adds some *random noise* to the exponential growth, such that instead of a smooth function the process evolves as a random walk with almost surely nowhere differentiable paths. If price movements are of exponential growth, then this is a very reasonable model. Figure 1.6 actually shows a simulation of two paths of a gBm.

Interpreting (3. 3) in a naive nonrigorous way, one can write

$$\frac{A_{t+dt} - A_t}{A_t} = \mu_A dt + \sigma_A dB_t .$$

The right side can be identified with the *relative return* of asset A w.r.t. an “infinitesimal” small time interval $[t, t + dt]$. The equation then says that this return has a *linear trend* with “slope” μ_A and some *random fluctuation term* $\sigma_A dB_t$. One therefore calls μ_A the *mean rate of return* and σ_A the *volatility* of asset A . For $\sigma_A = 0$ the process would be a deterministic exponential function, smooth and without any fluctuations. In this case any investment in A would yield a riskless profit only dependent on the time until payout. With increasing volatility σ_A , investments in A become more and more risky. The stronger fluctuations of the process bear a potential of higher wins (*upside potential*) but carry at the same time a higher risk of downturns respectively losses (*downside risk*). This is also expressed by the expectation and volatility functions of gBm, which are given by

$$\mathbb{E}[A_t] = A_0 \exp(\mu_A t) \tag{3. 4}$$

$$\mathbb{V}[A_t] = A_0^2 \exp(2\mu_A t) (\exp(\sigma_A^2 t) - 1) .$$

As a last remark we should mention that there are various other stochastic processes that could be used as a model for price movements. In fact, in most cases asset values will not evolve like a gBm but rather follow a process yielding fatter tails in their distribution of log-returns (see e.g. [33]).

3.2.2 Put and Call Options

An *option* is a contract written by an *option seller* or *option writer* giving the *option buyer* or *option holder* the *right* but not the obligation to buy or sell some specified asset at some specified time for some specified price. The time where the option can be *exercised* is called

the *maturity* or *exercise date* or *expiration date*. The price written in the option contract at which the option can be exercised is called the *exercise price* or *strike price*.

There are two basic types of options, namely a *call* and a *put*. A call gives the option holder the right to *buy* the underlying asset for the strike price, whereas a put guarantees the option holder the right to *sell* the underlying asset for the exercise price. If the option can be exercised only at the maturity of the option, then the contract is called a *European* option. If the option can be exercised at any time until the final maturity, it is called an *American* option.

There is another terminology in this context that we will frequently use. If someone wants to purchase an asset she or he does not possess at present, she or he currently is *short* in the asset but wants to go *long*. In general, every option contract has two sides. The investor who purchases the option takes a *long position*, whereas the option writer has taken a *short position*, because he sold the option to the investor.

It is always the case that the writer of an option receives cash up front as a compensation for writing the option. But receiving money today includes the potential liabilities at the time where the option is exercised. The question every option buyer has to ask is whether the right to buy or sell some asset by some later date for some price specified today is worth the price she or he has to pay for the option. This question actually is the basic question of *option pricing*.

Let us say the underlying asset of a European call option has price movements $(A_t)_{t \geq 0}$ evolving like a gBm, and the strike price of the call option is F . At the maturity time T one can distinguish between two possible scenarios:

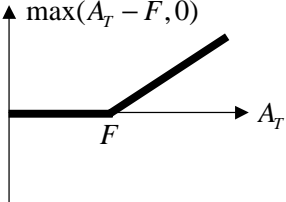
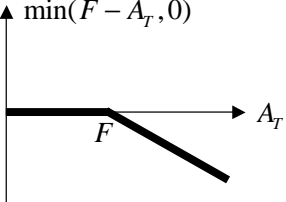
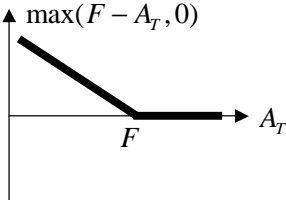
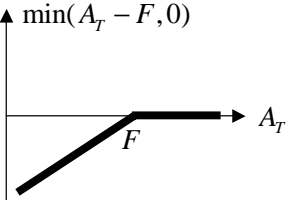
1. Case: $A_T > F$

In this case the option holder will definitely exercise the option, because by exercising the option he can get an asset worth A_T for the better price F . He will make a net profit in the deal, if the price C_0 of the call is smaller than the price advantage $A_T - F$.

2. Case: $A_T \leq F$

If the asset is cheaper or equally expensive in the market compared to the exercise price written in the option contract, the option holder will not exercise the option. In this case, the contract was good for nothing and the price of the option is the investor's loss.

TABLE 3.1: Four different positions are possible in plain-vanilla option trading.

	LONG	SHORT
CALL	<ul style="list-style-type: none"> ▪ buyer/holder of option ▪ payer of option price ▪ option to buy the asset ▪ payoff: $\uparrow \max(A_T - F, 0)$  	<ul style="list-style-type: none"> ▪ seller/writer of option ▪ receiver of option price ▪ obligation upon request of option holder to deliver the asset ▪ payoff: $\uparrow \min(F - A_T, 0)$ 
PUT	<ul style="list-style-type: none"> ▪ buyer/holder of option ▪ payer of option price ▪ option to sell the asset ▪ payoff: $\uparrow \max(F - A_T, 0)$  	<ul style="list-style-type: none"> ▪ seller/writer of option ▪ receiver of option price ▪ obligation upon request of option holder to buy the asset ▪ payoff: $\uparrow \min(A_T - F, 0)$ 

Both cases can be summarized in the *payoff* function of the option, which, in the case of a European call with strike F , is given by

$$\pi : \mathbb{R} \rightarrow \mathbb{R}, \quad A_T \mapsto \pi(A_T) = \max(A_T - F, 0) .$$

There are altogether four positions in option trading with calls and puts: long call, short call, long put, and short put. Table 3.1 summarizes these four positions and payoffs, clearly showing that for a fixed type of option the payoff of the seller is the reverse of the payoff of the buyer of the option. Note that in the table we have neglected the price of the option, which would shift the payoff diagram along the y -axis, namely into the negative for long positions (because the option price has to be paid) and into the positive for short positions (because the option price will be received as a compensation for writing the option).

It is interesting to mention that long positions have a *limited downside risk*, because the option buyer's worst case is that the money invested in the option is lost in total. The good news for option buyers is the *unlimited upside chance*. Correspondingly option writers have an *unlimited downside risk*. Moreover, the best case for option writers is that the option holder does not exercise the option. In this case the option price is the net profit of the option writer.

At first glance surprising, European calls and puts are related by means of a formula called the *put-call parity*.

3.2.2 Proposition *Let C_0 respectively P_0 denote the price of a European call respectively put option with strike F , maturity T , and underlying asset A . The risk-free rate is denoted by r . Then,*

$$C_0 + Fe^{-rT} = P_0 + A_0 .$$

This formula is called the put-call parity, connecting puts and calls.

Proof. For proving the proposition we compare two portfolios:

- a long call plus some investment Fe^{-rt} in the risk-free bond;
- a long put plus an investment of one share in asset A .

According to Proposition 3.2.1 we only have to show that the two portfolios have the same value at time $t = T$, because then their values at time $t = 0$ must also agree due to the no-arbitrage principle. We calculate their values at maturity T . There are two possible cases:

$A_T \leq F$: In this case the call option will not be exercised such that the value of the call is zero. The investment Fe^{-rT} in the bond at time $t = 0$ will payout exactly the amount F at $t = T$, such that the value of the first portfolio is F . But the value of the second portfolio is also F , because exercising the put will yield a payout of $F - A_T$, and adding the value of the asset A at $t = T$ gives a total pay out of $F - A_T + A_T = F$.

$A_T > F$: In the same manner as in the first case one can verify that now the value of the first and second portfolio equals A_T .

Altogether the values of the two portfolios at $t = T$ agree. \square

The put-call parity only holds for European options, although it is possible to establish some relationships between American calls and puts for a nondividend-paying stock as underlying.

Regarding call options we will now show that it is never optimal to exercise an American call option on a nondividend-paying stock before the final maturity of the option.

3.2.3 Proposition *The price of a European and an American call option are equal if they are written w.r.t. the same underlying, maturity, and strike price.*

Proof. Again we consider two portfolios:

- one American call option plus some cash amount of size Fe^{-rT} ;
- one share of the underlying asset A .

The value of the cash account at maturity is F . If we would force a payout of cash before expiration of the option, say at time t , then the value of the cash account would be $Fe^{-r(T-t)}$. Because American options can be exercised at any time before maturity, we can exercise the call in portfolio one in order to obtain a portfolio value of

$$A_t - F + Fe^{-r(T-t)} < A_t \quad \text{for } t < T .$$

Therefore, if the call option is exercised before the expiration date, the second portfolio will in all cases be of greater value than the first portfolio. If the call option is treated like a European option by exercising it at maturity T , then the value of the option is $\max(A_T - F, 0)$, such that the total value of the second portfolio equals $\max(A_T, F)$. This

shows that an American call option on a nondividend-paying stock never should be exercised before the expiration date. \square

In 1973 Fischer Black and Myron Scholes found a first analytical solution for the valuation of options. Their method is not too far from the method we used in Propositions 3.2.1 and 3.2.2: By constructing a riskless portfolio consisting of a combination of calls and shares of some underlying stock, an application of the no-arbitrage principle established an analytical price formula for European call options on shares of a stock. The pricing formula depends on five parameters:

- the share or asset price A_0 as of today;
- the volatility σ_A of the underlying asset A ;
- the strike price F of the option;
- the time to maturity T of the option;
- the risk-free interest rate $r > 0$.

Here we should mention that a key concept leading to the option pricing formulas presented below is the so-called *risk-neutral valuation*. In a world where all investors are *risk-neutral*, all securities earn the risk-free rate. This is the reason why the Black-Scholes formulas do not depend on the drift μ_A of $(A_t)_{t \geq 0}$. *In an arbitrage-free complete market, arbitrage prices of contingent claims equal their discounted expected values under the risk-neutral martingale measure.* Because we will just apply the option pricing formulas without being bothered about their deeper mathematical context, we refer to the literature for further reading. A comprehensive treatment of the mathematical theory of risk-neutral valuation is the book by Bingham and Kiesel [9].

The pricing formula for European calls is then given by

3.2.4 Proposition *The Black-Scholes price of a European call option with parameters (A_0, σ_A, F, T, r) is given by*

$$A_0 N[d_1] - e^{-rT} FN[d_2], \quad \text{where}$$
$$d_1 = \frac{\log(A_0/F) + (r + \sigma_A^2/2)T}{\sigma_A \sqrt{T}},$$

$$d_2 = \frac{\log(A_0/F) + (r - \sigma_A^2/2)T}{\sigma_A\sqrt{T}} = d_1 - \sigma_A\sqrt{T} .$$

As usual, $N[\cdot]$ denotes the cumulative standard normal distribution function. In the sequel we write $C_0(A_0, \sigma_A, F, T, r)$ to denote this price.

Proof. A proof can be found in the literature mentioned at the beginning of this chapter. \square

Because the prices of a European and an American call option agree due to Proposition 3.2.3, Proposition 3.2.4 also provides the pricing formula for American calls on a nondividend-paying stock. For European put options, the pricing formula follows by an application of the put-call parity.

3.2.5 Proposition *The Black-Scholes price of a European put option with parameters (A_0, σ_A, F, T, r) is given by*

$$e^{-rT} FN[-d_2] - A_0 N[-d_1] , \quad \text{where}$$

$$d_1 = \frac{\log(A_0/F) + (r + \sigma_A^2/2)T}{\sigma_A\sqrt{T}} ,$$

$$d_2 = \frac{\log(A_0/F) + (r - \sigma_A^2/2)T}{\sigma_A\sqrt{T}} = d_1 - \sigma_A\sqrt{T} .$$

In the sequel we write $P_0(A_0, \sigma_A, F, T, r)$ to denote this price.

Proof. The put-call parity from Proposition 3.2.2 yields

$$P_0(A_0, \sigma_A, F, T, r) = C_0(A_0, \sigma_A, F, T, r) + Fe^{-rT} - A_0 .$$

Evaluating the right side of the equation proves the proposition. \square

For American put option prices one has to rely on numerical methods, because no closed-form analytic formula is known.

3.3 Merton's Asset Value Model

In this chapter we describe the “classical” asset value model introduced by Merton. As always we assume all involved random variables

to be defined on a suitable common probability space. Additionally we make some typical economic assumptions. For example, we assume that markets are *frictionless* with no taxes and without bankruptcy costs. The no-arbitrage principle is assumed to hold. The complete set of conditions necessary for the Merton model can be found in the literature.

3.3.1 Capital Structure: Option-Theoretic Approach

Let's say we consider a firm with risky assets A , such that its asset value process $(A_t)_{t \geq 0}$ follows a gBm. The basic assumption now is that the firm is financed by means of a very simple *capital structure*, namely one *debt obligation* and one type of *equity*. In this case one can write

$$A_0 = E_0 + D_0 , \quad (3.5)$$

where $(E_t)_{t \geq 0}$ is a gBm describing the evolution of equity of the firm, and $(D_t)_{t \geq 0}$ is some stochastic process describing the market value of the debt obligation of the firm, which is assumed to have the cash profile of a *zero coupon bond* with maturity T and interest-adjusted *face value* F . By “interest-adjusted” we mean that F already includes some accrued interest at a rate reflecting the borrowing company's riskiness. The cash profile of debt is then very simple to describe: Debt holders pay a capital of D_0 to the firm at time $t = 0$, and at time $t = T$ they receive an amount equal to F , where F includes the principal D_0 plus the just-mentioned interest payment compensating for the credit risk associated with the credit deal. From the point of view of debt holders, credit risk arises if and only if

$$\mathbb{P}[A_T < F] > 0 ,$$

meaning that with positive probability the value of the borrowing company's assets at the debt's maturity is not sufficient for covering the payment F to debt holders. In case this *default probability* is greater than zero, one immediately can conclude that

$$D_0 < F e^{-rT} ,$$

where r denotes the risk-free interest rate. This inequality must hold because debt holders want some compensation for the credit respectively default risk of its obligor. Such a *risk premium* can be charged

TABLE 3.2: Credit protection by a suitable put option.

	asset value	debt holder's cash flows	debt holder's payout
$t = 0$	A_0	$-D_0$ (lend money) $-P_0$ (purchase put)	$-D_0 - P_0$
$t = T$	$A_T < F$	A_T (recovery) $F - A_T$ (apply put)	F
$t = T$	$A_T \geq F$	F (receive face value) 0	F

implicitly by means of discounting the face value F at a rate higher than the risk-free rate. The payout of debt to the obligor at time $t = 0$ will then be smaller the more risky the obligor's business is.

A typical strategy of debt holders (e.g., a lending bank) is the attempt to *neutralize* the credit risk by purchasing some kind of credit protection. In our case a successful strategy is to buy a suitable derivative. For this purpose, debt holders take a *long position in a put option on A with strike F and maturity T* ; see also Figure 3.1. Table 3.2 shows that purchasing the put option guarantees credit protection against the default risk of the borrowing company, because at the maturity date $t = T$ the debt holder's payout equals F no matter if the obligor defaults or not. Therefore, the credit risk of the loan is neutralized and completely hedged. In other words, buying the put transforms the risky corporate loan³ into a riskless bullet loan with face value F . This brings us to an important conclusion: Taking the hedge into account, the portfolio of debt holders consists of a put option and a loan. Its value at time $t = 0$ is $D_0 + P_0(A_0, \sigma_A, F, T, r)$. The risk-free payout of this portfolio at time $t = T$ is F . Because we assumed the no-arbitrage principle to hold, the payout of the portfolio has to be discounted to its present value at the risk-free rate r . This implies

$$D_0 + P_0(A_0, \sigma_A, F, T, r) = Fe^{-rT} ,$$

so that the present value of debt,

$$D_0 = Fe^{-rT} - P_0(A_0, \sigma_A, F, T, r) , \quad (3.6)$$

³Which will be a bond in most cases.

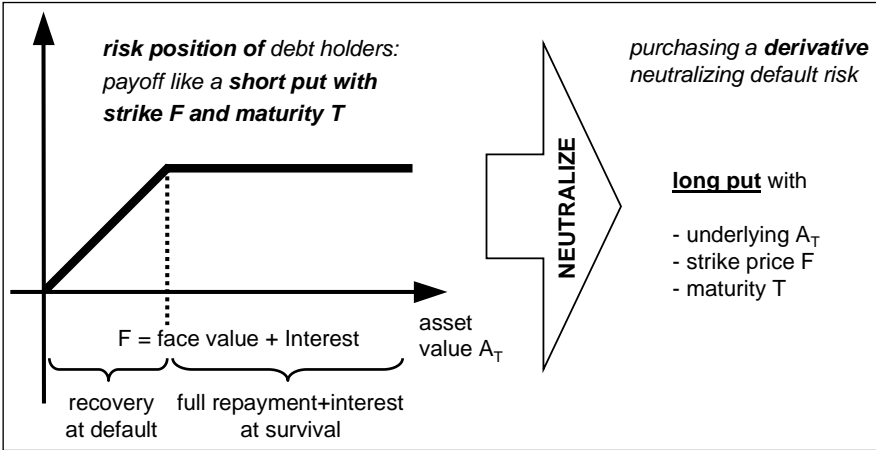


FIGURE 3.1
Hedging default risk by a long put.

is the present value of the face value F discounted at the risk-free rate r corrected by the price for hedging the credit risk by means of a put option.

3.3.1 Conclusion [Option-theoretic interpretation of debt]

From the company's point of view, the debt obligation can be described by taking a long position in a put option. From the debt holders point of view, the debt obligation can be described by writing a put option to the company.

Proof. Using the notation above, at time $t = T$ the company has to pay debt back to debt holders. This yields a cash flow

$$\underbrace{\max(F - A_T, 0)}_{\text{long put payoff}} - F$$

from the company's point of view. From the debt holder's point of view, the cash flow can be written as

$$F + \underbrace{\min(A_T - F, 0)}_{\text{short put payoff}}$$

units of money at time $t = T$. □

So we have found an interpretation of one component of the capital structure of the company in terms of options. But the other component, equity, can also be embedded in an option-theoretic concept: The equity or share holders of the firm have the right to *liquidate* the firm, i.e., paying off the debt and taking over the remaining assets. Let's say equity holders decide to liquidate the firm at the maturity date T of the debt obligation. There are two possible scenarios:

- $A_T < F$:

This is the default case, where the asset value at maturity is not sufficiently high for paying back debt holders in full. There are no assets left that could be taken over by the equity holders, such that their payoff is zero.

- $A_T \geq F$:

In this case, there is a net profit to equity holders of $A_T - F$ after paying back the debt.

Summarizing both cases we see that the total payoff to equity holders is $\max(A_T - F, 0)$, which is the payoff of a European call option on A with strike F and maturity T ; see [Table 3.1](#). Due to Proposition 3.2.1 the present value of equity therefore is given by

$$E_0 = C_0(A_0, \sigma_A, F, T, r) . \quad (3.7)$$

We conclude as follows:

3.3.2 Conclusion [Option-theoretic interpretation of equity]

From the company's point of view, equity can be described by selling a call option to equity holders. Consequently, the position of equity holders is a long call on the firm's asset values.

Proof. The proof follows from the discussion above. \square

Combining (3.5) with Conclusions 3.6 and 3.7 we obtain

$$A_0 = E_0 + D_0 = C_0(A_0, \sigma_A, F, T, r) + Fe^{-rT} - P_0(A_0, \sigma_A, F, T, r) .$$

Rearranging, we get

$$A_0 + P_0(A_0, \sigma_A, F, T, r) = C_0(A_0, \sigma_A, F, T, r) + Fe^{-rT} ,$$

which is nothing but the *put-call parity* we proved in Proposition 3.2.2.

Note that Conclusion 3.3.2 will not be harmed if one allowed equity holders to exercise the option before the maturity T . As a justification recall Proposition 3.2.3, saying that the price of a call option is the same no matter if it is European or American.

Our discussion also shows that equity and debt holders have *contrary risk preferences*. To be more explicit, consider

$$C_0(A_0, \sigma_A, F, T, r) = A_0 - Fe^{-rT} + P_0(A_0, \sigma_A, F, T, r) .$$

As can be found in the literature, increasing the riskiness of the investment by choosing some asset A with higher volatility σ_A will also increase the option premium C_0 and P_0 of the call and put options. Therefore, increased volatility (higher risk) is

- good for equity holders, because their natural risk position is a long call, and the value of the call increases with increasing volatility;
- bad for debt holders, because their natural risk position⁴ is a short put, whose value decreases with increasing volatility.

Note the unsymmetry in the position of equity holders: Their downside risk is limited, because they can not lose more than their invested capital. In contrast, their upside potential is unlimited. The better the firm performs, the higher the value of the firm's assets, the higher the remaining of assets after a repayment of debt in case the equity holders liquidate the firm.

3.3.2 Asset from Equity Values

The general problem with asset value models is that asset value processes are *not observable*. Instead, what people see every day in the stock markets are equity values. So the big question is how asset values can be derived from market data like equity processes. Admittedly, this is a very difficult question. We therefore approach the problem from two sides. In this section we introduce the classical concept of Merton, saying how one could solve the problem *in principle*. In the next section we then show a way how the problem can be tackled *in practice*.

⁴Which could only be neutralized by a long put.

We follow the lines of a paper by Nickell, Perraudin, and Varotto [101]. In fact, there are certainly more working approaches for the construction of asset values from market data. For example, in their published papers (see, e.g., Crosbie [19]) KMV incorporates the *classical* Merton model, but it is well known that in their commercial software (see [Section 1.2.3](#)) they have implemented a different, more complicated, and undisclosed algorithm for translating equity into asset values.

The classical approach is as follows: The process of a firm's equity is observable in the market and is given by the company's *market capitalization*, defined by

$$[\text{number of shares}] \times [\text{value of one share}] .$$

Also observable from market data is the *volatility* σ_E of the firm's equity process. Additional information we can get is the *book value* of the firm's *liabilities*. From these three sources,

- equity value of the firm,
- volatility of the firm's equity process, and
- book value of the firm's liabilities,

we now want to infer the *asset value process* $(A_t)_{t \geq 0}$ (as of today). Once more we want to remark that the following is more a "schoolbook model" than a working approach. In contrast, the next paragraph will show a more applicable solution.

Let us assume we consider a firm with the same simple capital structure⁵ as introduced in (3. 5). From Conclusion 3.3.2 we already know that the firm's equity can be seen as a call option on the firm's assets, written by the firm to the equity or share holders of the firm. The strike price F is determined by the book value of the firm's liabilities, and the maturity T is set to the considered planning horizon, e.g., one year. According to (3. 7) this option-theoretic interpretation of equity yields the functional relation

$$E_t = C_t(A_t, \sigma_A, F, (T - t), r) \quad (t \in [0, T]) \quad (3. 8)$$

⁵Actually it is in part due to the assumption of a simple capital structure that the classical Merton model is not really applicable in practice.

This functional relation can be locally inverted, due to the *implicit function theorem*, in order to solve (3. 8) for A_t . Therefore, the asset value of the firm can be calculated as a function of the firm's equity and the parameters F , t , T , r , and the asset volatility σ_A . If, as we already remarked, asset value processes are not observable, the asset volatility also is not observable. It therefore remains to determine the asset volatility σ_A in order to obtain A_t from (3. 8).

Here, we actually need some insights from stochastic calculus, such that for a brief moment we are now forced to use results for which an exact and complete explanation is beyond the scope of the book. However, in the next section we will provide some "heuristic" background on *pathwise stochastic integrals*, such that at least some open questions will be answered later on. As always we assume for the sequel that all random variables respectively processes are defined on a suitable common probability space.

Recall that we assumed that the asset value process $(A_t)_{t \geq 0}$ is assumed to evolve like a *geometric Brownian motion* (see [Section 3.2.1](#)), meaning that A solves the stochastic differential equation

$$A_t - A_0 = \mu_A \int_0^t A_s ds + \sigma_A \int_0^t A_s dB_s^{(A)}.$$

Following almost literally the arguments in Merton's approach, we assume for the equity of the firm that $(E_t)_{t \geq 0}$ solves the stochastic differential equation

$$E_t - E_0 = \int_0^t \mu_E(s) E_s ds + \int_0^t \sigma_E(s) E_s dB_s^{(E)}. \quad (3. 9)$$

Here, $(B_t^{(A)})_{t \geq 0}$ and $(B_t^{(E)})_{t \geq 0}$ denote *standard Brownian motions*. Applying *Itô's lemma* to the function⁶

$$f(t, A_t) = C_t(A_t, \sigma_A, F, (T - t), r)$$

⁶We refer to the literature for checking that the conditions necessary for applying Itô's lemma are satisfied in our case.

and comparing⁷ the martingale part of the resulting equation with the martingale part of Equation (3. 9) yield in informal differential notation

$$\sigma_E E_t dB_t^{(E)} = f_2(t, A_t) \sigma_A A_t dB_t^{(A)}, \quad (3. 10)$$

where $f_2(\cdot, \cdot)$ denotes the partial derivative w.r.t. the second component. But the coefficients of stochastic differential equations are uniquely determined, such that from (3. 10) we can conclude

$$\frac{\sigma_E}{\sigma_A} = \frac{A_t f_2(t, A_t)}{E_t}. \quad (3. 11)$$

Solving (3. 11) for σ_A and inserting the solution into Equation (3. 8) yields A_t for $t \in [0, T]$.

This concludes our discussion of the classical Merton model. We now proceed to a more mathematical as well as more applicable approach. For this purpose, we explicitly define the stochastic integral for a specific class of integrands in [Section 3.4.1](#). Then, in [Section 3.4.2](#), we present a more accurate derivation of the Black-Scholes partial differential equation due to Duffie [28]. Additionally, we introduce a boundary condition going back to Perraudin et al. [101] which specifies a reasonable relation between asset values and equities.

3.4 Transforming Equity into Asset Values: A Working Approach

Let us begin with a few words on pathwise Itô Calculus (see Revuz and Yor [108], and Foellmer [42]). The following treatment is rather self-contained because no difficult prerequisites from measure theory are required. Unfortunately, the pathwise calculus is only valid for a specific type of *trading strategies*, as we will later see.

3.4.1 Itô’s Formula “Light”

In this paragraph we want to establish the existence of a *pathwise stochastic integral* by an argument based on elementary calculus, thereby avoiding the usual requirements from measure theory.

⁷Such a comparison is justified, because the components of so-called *Itô processes* are uniquely determined.

Let ω be a real-valued continuous function of time t with finite quadratic variation $\langle \omega \rangle$, and $F \in \mathcal{C}^2$. Denoting by \mathcal{Z}_n a sequence of partitions of the interval $[0, t)$ with $\text{mesh}(\mathcal{Z}_n) \rightarrow 0$, a Taylor expansion up to second order yields

$$\begin{aligned}
 F(\omega_t) - F(\omega_0) &= \lim_{n \rightarrow \infty} \left(\sum_{(t_i) \in \mathcal{Z}_n^t} F'(\omega_{t_i})(\omega_{t_{i+1}} - \omega_{t_i}) \right. & (3.12) \\
 &+ \left. \sum_{(t_i) \in \mathcal{Z}_n^t} \frac{1}{2} F''(\omega_{t_i})(\omega_{t_{i+1}} - \omega_{t_i})^2 + o((\Delta\omega)^2) \right).
 \end{aligned}$$

From the existence of the quadratic variation of ω we conclude that the second term in (3.12) converges to

$$\frac{1}{2} \int_0^t F''(\omega_s) d\langle \omega \rangle_s.$$

Hence the limit of the first term in (3.12) also exists. It is denoted by

$$\int_0^t F'(\omega_s) d\omega_s$$

and called a *stochastic integral*. In this context, the *Itô formula* is just a by-product of the Taylor expansion (3.12), and can be obtained by writing (3.12) in the limit form

$$F(\omega_t) - F(\omega_0) = \int_0^t F'(\omega_s) d\omega_s + \frac{1}{2} \int_0^t F''(\omega_s) d\langle \omega \rangle_s. \quad (3.13)$$

The just-derived stochastic integral can be interpreted in terms of trading gains. The discrete approximation

$$\sum_{t_i \in \mathcal{Z}_n^t} F'(\omega_{t_i})(\omega_{t_{i+1}} - \omega_{t_i})$$

of the stochastic integral is the *gain* of the following trading strategy:

Buy $F'(\omega_{t_i})$ shares of a financial instrument with value ω_t at time t_i .

The gain over the time interval $[t_i, t_{i+1})$ then equals

$$F'(\omega_{t_i})(\omega_{t_{i+1}} - \omega_{t_i}).$$

The stochastic integral is just the limit of the sum over all these trading gains in the interval $[0, t)$. From these observations it becomes

also clear why the stochastic integral as introduced above sometimes is called *non-anticipating*. This terminology just refers to the fact that the investment took place at the beginning of the intervals $[t_i, t_{i+1})$.

For a thorough introduction to the stochastic integral in the more general measure-theoretic setting we refer to the literature mentioned at the beginning of this chapter. However, the intuitive interpretation of the stochastic integral as the gain of a (non-anticipating) trading strategy and the basic structure of the Itô formula remain both valid in the measure-theoretic approach.

3.4.2 Black-Scholes Partial Differential Equation

In this paragraph we follow the approach outlined in Duffie [28]. As in the previous paragraphs, we assume that the asset value process $A = (A_t)_{t \geq 0}$ follows a geometric Brownian motion driven by some Brownian motion B . But this time we include *dividend payments*, such that A is the solution of the stochastic differential equation

$$A_t - A_0 = \int_0^t (\mu_A A_s - C_{A,s}) ds + \sigma_A \int_0^t A_s dB_s, \quad (3.14)$$

where $C_{A,s}$ is the dividend paid by the firm at time s . In the literature the following intuitive differential notation of (3.14) also is used

$$dA_t = (\mu_A A_t - C_{A,t}) dt + \sigma_A A_t dB.$$

In previous paragraphs the capital structure of the considered firm contained one debt obligation. Here we assume that the market value of debt D_t at time t is just a nonstochastic exponential function,

$$D_s = D_0 e^{\mu_D s}.$$

By Itô's formula (3.13), any process $(E_t)_{t \geq 0}$ represented by a smooth function $E(x, y, t)$ applied to the processes A and D ,

$$E_t = E(A_t, D_t, t), \quad E \in \mathcal{C}^{2,1,1},$$

solves the integral equation

$$\begin{aligned} E_t - E_0 &= \int_0^t [\partial_t E(A_s, D_s, s) + (\mu_A A_s - C_{A,s}) \partial_x E(A_s, D_s, s) \\ &\quad + \mu_D D \partial_y E(A_s, D_s, s) + \frac{1}{2} \sigma_A^2 A_s^2 \partial_{xx} E(A_s, D_s, s)] ds \end{aligned}$$

$$+ \sigma_A \int_0^t A \partial_x E(A_s, D_s, s) dB_s .$$

We now want to construct a so-called *self-financing* trading strategy (η_t, θ_t) such that

- $\eta_t A_t + \theta_t K_t = E_t$
- $K_t = e^{rt}$,

where K_t denotes the value of a risk-free investment (e.g., some treasury bond) earning interest at the risk-free rate r . The attribute *self-financing* means that the value of the portfolio, at time t consisting of η_t shares of A and θ_t share of K , has a value equal to the initial investment plus trading gains. More explicitly,

$$\eta_t A_t + \theta_t K_t = \eta_0 A_0 + \theta_0 K_0 + \int_0^t \eta_s dA_s + \int_0^t \theta_t dK_t . \quad (3. 15)$$

The assumption that there is a self-financing strategy⁸ that perfectly replicates E_t leads to

$$E_t = \int_0^t [\eta_s \mu_A A_s + \theta_s K_s r] ds + \int_0^t \eta_s \sigma_A A_s dB_s \quad (3. 16)$$

$$\begin{aligned} &= \int_0^t [\partial_t E(A_s, D_s, s) + (\mu_A A_s - C_{A,s}) \partial_x E(A_s, D_s, s) \\ &+ \mu_D D \partial_y E(A_s, D_s, s) + \frac{1}{2} \sigma_A^2 A_s^2 \partial_{xx} E(A_s, D_s, s) ds \\ &+ \sigma_A \int_0^t \partial_x E(A_s, D_s, s) A_s dB_s . \end{aligned} \quad (3. 17)$$

The unique decomposition of an Itô process into a stochastic integral with respect to K and a drift leads to

$$\eta_t = \partial_x E(A_t, D_t, t)$$

⁸A straightforward application of Itô's formula would imply that

$$\eta_t A_t - \eta_0 A_0 = \int_0^t A_s d\eta_s + \int_0^t \eta_s dA_s + \langle \eta, A \rangle_t .$$

This result would not lead to the Black-Scholes PDE. "Self-financing" therefore is essential from a mathematical point of view.

taking (3. 16) and (3. 17) into account. Since the trading strategy replicates E_t we necessarily have

$$\theta_t = \frac{1}{K_t} [E_t - \partial_x E(A_t, D_t, t) A_t] .$$

The comparison of the coefficient of dt implies the equation

$$0 = \partial_t E(A_s, D_s, s) + (rA_s - C_{A,s}) \partial_x E(A_s, D_s, s) \tag{3. 18}$$

$$+ \mu_D D \partial_y E(A_s, D_s, s) + \frac{1}{2} \sigma_A^2 A_s^2 \partial_{xx} E(A_s, D_s, s) - rE_s .$$

As in Nickell, Perraudin and Varotto [101], let us now specify the dividend by $C_{A,s} = \delta A_s$. Then, E solves Equation (3. 18) if it solves the partial differential equation

$$0 = \partial_t E(x, y, s) + (rx - \delta x) \partial_x E(x, y, s) + \mu_D D \partial_y E(x, y, s) \tag{3. 19}$$

$$+ \frac{1}{2} \sigma_A^2 x^2 \partial_{xx} E(x, y, s) - rE(x, y, s) .$$

For $\delta = 0$ and $D_0 = 0$, the last equation becomes the celebrated *Black-Scholes formula*. It should be clear that we have to specify boundary conditions for (3. 19).

As a first approach, let us assume that the function E is independent of the third component, i.e., $E_t = E(A_t, D_t)$. Then, $\partial_t E = 0$ and Equation (3. 19) becomes an ordinary differential equation.

Analogous to the lines in [101] we now assume that the firm is declared to be in bankruptcy as soon as the ratio of *assets to liabilities* A_t/D_t hits some low level for the very first time. We call this critical threshold γ and assume the equity-holders to receive no money in case of a bankruptcy settlement. Then, the value of the firm's equity, E , satisfies the differential equation (3. 19) subject to the following boundary conditions:

$$E(A, D)|_{A/D=\gamma} = 0 \quad \text{and} \quad \lim_{A/D \rightarrow \infty} E(A, D) = A - \frac{\delta}{r} D .$$

For some background on differential equations, refer to [129].

Now let us present the solution of (3. 19) under these boundary conditions. It is given by

$$E(A, D) = D \left[\frac{A}{D} - \frac{\delta}{r - \mu_D} - \left(\gamma - \frac{\delta}{r - \mu_D} \right) \left(\frac{A/D}{\gamma} \right)^\lambda \right] , \tag{3. 20}$$

where λ is defined in dependence on σ_A by

$$\begin{aligned} \lambda &= \lambda(\sigma_A) = \\ &= \frac{1}{\sigma_A^2} \left[\left(\frac{\sigma_A^2}{2} + \delta + \mu_D - r \right) - \sqrt{(r - \sigma_A^2 - \delta - \mu_D)^2 + 2\sigma_A^2(r - \mu_D)} \right]. \end{aligned}$$

In this model, the level of the bankruptcy trigger γ is chosen by the equity holders, since the firm will continue to operate until equity holders are unwilling to absorb more than the already occurred losses. The threshold γ therefore is determined by the first order condition $\partial_\gamma E = 0$, from which it follows that

$$\gamma = \frac{\lambda}{\lambda - 1} \frac{\delta}{r - \mu_D}.$$

Figure 3.2 shows the asset-equity relation from Equation (3. 20) for some parameter sets.

Now, if A would be an observable variable and σ_A would be known, E would be specified by (3. 20). But in “option terminology” we observe the price of an option, namely the equity price. Therefore, we can only estimate the volatility of the changes in the price of the option. From this we have to find the value of the underlying instrument A and its volatility σ_A . This means for determining A and σ_A we need a second equation. Since E is an Itô process, its quadratic variation $\langle E \rangle$ can be read off from Equation (3. 17) as

$$\langle E \rangle_t = \int_0^t \sigma_{E,s} ds = \sigma_A^2 \int_0^t A_s^2 \partial_A^2 E(A_s, D_s, s) ds. \quad (3. 21)$$

Therefore A and σ_A have to solve the two equations (3. 19 and 3. 21), which are strictly speaking *pathwise* equations, since E has a stochastic volatility. Nevertheless Equation (3. 21) can be replaced by

$$\sigma_{E,t}^2 = \sigma_A^2 A_t^2 \partial_x^2 E(A_t, D_t).$$

Let us define

$$\begin{aligned} E'(A, D, \sigma_A) &= \partial_x E(A, D, \sigma_A) \\ &= \left[\left(\gamma - \frac{\delta}{r - \mu_D} \right) \left(\frac{1}{D\gamma} \right)^{\lambda(\sigma_A)} \lambda(\sigma_A) A^{\lambda(\sigma_A)-1} \right]. \end{aligned} \quad (3. 22)$$




FIGURE 3.2

Asset-Equity relation, Equation (3. 20), for parameter sets $(\delta, r, \gamma, \mu, \sigma_A)$ and $D = 1$: (-) solid (0.1, 0.05, 1, 0.0, 0.1), (-) dashed (0.1, 0.05, 1, 0.03, 0.1), (-.) dashed-dotted (0.0, 0.05, 1, 0.03, 0.1).

If we observe E at time t and know the estimate $\sigma_{E,t}$ for the equity volatility, then A and σ_A have to solve the equations

$$E = D \left[\frac{A}{D} - \frac{\delta}{r - \mu_D} - \left(\gamma - \frac{\delta}{r - \mu_D} \right) \left(\frac{A/D}{\gamma} \right)^{\lambda(\sigma_A)} \right] \quad (3. 23)$$

$$\sigma_{E,t} = \sigma_A A \left[\left(\gamma - \frac{\delta}{r - \mu_D} \right) \left(\frac{1}{D\gamma} \right)^{\lambda(\sigma_A)} \lambda(\sigma_A) A^{\lambda(\sigma_A)-1} \right] . \quad (3. 24)$$

As a further simplification it is often assumed that E locally evolves like a geometric Brownian motion, which leads to $\sigma_{E,t} = \sigma_E E$ for some σ_E .

In the implementation one usually starts with some $\sigma_A = \sigma_A^0$. For example, the equity volatility is used to generate two time series $(A_s)_{s \geq 0}$ and $(E_s)_{s \geq 0}$. Then, the volatility of E is estimated, and the parameter σ_A^1 is adjusted to a higher or lower level, trying to best match the estimated volatility of E with the observed equity volatility. One proceeds that way until the σ_E^n , implied by σ_A^n , is close to the observed σ_E . Observe also that the set of equations (3. 23) and (3. 24) can be generalized to any contingent claim approach for the asset values, once a functional relationship $E = E(A, D, \sigma_A, t)$ is specified between assets A , debt D , and equity E . Conceptually, they look like

$$E = E(A, D, \sigma_A) , \quad \sigma_E E = \sigma_A A E'(A, D, \sigma_A) .$$

This concludes are discussion of asset value models.

Chapter 4

The CreditRisk⁺ Model

In [Section 2.4.2](#) we already described the CreditRisk⁺ model as a Poissonian mixture with gamma-distributed random intensities for each sector. In this section we will explain CreditRisk⁺ in some greater detail. The justification for another and more exhaustive chapter on CreditRisk⁺ is its broad acceptance by many credit risk managing institutes. Even in the new Capital Accord (some references regarding the Basel II approach are Gordy [52], Wilde [126], and the IRB consultative document [103]), CreditRisk⁺ was originally applied for the calibration of the so-called *granularity adjustment* in the context of the *Internal Ratings-based Approach* (IRB) of regulatory capital risk weights. The popularity of CreditRisk⁺ has two major reasons:

- It seems easier to calibrate data to the model than is the case for multi-factor asset value models. Here we intentionally said “it seems” because from our point of view the calibration of bank-internal credit data to a multi-sector model is in general neither easier nor more difficult than the calibration of a multi-factor model on which an asset value model can be based.
- The second and maybe most important reason for the popularity of CreditRisk⁺ is its *closed-form* loss distribution. Using *probability generating functions*, the CreditRisk⁺ model offers (even in case of more than one sector) a full analytic description of the portfolio loss of any given credit portfolio. This enables users of CreditRisk⁺ to compute loss distributions in a quick and still “exact” manner. For many applications of credit risk models, this is a “nice-to-have” feature, e.g., in pricing or ABS structuring.

Before going into the details of the CreditRisk⁺ model, we like to present a quotation from the CreditRisk⁺ Technical Document [18] on page 8. There we find that

CreditRisk⁺ focuses on modeling and managing credit default risk.

In other words, CreditRisk⁺ helps to quantify the potential risk of defaults and resulting losses in terms of exposure in a given portfolio. Although it incorporates a *term structure of default rates* (more explicitly yearly *marginal* default rates) for implementing multi-year loss distributions (see [18], A5.2), it is not an appropriate choice if one is interested in a mark-to-market model of credit risk.

4.1 The Modeling Framework of CreditRisk⁺

Crucial in CreditRisk⁺ is the use of probability-generating functions¹. Recall that the generating function of a Poisson random variable L' with intensity λ is given by

$$G(z) = \sum_{k=0}^{\infty} \mathbb{P}[L' = k] z^k = e^{-\lambda} \sum_{k=0}^{\infty} \frac{\lambda^k}{k!} z^k = e^{\lambda(z-1)}. \quad (4.1)$$

In order to reduce the computational effort, CreditRisk⁺ groups the individual exposures of the obligors in a considered portfolio into *exposure bands*. This is done as follows:

Choose an exposure *unit* amount E . Analogously to Chapter 1, denote for any obligor i its *Expected Loss* by EL_i , its *Exposure At Default* by EAD_i , and its *Loss Given Default* by LGD_i . The exposure that is subject to be lost after an obligor's default is then

$$E_i = EAD_i \times LGD_i, \quad (4.2)$$

assuming a nonrandom LGD. The exposure ν_i respectively the Expected Loss ε_i of obligor i in multiples of the exposure unit E is given by

$$\nu_i = \frac{E_i}{E}, \quad \varepsilon_i = \frac{EL_i}{E}.$$

From this point on, CreditRisk⁺ “forgets” the exact exposures from the original portfolio and uses an approximation by means of exposure

¹In probability theory there are three concepts of translating a probability distribution into a functional context, namely the *Fourier transform*, the *Laplace transform* (which is in case of distributions on \mathbb{R}_+^d often more convenient), and the *probability-generating function* (often preferably used for distributions on \mathbb{Z}_+). The latter is defined by the function $z \mapsto \mathbb{E}[z^X]$ for a random variable X . Regarding basic properties of generating functions we refer to [53].

bands by rounding the exposures ν_i to the nearest integer number. In other words, every exposure E_i is replaced by the closest integer multiple of the unit exposure E . Already one can see that an appropriate choice of E is essential in order to end up at an approximation that is on one hand “close” enough to the original exposure distribution of the portfolio in order to obtain a loss distribution applicable to the original portfolio, and on the other hand efficient enough to really partition the portfolio into m_E exposure bands, such that m_E is significantly smaller than the original number of obligors m . An important “rule-of-thumb” for making sure that not too much precision is lost is to at least take care that the width of exposure bands is “small” compared to the average exposure size in the portfolio. Under this rule, *large* portfolios (containing many loans) should admit a good approximation by exposure bands in the described manner.

In the sequel we write $i \in [j]$ whenever obligor² i is placed in the exposure band j . After the exposure grouping process, we have a partition of the portfolio into m_E exposure bands, such that obligors in a common band $[j]$ have the common exposure $\nu_{[j]}E$, where $\nu_{[j]} \in \mathbb{N}_0$ is the integer multiple of E representing all obligors i with

$$\min\{|\nu_i - n| : n \in \mathbb{N}_0\} = |\nu_i - \nu_{[j]}|$$

$$\text{where } i = 1, \dots, m; i \in [j]; j = 1, \dots, m_E.$$

In cases where ν_i is an odd-integer multiple of 0.5, the above minimum is not uniquely defined. In such cases (which are obviously not very likely) one has to make a decision, if an up- or down-rounding would be appropriate.

Now let us discuss how to assign a default intensity to a given exposure band. Because CreditRisk⁺ plays in a Poissonian world, every obligor in the portfolio has its individual (one-year) *default intensity* λ_i , which can be calibrated from the obligor’s one-year default probability DP_i by application of (2. 12),

$$\lambda_i = -\log(1 - DP_i) \quad (i = 1, \dots, m). \quad (4. 3)$$

²Here we make the simplifying assumption that the number of loans in the portfolio equals the number of obligors involved. This can be achieved by aggregating different loans of a single obligor into one loan. Usually the DP, EAD, and LGD of such an aggregated loan are exposure-weighted average numbers.

Because the expectation of $L'_i \sim Pois(\lambda_i)$ is $\mathbb{E}[L'_i] = \lambda_i$, the *expected number of defaults* in exposure band $[j]$ (using the additivity of expectations) is given by

$$\lambda_{[j]} = \sum_{i \in [j]} \lambda_i . \tag{4. 4}$$

The *Expected Loss* in band $[j]$ will be denoted by $\varepsilon_{[j]}$ and is calculated by multiplying the expected number of defaults in band $[j]$ with the band's exposure,

$$\varepsilon_{[j]} = \lambda_{[j]} \nu_{[j]} . \tag{4. 5}$$

Here, the CreditRisk⁺ Technical Document suggests making an *adjustment of the default intensities* λ_i (which so far have not been affected by the exposure band approximation process) in order to preserve the original value of the obligor's Expected Losses. This could be done by defining an adjustment factor γ_i for every obligor i by

$$\gamma_i = \frac{E_i}{\nu_{[j]} E} \quad (i \in [j], j = 1, \dots, m_E) . \tag{4. 6}$$

Replacing for every obligor i the original default intensity λ_i by $\gamma_i \lambda_i$ with γ_i as defined in (4. 6) preserves the original ELs after approximating the portfolio's exposure distribution by a partition into exposure bands. In the following we assume without loss of generality that the default intensities λ_i already include the adjustment (4. 6). From (4. 4) respectively (4. 5) it is straightforward to write down the *portfolio's expected number of default events* (respectively the *portfolio's overall default intensity*), namely

$$\lambda_{PF} = \sum_{j=1}^{m_E} \lambda_{[j]} = \sum_{j=1}^{m_E} \frac{\varepsilon_{[j]}}{\nu_{[j]}} . \tag{4. 7}$$

After these preparations we are now ready to describe the construction of the CreditRisk⁺ loss distribution. We will proceed in two steps, starting with a portfolio of independent obligors and then mixing the involved Poisson distributions by means of a sector model as indicated in [Section 2.4.2](#).

4.2 Construction Step 1: Independent Obligors

We begin with a portfolio of m independent obligors whose default risk is modeled by Poisson variables L_i . As already mentioned in [Section 2.2.1](#), Poisson models allow for multiple defaults of a single obligor. This is an unpleasant, but due to the small occurrence probability, mostly ignored feature of all Poisson approaches to default risk.

Involving the (nonrandom) exposures E_i as defined in (4. 2), we obtain loss variables

$$E_i L'_i \quad \text{where} \quad L'_1 \sim Pois(\lambda_1) , \dots , L'_m \sim Pois(\lambda_m) \quad (4. 8)$$

are independent Poisson random variables. Grouping the individual exposures E_i into exposure bands $[j]$ and assuming the intensities λ_i to incorporate the adjustments by the factors γ_i as described in the introduction, we obtain new loss variables $\nu_{[j]} L'_i$, where losses are measured in multiples of the exposure unit E . Because obligors are assumed to be independent, the number of defaults L' in the portfolio respectively $L'_{[j]}$ in exposure band j also follow a Poisson distribution, because the convolution of independent Poisson variables yields a Poisson distribution. We obtain

$$L'_{[j]} = \sum_{i \in [j]} L'_i \sim Pois(\lambda_{[j]}) , \quad \lambda_{[j]} = \sum_{i \in [j]} \lambda_i , \quad (4. 9)$$

for the number of defaults in exposure band $[j]$, $j = 1, \dots, m_E$, and

$$L' = \sum_{j=1}^{m_E} \sum_{i \in [j]} L'_i \sim Pois\left(\sum_{j=1}^{m_E} \lambda_{[j]}\right) = Pois(\lambda_{PF}) \quad (4. 10)$$

(see (4. 7)), for the portfolio's number of defaults. The corresponding losses (counted in multiples of the exposure unit E) are given by

$$\tilde{L}'_{[j]} = \nu_{[j]} L'_{[j]} \quad \text{respectively} \quad \tilde{L}' = \sum_{j=1}^{m_E} \nu_{[j]} L'_{[j]} = \sum_{j=1}^{m_E} \tilde{L}'_{[j]} . \quad (4. 11)$$

Due to grouping the exposures $\nu_{[j]} \in \mathbb{N}_0$ together, we can now conveniently describe the portfolio loss by the probability-generating func-

tion of the random variable \tilde{L}' defined in (4. 11), applying the convolution theorem³ for generating functions,

$$\begin{aligned}
 G_{\tilde{L}'}(z) &= \prod_{j=1}^{m_E} G_{\tilde{L}'_{[j]}} = \prod_{j=1}^{m_E} \sum_{k=0}^{\infty} \mathbb{P}[\tilde{L}'_{[j]} = \nu_{[j]}k] z^{\nu_{[j]}k} & (4. 12) \\
 &= \prod_{j=1}^{m_E} \sum_{k=0}^{\infty} \mathbb{P}[L'_{[j]} = k] z^{\nu_{[j]}k} = \prod_{j=1}^{m_E} \sum_{k=0}^{\infty} e^{-\lambda_{[j]}} \frac{\lambda_{[j]}^k}{k!} z^{\nu_{[j]}k} \\
 &= \prod_{j=1}^{m_E} e^{-\lambda_{[j]} + \lambda_{[j]} z^{\nu_{[j]}}} = \exp\left(\sum_{j=1}^{m_E} \lambda_{[j]} (z^{\nu_{[j]}} - 1)\right).
 \end{aligned}$$

So far we assumed independence among obligors and were rewarded by the nice closed formula (4. 12) for the generating function of the portfolio loss. In the next section we drop the independence assumption, but the nice feature of CreditRisk⁺ is that, nevertheless, it yields a closed-form loss distribution, even in the case of correlated defaults.

4.3 Construction Step 2: Sector Model

A key concept of CreditRisk⁺ is *sector analysis*. The rationale underlying sector analysis is that the volatility of the default intensity of obligors can be related to the volatility of certain underlying factors incorporating a common systematic source of credit risk. Associated with every such background factor is a so-called *sector*, such that every obligor i admits a breakdown into *sector weights* $w_{is} \geq 0$, $\sum_{s=1}^{m_S} w_{is} = 1$, expressing for every $s = 1, \dots, m_S$ that sector s contributes with a fraction w_{is} to the default intensity of obligor i . Here m_S denotes the number of involved sectors. Obviously the calibration of sectors and sector weights is the crucial challenge in CreditRisk⁺. For example, sectors could be constructed w.r.t. industries, countries, or rating classes.

³For independent variables, the generating function of their convolution equals the product of the corresponding single generating functions.

In order to approach the sector model of CreditRisk⁺ we rewrite Equation (4. 12):

$$\begin{aligned}
 G_{\tilde{L}'}(z) &= \exp\left(\sum_{j=1}^{m_E} \lambda_{[j]}(z^{\nu_{[j]}} - 1)\right) & (4. 13) \\
 &= \exp\left(\lambda_{PF} \left(\sum_{j=1}^{m_E} \frac{\lambda_{[j]}}{\lambda_{PF}} z^{\nu_{[j]}} - 1\right)\right),
 \end{aligned}$$

where λ_{PF} is defined as in (4. 7). Defining functions

$$G_{L'}(z) = e^{\lambda_{PF}(z-1)} \quad \text{and} \quad G_N(z) = \sum_{j=1}^{m_E} \frac{\lambda_{[j]}}{\lambda_{PF}} z^{\nu_{[j]}} , \quad (4. 14)$$

we see that the generating function of the portfolio loss variable \tilde{L}' can be written as

$$G_{\tilde{L}'}(z) = G_{L'} \circ G_N(z) = e^{\lambda_{PF}(G_N(z)-1)} . \quad (4. 15)$$

Therefore, the portfolio loss \tilde{L}' has a so-called *compound* distribution, essentially meaning that the randomness inherent in the portfolio loss is due to the compound effect of two independent sources of randomness. The first source of randomness arises from the uncertainty regarding the number of defaults in the portfolio, captured by the Poisson random variable L' with intensity λ_{PF} defined in (4. 10). The function $G_{L'}(z)$ is the generating function of L' ; recall (4. 1). The second source of randomness is due to the uncertainty about the exposure bands affected by the L' defaults. The function $G_N(z)$ is the generating function of a random variable N taking values in $\{\nu_{[1]}, \dots, \nu_{[m_E]}\}$ with distribution

$$\mathbb{P}[N = \nu_{[j]}] = \frac{\lambda_{[j]}}{\lambda_{PF}} \quad (j = 1, \dots, m_E). \quad (4. 16)$$

For some more background on compound⁴ distributions, refer to the literature. For example in [53] the reader will find theory as well as some

⁴Compound distributions arise very naturally as follows: Assume X_0, X_1, X_2, \dots be i.i.d. random variables with generating function G_X . Let $N \in \mathbb{N}_0$ be a random variable, e.g., $N \sim Pois(\lambda)$, independent of the sequence $(X_i)_{i \geq 0}$. Denote the generating function of N by G_N . Then, the generating function of $X_1 + \dots + X_N$ is given by $G = G_N \circ G_X$. In the case where the distribution of N is degenerate, e.g., $\mathbb{P}[N = n] = 1$, we obtain $G_N(z) = z^n$ and therefore $G(z) = [G_X(z)]^n$, confirming the convolution theorem for generating functions in its most basic form.

interesting examples. Later on we will obtain the generating function of sector losses in form of an equation that, conditional on the sector's default rate, replicates Equation (4. 15).

Let us assume that we have parametrized our portfolio by means of m_S sectors. CreditRisk⁺ assumes that a gamma-distributed random variable

$$\Lambda^{(s)} \sim \Gamma(\alpha_s, \beta_s) \quad (s = 1, \dots, m_S)$$

is assigned to every sector; see Figure 2.2 for an illustration of gamma densities. The number of defaults in any sector s follows a gamma-mixed Poisson distribution with random intensity $\Lambda^{(s)}$; see also Section 2.2.2. Hereby it is always assumed that the *sector variables* $\Lambda^{(1)}, \dots, \Lambda^{(m_S)}$ are independent.

For a calibration of $\Lambda^{(s)}$ recall from (2. 38) that the first and second moment of $\Lambda^{(s)}$ are

$$\mathbb{E}[\Lambda^{(s)}] = \alpha_s \beta_s, \quad \mathbb{V}[\Lambda^{(s)}] = \alpha_s \beta_s^2 . \quad (4. 17)$$

We denote the expectation of the random intensity $\Lambda^{(s)}$ by $\lambda_{(s)}$. The volatility of $\Lambda^{(s)}$ is denoted by $\sigma_{(s)}$. Altogether we have from (4. 17)

$$\lambda_{(s)} = \alpha_s \beta_s, \quad \sigma_{(s)} = \sqrt{\alpha_s \beta_s^2} . \quad (4. 18)$$

Knowing the values of $\lambda_{(s)}$ and $\sigma_{(s)}$ determines the parameters α_s and β_s of the sector variable $\Lambda^{(s)}$.

For every sector we now follow the approach that has taken us to Equation (4. 15). More explicitly, we first find the generating function of the number of defaults in sector s , then obtain the generating function for the distribution of default events among the exposures in sector s , and finally get the portfolio-loss-generating function as the product⁵ of the compound sector-generating functions.

4.3.1 Sector Default Distribution

Fix a sector s . The defaults in all sectors are gamma-mixed Poisson. Therefore, conditional on $\Lambda^{(s)} = \theta_s$ the sector's conditional generating function is given by (4. 1),

$$G_s |_{\Lambda^{(s)}=\theta_s}(z) = e^{\theta_s(z-1)} . \quad (4. 19)$$

⁵Recall that we assumed independence of sector variables.

The unconditional generating function also is explicitly known, because fortunately it is a standard fact from elementary statistics that gamma-mixed Poisson variables follow a *negative binomial distribution* (see, e.g., [109], 8.6.1.) The negative binomial distribution usually is taken as a suitable model for a counting variable when it is known that the variance of the counts is larger than the mean. Recalling our discussion in Section 2.2.2 we know that the dispersion of Poisson variables is equal to 1 due to the agreement of mean and variance. Mixing Poisson variables with gamma distributions will always result in a distribution with a conditional dispersion of 1 but *unconditionally overdispersed*.

At this point we make a brief detour in order to provide the reader with some background knowledge on negative binomial distributions. There are two major reasons justifying this. First, the negative binomial distribution is probably not as well known to all readers as the (standard) binomial distribution. Second, the negative binomial distribution is differently defined in different textbooks. We therefore believe that some clarification about our view might help to avoid misunderstandings.

One approach to the negative binomial distribution (see, e.g., [53]) is as follows: Start with a sequence of independent Bernoulli default indicators $X_i \sim B(1; p)$. Let T be the *waiting time* until the first default occurs, $T = \min\{i \in \mathbb{N} \mid X_i = 1\}$. We have

$$\begin{aligned} \mathbb{P}[T = k] &= \mathbb{P}[T > k - 1] - \mathbb{P}[T > k] \\ &= (1 - p)^{k-1} - (1 - p)^k = p(1 - p)^{k-1}. \end{aligned}$$

Therefore, T has a *geometric distribution*. If more generally we ask for the waiting time T_q until the q -th default occurs, then we obtain the *negative binomial distribution with parameters p and q* . The mass function of T_q obviously is given by

$$\mathbb{P}[T_q = k] = \binom{k-1}{q-1} p^q (1-p)^{k-q} \quad (k \geq q). \quad (4. 20)$$

For $q = 1$ the negative binomial and the geometric distributions agree. Moreover,

$$T_q = \sum_{i=1}^q T'_i \quad \text{where} \quad T'_1 = T_1, \quad T'_i = T_i - T_{i-1} \quad \text{for } i = 2, \dots, q, \quad (4. 21)$$

where T'_1, \dots, T'_q are *independent* geometric variables with parameter p . For $i \geq 2$ the variable T'_i is the waiting time until the next default following the $(i - 1)$ -th default. Because the mean and the variance of a geometric random variable T with parameter p are $\mathbb{E}[T] = 1/p$ respectively $\mathbb{V}[T] = (1 - p)/p^2$, (4. 21) yields

$$\mathbb{E}[T_q] = \frac{q}{p} \quad \text{and} \quad \mathbb{V}[T_q] = \frac{q(1 - p)}{p^2} . \quad (4. 22)$$

The generating function of an exponential variable T with parameter p is

$$G_T(z) = \sum_{k=1}^{\infty} p(1 - p)^{k-1} z^k = \frac{pz}{1 - (1 - p)z} \quad (|z| < 1/(1 - p)).$$

Therefore the convolution theorem for generating functions immediately implies

$$G_{T_q}(z) = \left(\frac{pz}{1 - (1 - p)z} \right)^q \quad (|z| < 1/(1 - p)). \quad (4. 23)$$

Application of the relation $\binom{x}{k} = (-1)^k \binom{k-x-1}{k}$ ($x \in \mathbb{R}$, $k \in \mathbb{N}_0$) and the symmetry property $\binom{n}{m} = \binom{n}{n-m}$ yields

$$\mathbb{P}[T_q = k] = \binom{-q}{k - q} p^q (p - 1)^{k-q} ,$$

which explains the name *negative binomial distribution*.

So far this is what many authors consider to be a negative binomial distribution. Now, some people consider it a technical disadvantage that the (according to our discussion above very naturally arising) negative binomial distribution ranges in $\{k \in \mathbb{N} \mid k \geq q\}$. For reasons also applying to the situation in CreditRisk⁺ one would rather like to see T_q ranging in \mathbb{N}_0 . We can adopt this view by replacing T_q by $\tilde{T}_q = T_q - q$, again applying the symmetry property $\binom{n}{m} = \binom{n}{n-m}$, and substituting $n = k - q$ in Equation (4. 20):

$$\begin{aligned} \mathbb{P}[\tilde{T}_q = n] &= \mathbb{P}[T_q = n + q] \\ &= \binom{n + q - 1}{n} p^q (1 - p)^n \quad (n \geq 0). \end{aligned} \quad (4. 24)$$

The variable \tilde{T}_q obviously describes the number of survivals until the q -th default has occurred.

Because CreditRisk⁺ requires it in this way, we from now on mean by a negative binomial distribution with parameters q and p the distribution of \tilde{T}_q defined by (4. 24). It is well known (see, e.g., [109], 8.6.1) that any $\Gamma(\alpha, \beta)$ -mixed Poisson variable L' follows a negative binomial distribution with parameters α and $1/(1 + \beta)$. This concludes our detour and we return to the actual topic of this section.

The conditional distribution of the sector defaults is given by (4. 19). The mixing variable is $\Lambda^{(s)} \sim \Gamma(\alpha_s, \beta_s)$. According to our discussion above, the unconditional distribution of sector defaults (denoted by $L'_{(s)}$) is negative binomial with parameters α_s and $1/(1 + \beta_s)$; in short: $L'_{(s)} \sim NB(\alpha_s, 1/(1 + \beta_s))$. We can now easily obtain the unconditional generating function $G_{L'_{(s)}}(z) = G_s(z)$ of the sector defaults by evaluating Formula (4. 23) with T_q replaced by $\tilde{T}_q = T_q - q$ and taking the parametrization $q = \alpha_s$ and $p = 1/(1 + \beta_s)$ into account. Replacing T_q by $\tilde{T}_q = T_q - q$ changes (4. 23) to

$$G_{\tilde{T}_q}(z) = \left(\frac{pz}{1 - (1 - p)z} \right)^q \frac{1}{z^q} = \left(\frac{p}{1 - (1 - p)z} \right)^q.$$

Inserting $q = \alpha_s$ and $p = 1/(1 + \beta_s)$ finally yields

$$\begin{aligned} G_{L'_{(s)}}(z) &= \int_0^\infty [G_{L'_{(s)}} |_{\Lambda^{(s)}=\theta_s}](z) \gamma_{\alpha_s, \beta_s}(\theta_s) d\theta_s & (4. 25) \\ &= \left(\frac{1 - \frac{\beta_s}{1 + \beta_s}}{1 - \frac{\beta_s}{1 + \beta_s} z} \right)^{\alpha_s}, \end{aligned}$$

where $\gamma_{\alpha_s, \beta_s}$ denotes the density of $\Gamma(\alpha_s, \beta_s)$ and α_s, β_s are calibrated to the sector by means of (4. 18). We included the integral in the center of (4. 25) in order to explicitly mention the link to [Section 2.2.1](#).

Formula (4. 25) can be found in the CreditRisk⁺ Technical Document [18] (A8.3, Equation (55)). The probability mass function of $L'_{(s)}$ follows from (4. 24),

$$\mathbb{P}[L'_{(s)} = n] = \binom{n + \alpha_s - 1}{n} \left(1 - \frac{\beta_s}{1 + \beta_s} \right)^{\alpha_s} \left(\frac{\beta_s}{1 + \beta_s} \right)^n. \quad (4. 26)$$

The first and second moments of the distribution of defaults in sector s directly follow from the general results on Poisson mixtures; see (2.

15) in [Section 2.2](#). They depend on the mixture distribution only and are given by

$$\mathbb{E}[L'_{(s)}] = \mathbb{E}[\Lambda^{(s)}] = \alpha_s \beta_s \quad \text{and} \quad (4. 27)$$

$$\mathbb{V}[L'_{(s)}] = \mathbb{V}[\Lambda^{(s)}] + \mathbb{E}[\Lambda^{(s)}] = \alpha_s \beta_s (1 + \beta_s)$$

(see also (4. 17)), hereby confirming our previous remark that the unconditional distribution of sector defaults is *overdispersed*. In (4. 28) we see that we always have $\beta_s \geq 0$ and that $\beta_s > 0$ if and only if the volatility of the sector's default intensity does not vanish to zero. [Figure 2.7](#) in [Section 2.5.2](#) graphically illustrates (4. 26).

Alternatively, the first and second moments of $L'_{(s)}$ could have been calculated by application of (4. 22), taking the shift $T_q \rightarrow T_q - q$ and the parametrization of q and p into account. It is a straightforward calculation to show that the result of such a calculation agrees with the findings in (4. 27).

4.3.2 Sector Compound Distribution

As a preparation for the compound approach on the sector level we begin this section with a remark regarding the calibration of α_s and β_s . Solving Equations (4. 18) for α_s and β_s gives us the mixing parameters in terms of the sector parameters $\lambda_{(s)}$ and $\sigma_{(s)}$:

$$\alpha_s = \frac{\lambda_{(s)}^2}{\sigma_{(s)}^2} \quad \text{and} \quad \beta_s = \frac{\sigma_{(s)}^2}{\lambda_{(s)}}. \quad (4. 28)$$

Most often $\lambda_{(s)}$ and $\sigma_{(s)}$ will be calibrated from obligor data. As already mentioned before, in the sector model every obligor i admits a breakdown into sector weights w_{is} such that

$$\sum_{s=1}^{m_S} w_{is} = 1 \quad (w_{is} \geq 0; i = 1, \dots, m).$$

Moreover, any obligor i admits a random default intensity defined by

$$\Lambda_i = \sum_{s=1}^{m_S} w_{is} \lambda_i \frac{\Lambda^{(s)}}{\lambda_{(s)}} \quad (i = 1, \dots, m). \quad (4. 29)$$

The expected intensity then obviously equals λ_i , which is consistent with the case of independent obligors where λ_i denoted the nonrandom default intensity of obligor i . The expected intensities λ_i can be

calibrated to one-year default probabilities by application of Formula (4. 3). Due to the additivity of expectations it is then very natural to define the expected sector intensity $\lambda_{(s)}$ by

$$\lambda_{(s)} = \sum_{i=1}^m w_{is}\lambda_i = \sum_{j=1}^{m_E} \sum_{i \in [j]} w_{is}\lambda_i, \quad (4. 30)$$

where the right side expresses the grouping into exposure bands. Note that an exposure band j takes part in sector s if and only if there exists some obligor $i \in [j]$ such that $w_{is} > 0$. The sector volatility $\sigma_{(s)}$ can be either calibrated from empirical data related to the meaning of sector s or calculated from the single obligor's default intensity volatilities. An example for the first case would be an industry sector, where the volatility of a historical time series of insolvency quotes for that particular industry could be taken as a proxy for $\sigma_{(s)}$. An example for the latter case would be a portfolio where, in addition to the default rate, the default rate volatility is also known for every obligor. Such estimates usually depend on the creditworthiness of obligors. For example, denoting the default rate volatility for obligor i by σ_i and assuming that the sectors perform a *partition* of the portfolio's set of obligors (more explicitly: $w_{is} = 1$ for a unique sector $s = s(i)$ for every obligor i) one obtains from (4. 29) and (4. 30) for every sector s the following identity:

$$\sum_{i:w_{is}=1} \sigma_i = \sum_{i:w_{is}=1} \sqrt{\mathbb{V}\left[w_{is}\lambda_i \frac{\Lambda^{(s)}}{\lambda_{(s)}}\right]} = \sum_{i:w_{is}=1} w_{is}\lambda_i \frac{\sigma_{(s)}}{\lambda_{(s)}} = \sigma_{(s)},$$

where the sum takes all obligors i in sector s into account. So in this particular example, the volatility of the sector default intensity can be directly estimated from the volatility of the default intensities of obligors collected into that sector. The calibration of a sector variable $\Lambda^{(s)}$ can then be finalized by applying (4. 28).

For the general case where obligors are allowed to be influenced by more than one sector, the CreditRisk⁺ Technical Document [18] (A12.2) suggests an analogous approach by estimating the sector volatility $\sigma_{(s)}$ by the weighted contribution of the default rate volatilities of obligors influenced by the sector, namely $\hat{\sigma}_{(s)} = \sum_{i=1}^m w_{is}\sigma_i$. Again note that only obligors i with $w_{is} > 0$ contribute to sector s .

Based on our calculations above we can now just follow the lines of (4. 14), (4. 15), and (4. 16). Analogously to (4. 16) we first of all

define a random variable N_s by

$$\mathbb{P}[N_s = \nu_{[j]}] = \frac{1}{\lambda^{(s)}} \sum_{i \in [j]} w_{is} \lambda_i \quad (j = 1, \dots, m_E). \quad (4. 31)$$

Equation (4. 30) shows that (4. 16) really defines a probability distribution on the set of exposures $\{\nu_{[1]}, \dots, \nu_{[m_E]}\}$. The generating function of N_s is given by a polynomial analogous to (4. 14),

$$G_{N_s}(z) = \sum_{j=1}^{m_E} \left(\frac{1}{\lambda^{(s)}} \sum_{i \in [j]} w_{is} \lambda_i \right) z^{\nu_{[j]}}. \quad (4. 32)$$

Instead of $G_{L'}$ as in (4. 14) we now use the generating function $G_{L'_{(s)}}$ of the sector defaults as described in (4. 25). Because the generating function G_{N_s} does not depend on realizations of the random intensity $\Lambda^{(s)}$, the arguments leading to Formula (4. 25) are not affected when replacing the variable z in (4. 25) by $G_{N_s}(z)$. We therefore obtain the compound generating function of the distribution of losses $\tilde{L}'_{(s)}$ in sector s by writing

$$\begin{aligned} G_{\tilde{L}'_{(s)}}(z) &= G_{L'_{(s)}} \circ G_{N_s}(z) = \left(\frac{1 - \frac{\beta_s}{1+\beta_s}}{1 - \frac{\beta_s}{1+\beta_s} G_{N_s}(z)} \right)^{\alpha_s} \\ &= \left(\frac{1 - \frac{\beta_s}{1+\beta_s}}{1 - \frac{\beta_s}{1+\beta_s} \frac{1}{\lambda^{(s)}} \sum_{j=1}^{m_E} \sum_{i \in [j]} w_{is} \lambda_i z^{\nu_{[j]}}} \right)^{\alpha_s}. \end{aligned} \quad (4. 33)$$

Note that this is the same two-step randomness we previously derived in the case of independent obligors leading to Formula (4. 15). The bridge between the independent case and (4. 33) is just a simple conditioning argument. Conditional on a sector intensity realization θ_s , the conditional compound probability-generating function analogous to (4. 15) is given by

$$[G_{\tilde{L}'_{(s)}} |_{\Lambda^{(s)}=\theta_s}](z) = e^{\theta_s(G_{N_s}(z)-1)}. \quad (4. 34)$$

Integration of the right side w.r.t. the mixing gamma distribution gives (4. 33), taking into account that $G_{N_s}(z)$ does not depend on the integration variable θ_s . Therefore the integral can be calculated by means of exactly the same argument as in (4. 25).

4.3.3 Sector Convolution

The portfolio loss $L' = L'_1 + \dots + L'_m$ is a mixed Poisson variable with random intensity $\Lambda = \Lambda_1 + \dots + \Lambda_m$. Grouped into sectors, the intensity Λ of L' can also be written as the sum of sector intensities,

$$\Lambda = \Lambda^{(1)} + \dots + \Lambda^{(m_S)}.$$

This follows from Formulas (4. 29) and (4. 30). Because sectors are assumed to be independent, the distribution of defaults in the portfolio is the convolution of the sector's default distributions. Therefore, due to (4. 25) the generating function of L' is given by

$$G_{L'}(z) = \prod_{s=1}^{m_S} \left(\frac{1 - \frac{\beta_s}{1+\beta_s}}{1 - \frac{\beta_s}{1+\beta_s} z} \right)^{\alpha_s}. \quad (4. 35)$$

The generating function of the portfolio losses is determined by the convolution of compound sector distributions as elaborated in (4. 33),

$$G_{\tilde{L}'}(z) = \prod_{s=1}^{m_S} \left(\frac{1 - \frac{\beta_s}{1+\beta_s}}{1 - \frac{\beta_s}{1+\beta_s} \frac{1}{\lambda_{(s)}} \sum_{j=1}^{m_E} \sum_{i \in [j]} w_{is} \lambda_i z^{\nu_{[j]}}} \right)^{\alpha_s}. \quad (4. 36)$$

So we see that in CreditRisk⁺ the portfolio loss distribution can be described in an analytical manner by means of a closed-form generating function. Remarkable is the fact that this nice property even holds in the most general case of a sector model where complex dependence structures are allowed. In the general sector model of CreditRisk⁺ leading to Formula (4. 36), obligors i_1 and i_2 will be correlated if and only if there exists at least one sector s such that $w_{i_1 s} > 0$ and $w_{i_2 s} > 0$.

Because probability distributions and generating functions are uniquely associated with each other, Formula (4. 36) allows for very quick computations of loss distributions. In the CreditRisk⁺ Technical Document [18] the reader will find a calculation scheme for the loss amount distribution of portfolios based on a certain *recurrence relation*. The formulas are slightly technical but allow the calculation of loss distributions in fractions of a second, using a standard personal computer with some suitable numerical software. Alternatively a demo version of CreditRisk⁺ in form of a spreadsheet can be downloaded from www.csfb.com/creditrisk. For readers interested in working with CreditRisk⁺ we recommend playing with the spreadsheet in order to

develop some good intuition regarding the impact of driving variables and parameters. Anyway, the recurrence formulas in [18] can be easily implemented so that everyone is free to program his or her “individual” version of CreditRisk⁺.

There is much more to say about CreditRisk⁺, but due to the introductory character of this book we will not go any further. The Technical Document [18] contains some more detailed information on the management of credit portfolios, the calibration of the model, and the technical implementation. For example, [18], A13, addresses the question of risk contributions (see also [Chapter 5](#)) and the role of correlations in CreditRisk⁺. Risk contributions in CreditRisk⁺ are also extensively studied in Tasche [121]. In [18], A12.3, the introduction of a sector for incorporating *specific risk* is discussed. As a last remark we should mention that because the sector distributions are Poisson mixtures, the general results from [Section 2.2](#) can also be applied.

Chapter 5

Alternative Risk Measures and Capital Allocation

The definition of economic capital as introduced in [Chapter 1](#) appears fully satisfactory at first glance. Starting with the path-breaking paper by Artzner et al. [5], several studies revealed a number of methodological weaknesses in the VaR concept by drawing up a catalogue of mathematical and material attributes that a risk measure should fulfill, and proving that the VaR concept only partly fulfills them. Risk measures that satisfy these axioms are called *coherent*. Before we describe their basic features in the next section we briefly reiterate the main notations (cf. [Chapters 1](#) and [2](#)):

The portfolio loss variable (compare Equation (2. 51)) is given by

$$L = \sum_{i=1}^m w_i \eta_i L_i,$$

where

$$w_i = \frac{\text{EAD}_i}{\sum_{i=1}^m \text{EAD}_i}$$

are the exposure weights, so all portfolio quantities are calculated in percent of the portfolio's total exposure. The severity is abbreviated to $\eta_i = \text{LGD}_i$; $L_i = \mathbf{1}_{D_i}$ are Bernoulli random variables with default events D_i . Accordingly, the default probability (e.g., for a one-year horizon) for obligor i is given by

$$\text{DP}_i = \mathbb{P}(D_i) .$$

Default correlations are denoted as

$$\rho_{ij} = \text{Corr}[L_i, L_j] = \text{Corr}[\mathbf{1}_{D_i}, \mathbf{1}_{D_j}] .$$

In this chapter we assume a deterministic severity, η_i , but, if severities are assumed to be independent of defaults, then it is straightforward to

extend the following to random severities. For notational convenience we denote by

$$\hat{L}_i := \eta_i L_i$$

the loss variable of some obligor i .

5.1 Coherent Risk Measures and Conditional Shortfall

Denote by $L^\infty = L^\infty(\Omega, \mathcal{F}, \mathbb{P})$ the space of bounded real random variables, defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The definition of a *coherent risk measure* as suggested by Artzner et al. [4,5] can then be stated in the following way:

5.1.1 Definition A mapping $\gamma : L^\infty(\Omega, \mathcal{F}, \mathbb{P}) \rightarrow \mathbb{R}$ is called a *coherent risk measure* if the following properties hold:

- (i) **Subadditivity** $\forall X, Y \in L^\infty: \gamma(X + Y) \leq \gamma(X) + \gamma(Y)$
- (ii) **Monotonicity** $\forall X, Y \in L^\infty$ with $X \leq Y$ a.s.: $\gamma(X) \leq \gamma(Y)$
- (iii) **Positive homogeneity** $\forall \lambda > 0, \forall X \in L^\infty: \gamma(\lambda X) = \lambda \gamma(X)$
- (iv) **Translation invariance** $\forall x \in \mathbb{R}, \forall X \in L^\infty:$
 $\gamma(X + x) = \gamma(X) + x.$

Note that the definition here slightly differs from the original set of axioms as they were introduced by Artzner et al. [4,5]. Our definition equals the one given by Frey and McNeil in [46], because we want to think about X in terms of a portfolio loss and about $\gamma(X)$ as the amount of capital required as a cushion against the loss X , according to the credit management policy of the bank. In the original approach by Artzner et al., X was interpreted as the future value of the considered portfolio. Let us now briefly explain the four axioms in an intuitive manner:

Subadditivity: Axiom (i) reflects the fact that due to diversification effects the risk inherent in the union of two portfolios should be less than the sum of risks of the two portfolios considered separately. We will later see that quantiles are not subadditive in general, such that the economic capital (EC) as introduced in [Chapter 1](#) turns out to be not coherent.

Monotonicity: Let us say we are considering portfolios A and B with losses X_A and X_B . If almost surely the losses of portfolio A are lower than the losses of portfolio B, i.e., $X_A \leq X_B$ a.s., then the required risk capital $\gamma(X_A)$ for portfolio A should be less than the required risk capital $\gamma(X_B)$ of portfolio B. Seen from this perspective, monotonicity is a very natural property.

Homogeneity: Axiom (iii) can best be illustrated by means of the following example. Consider a credit portfolio with loss X and scale all exposures by a factor λ . Then, of course, the loss X changes to a scaled loss λX . Accordingly, the originally required risk capital $\gamma(X)$ will also change to $\lambda\gamma(X)$.

Translation invariance: If x is some capital which will be lost/gained on a portfolio with certainty at the considered horizon, then the risk capital required for covering losses in this portfolio can be increased/reduced accordingly. Translation invariance implies the natural property $\gamma(X - \gamma(X)) = 0$ for every loss $X \in L^\infty$.

Typical risk measures discussed by Artzner et al. are the *value-at-risk* and *expected shortfall capital*, which will be briefly discussed in the sequel.

Value-at-Risk Value-at-risk (VaR) has already been mentioned in [Section 1.2.1](#) as a synonymous name for EC. Here, VaR will be defined for a probability measure \mathbb{P} and some confidence level α as the α -quantile of a loss random variable X ,

$$\text{VaR}_\alpha(X) = \inf\{x \geq 0 \mid \mathbb{P}[X \leq x] \geq \alpha\} .$$

VaR as a risk measure defined on L^∞ is

- translation invariant, because shifting a loss distribution by a fixed amount will shift the quantile accordingly,
- positively homogeneous, because scaling a loss variable will scale the quantile accordingly,
- monotone, because quantiles preserve monotonicity, but
- not subadditive, as the following example will show.

Because VaR is not subadditive, it is *not coherent*. Now let us give a simple example showing that VaR is not subadditive.

Consider two independent loans, represented by two loss indicator variables $\mathbf{1}_{D_A}, \mathbf{1}_{D_B} \sim B(1; p)$ with, e.g., $0.006 \leq p < 0.01$. Assume LGDs equal to 100% and exposures equal to 1. Define two portfolios A and B, each consisting of one single of the above introduced loans. Then, for the portfolio losses $X_A = \mathbf{1}_{D_A}$ and $X_B = \mathbf{1}_{D_B}$ we have

$$\text{VaR}_{99\%}(X_A) = \text{VaR}_{99\%}(X_B) = 0 ,$$

Now consider a portfolio C defined as the union of portfolios A and B, and denote by $X_C = X_A + X_B$ the corresponding portfolio loss. Then

$$\mathbb{P}[X_C = 0] = (1 - p)^2 < 99\% .$$

Therefore, $\text{VaR}_{99\%}(X_C) > 0$, so that

$$\text{VaR}_{99\%}(X_A + X_B) > \text{VaR}_{99\%}(X_A) + \text{VaR}_{99\%}(X_B) .$$

This shows that in general VaR is not a coherent risk measure.

Expected Shortfall The *tail conditional expectation*, or *expected shortfall*, w.r.t. a confidence level α is defined as

$$\text{TCE}_\alpha(X) = \mathbb{E}[X \mid X \geq \text{VaR}_\alpha(X)] .$$

In the literature one can find several slightly different versions of TCE definitions. Tasche [122] showed that expected shortfall to a great extent enjoys the coherence properties; for further reading see also [5,66]. For example, expected shortfall is coherent when restricted to loss variables X with continuous distribution function¹; see [122]. Therefore, TCE provides a reasonable alternative to the classical VaR measure. [Figure 5.1](#) illustrates the definition of expected shortfall capital. From an

¹In [1], Acerbi and Tasche introduce a so-called *generalized expected shortfall* measure, defined by

$$\text{GES}_\alpha = \frac{1}{1 - \alpha} \left(\mathbb{E}[X \mid X \geq \text{VaR}_\alpha(X)] + \text{VaR}_\alpha(X)(1 - \alpha - \mathbb{P}[X \geq \text{VaR}_\alpha(X)]) \right) .$$

If the distribution of X is continuous, the second term in the definition of GES_α vanishes, and the generalized expected shortfall coincides with (normalized) expected shortfall in this case. It is shown in [1] that generalized expected shortfall is a coherent risk measure.




FIGURE 5.1
Expected Shortfall $\mathbb{E}[X \mid X \geq \text{VaR}_\alpha(X)]$.

insurance point of view, expected shortfall is a very reasonable measure: Defining by $c = \text{VaR}_\alpha(X)$ a *critical loss threshold* corresponding to some confidence level α , expected shortfall capital provides a cushion against the mean value of losses exceeding the critical threshold c . In other words, TCE focusses on the expected loss in the tail, starting at c , of the portfolio's loss distribution. The critical threshold c , driven by the confidence level α , has to be fixed by the senior management of the bank and is part of the bank's credit management policy.

The following proposition provides another interpretation of coherency and can be found in [5].

5.1.2 Proposition *A risk measure γ is coherent if and only if there exists a family \mathcal{P} of probability measures such that*

$$\gamma = \gamma_{\mathcal{P}} ,$$

where $\gamma_{\mathcal{P}}$ is defined by

$$\gamma_{\mathcal{P}}(X) = \sup_{Q \in \mathcal{P}} \mathbb{E}_Q[X] \quad \text{for all } X \in L^\infty .$$

The probability measures in \mathcal{P} are called *generalized scenarios*.

The challenge underlying Proposition 5.1.2 is to find a suitable set \mathcal{P} of probability distributions matching a given coherent risk measure γ such that $\gamma = \gamma_{\mathcal{P}}$.

Economic capital, based on shortfall risk, can be defined as the mean loss above a threshold c minus the expected loss:

$$\text{EC}_{\text{TCE}}(c) = \mathbb{E}[X \mid X \geq c] - \mathbb{E}[X].$$

This calculation method for risk capital also includes events above the critical loss threshold c , e.g., $c = \text{VaR}_\alpha$, and answers the question “how bad is bad” from a mean value's perspective. If $c = \text{VaR}_\alpha(X)$, we write

$$\text{EC}_{\text{TCE}_\alpha} = \text{EC}_{\text{TCE}}(\text{VaR}_\alpha(X))$$

in the sequel in order to keep the notation simple.

Quantiles The following table shows a comparison of expected short-fall EC and VaR-EC over the 99%-quantile for different distributions. For example, one can see that if X had a normal distribution, the expected shortfall EC would not much differ from the VaR-EC. In contrast, for a t -distribution and the loss distribution $F_{p,\varrho}$ defined by a uniform (limit) portfolio as introduced right before Proposition 2.5.7, the difference between the ECs is quite significant.

	t(3)	N(0,1.73)	LN(0,1)	N(1.64,2.16)	Weil(1,1)	N(1,1)
std	1.73	1.73	2.16	2.16	1	1
EC _{VaR} (0.99)	4.54	4.02	8.56	5.02	3.6	2.32
EC _{TCE} (0.99)	6.99	4.61	13.57	5.76	4.6	2.66
rel.diff. (%)	54	15	58	15	27	14

	$F_{0.003,0.12}$
std	0.0039
EC _{VaR} (0.99)	0.0162
EC _{TCE} (0.99)	0.0237
rel.diff. (%)	46

This table highlights the sensitivity of the determination of economic capital w.r.t. its definition (VaR-based or shortfall-based) and the choice of the underlying probability measure. Here, the probability measures are given by the following distributions:

- the t -distribution $t(3)$ with 3 degrees of freedom,
- three normal distributions with different parameter pairs,
- the $(0, 1)$ -log-normal distribution,
- the $(1, 1)$ -Weibull distribution, and
- the uniform portfolio distribution $F_{p,\varrho}$ with uniform default probability $p = 30\text{bps}$ and uniform asset correlation $\varrho = 12\%$.

In the table, the first row contains the respective standard deviations, the second row shows the EC based on VaR at the 99%-confidence level, and the third row shows the EC based on expected shortfall with threshold $c = \text{VaR}_{99\%}$. The fourth row shows the relative difference between the two EC calculations.

5.2 Contributory Capital

If the economic capital EC of the bank is determined, one of the main tasks of risk management is to calculate *contributory economic capital* for each business division of the bank. Such a contributory EC could be interpreted as the *marginal capital* a single transaction respectively business unit adds or contributes to the overall required capital. Tasche [120] showed that there is only one definition for *risk contributions* that is suitable for *performance measurement*, namely the *derivative of the underlying risk measure in direction of the asset weight of the considered entity*. Additionally, Denault [24] showed by arguments borrowed from game theory that in the case of a so-called *1-homogeneous risk measure* its gradient is the only “allocation principle” that satisfies some “coherency” conditions.

So when talking about capital allocation, we are therefore faced with the problem of *differentiability* of some considered risk measure. To allow for differentiation, the risk measure must be “sufficiently smooth” for guaranteeing that the respective derivative exists. It turns out that the standard deviation as a risk measure has nice differentiability properties. This observation is the “heart” of Markovitz’s classical portfolio theory (var/covar approach). Now, VaR-based EC is a *quantile-based* quantity, and only in case of normal distributions quantiles and standard deviations are reconcilable from a portfolio optimization point of view. Because we know from previous chapters that credit portfolio loss distributions are typically skewed with fat tails, the “classical” approach can not be applied to credit risk in a straightforward manner without “paying” for the convenience of the approach by allowing for inconsistencies.

Fortunately, the var/covar approach for *capital allocation*, adapted to credit risk portfolios, in many cases yields acceptable results and is, due to its simplicity, implemented in most standard software packages. The following section explains the details.

5.2.1 Variance/Covariance Approach

At the core of var/covar is the question of what an individual credit or business unit contributes to the portfolio’s standard deviation UL_{PF} . To answer this question, the classical var/covar approach splits the

portfolio risk UL_{PF} into risk contributions RC_i in a way such that

$$\sum_{i=1}^m w_i \times RC_i = UL_{PF} .$$

In this way, the weighted risk contributions sum-up to the total risk of the portfolio, where “risk” is identified with volatility. It follows from Proposition 1.2.3 that

$$\begin{aligned} UL_{PF} &= \frac{1}{UL_{PF}} \sum_{i=1}^m w_i \cdot \sum_{j=1}^m w_j \eta_i \eta_j \sigma_{D_i} \sigma_{D_j} \rho_{ij} \\ &= \frac{1}{UL_{PF}} \sum_{i=1}^m w_i \cdot \sum_{j=1}^m w_j UL_i UL_j \rho_{ij}, \end{aligned}$$

where $\sigma_{D_i} = \sqrt{DP_i(1 - DP_i)}$ denotes the standard deviation of the default indicator and $UL_i = \eta_i \sigma_{D_i}$. Thus,

$$RC_i = \frac{UL_i}{UL_{PF}} \sum_{j=1}^m w_j UL_j \rho_{ij}$$

is a plausible quantity measuring the “risk portion” of credit i in a way such that all weighted risks sum-up to the portfolio’s UL. It is straightforward to show that the quantity RC_i corresponds to the covariance of credit (business unit) i and the total portfolio loss, divided by the portfolio’s volatility respectively UL. The definition of RC_i obviously is in analogy to *beta-factor models* used in market risk. Furthermore, RC_i is equal to the *partial derivative* of UL_{PF} w.r.t. w_i , the weight of the i -th credit in the portfolio, i.e.,

$$RC_i = \frac{\partial UL_{PF}}{\partial w_i} .$$

In other words, an increase in the weight of the considered credit by a small amount h in the portfolio, implies a growth of UL_{PF} by $h \times RC_i$. Coming from this side, it can in turn be easily shown that the total sum of weighted partial derivatives again equals UL_{PF} .

Regarding the ratio between risk contributions and the standard deviation of the individual exposure, it is true in most cases that

$$\frac{RC_i}{UL_i} \leq 1.$$

This quantity is known as the *retained risk* of unit i . It is the portion of risk of the i -th entity that has not been absorbed by diversification in the portfolio. In contrast, the quantity

$$1 - \frac{RC_i}{UL_i}$$

is often called the corresponding *diversification effect*.

Capital multiplier Since the total risk capital is typically determined via quantiles, i.e.,

$$EC_{VaR_\alpha} = VaR_\alpha(L) - \mathbb{E}[L] ,$$

the individual risk contributions have to be rescaled with the so-called *capital multiplier*

$$CM_\alpha = \frac{EC_{VaR_\alpha}}{UL_{PF}}$$

in order to imitate the classical approach from market risk. The contributory capital for credit i then equals

$$\delta_i = CM_\alpha \times RC_i, \quad \text{with} \quad \sum_{i=1}^m w_i \delta_i = EC_{VaR_\alpha} .$$

The quantity δ_i is called the *analytic capital contribution* of transaction i to the portfolio capital. For a business unit in charge for credits 1 to l , where $l < m$, the capital requirement is

$$CM_\alpha \sum_{j=1}^l RC_j .$$

Note that the capital multiplier is an auxiliary quantity depending on the particular portfolio, due to the fact that, in contrast to the normal distribution, the quantiles of credit portfolio loss distributions not only depend on the standard deviation, but also on other influences like correlations, default probabilities, and exposure weights. Therefore it is unrealistic, after changing the portfolio, to obtain the same capital multiplier CM_α as originally calculated.

In the next two sections we discuss capital allocation w.r.t. EC based on VaR and expected shortfall respectively.

5.2.2 Capital Allocation w.r.t. Value-at-Risk

Calculating risk contributions associated with the VaR risk measure is a natural but difficult attempt, since in general the quantile function will not be differentiable with respect to the asset weights. Under certain continuity assumptions on the joint density function of the random variables X_i , differentiation of $\text{VaR}_\alpha(X)$, where $X = \sum_i w_i X_i$, is guaranteed. One has (see [122])

$$\frac{\partial \text{VaR}_\alpha}{\partial w_i}(X) = \mathbb{E}[X_i \mid X = \text{VaR}_\alpha(X)]. \quad (5.1)$$

Unfortunately, the distribution of the portfolio loss $L = \sum w_i \hat{L}_i$, as specified at the beginning of this chapter, is purely discontinuous. Therefore the derivatives of VaR_α in the above sense will either not exist or vanish to zero. In this case we could still define risk contributions via the right-hand-side of Equation (5.1) by writing

$$\gamma_i = \mathbb{E}[\hat{L}_i \mid L = \text{VaR}_\alpha(L)] - \mathbb{E}[\hat{L}_i]. \quad (5.2)$$

For a clearer understanding, note that

$$\frac{\partial \mathbb{E}[L]}{\partial w_i} = \mathbb{E}[\hat{L}_i] \quad \text{and} \quad \sum_{i=1}^m w_i \gamma_i = \text{EC}_{\text{VaR}_\alpha}.$$

Additionally observe, that for a large portfolio and on an appropriate scale, the distribution of L will appear to be “close to continuous”. Unfortunately, even in such “approximately good” cases, the loss distribution often is not given in an analytical form in order to allow for differentiations.

Remark For the CreditRisk⁺ model, an analytical form of the loss distribution can be found; see [Section 2.4.2](#) and [Chapter 4](#) for a discussion of CreditRisk⁺. Tasche [121] showed that in the CreditRisk⁺ framework the VaR contributions can be determined by calculating the corresponding loss distributions several times with different parameters. Martin et al. [82] suggested an approximation to the partial derivatives of VaR via the so-called *saddle point method*.

Capital allocation based on VaR is not really satisfying, because in general, although $(\text{RC}_i)_{i=1, \dots, m}$ might be a reasonable partition of the portfolio’s standard deviation, it does not really say much about the

tail risks captured by the quantile on which VaR-EC is relying. Even if in general one views capital allocation by means of partial derivatives as useful, the problem remains that the var/covar approach completely neglects the dependence of the quantile on correlations. For example, var/covar implicitly assumes

$$\frac{\partial \text{VaR}_\alpha(X)}{\partial \text{UL}_{PF}} = \text{const} = \text{CM}_\alpha,$$

for the specified confidence level α . This is true for (multivariate) normal distributions, but generally not the case for loss distributions of credit portfolios. As a consequence it can happen that transactions require a contributory EC exceeding the original exposure of the considered transaction. This effect is very unpleasant. Therefore we now turn to expected shortfall-based EC instead of VaR-based EC.

5.2.3 Capital Allocations w.r.t. Expected Shortfall

At the beginning we must admit that shortfall-based risk contributions bear the same “technical” difficulty as VaR-based measures, namely the quantile function is not differentiable in general. But, we find in Tasche [122] that if the underlying loss distribution is “sufficiently smooth”, then TCE_α is partially differentiable with respect to the exposure weights. One finds that

$$\frac{\partial \text{TCE}_\alpha}{\partial w_i}(X) = \mathbb{E}[X_i \mid X \geq \text{VaR}_\alpha(X)].$$

In case the partial derivatives do not exist, one again can rely on the right-hand side of the above equation by defining shortfall contributions for, e.g., discontinuous portfolio loss variables $L = \sum w_i \hat{L}_i$ by

$$\zeta_i = \mathbb{E}[\hat{L}_i \mid L \geq \text{VaR}_\alpha(L)] - \mathbb{E}[\hat{L}_i], \tag{5. 3}$$

which is consistent with expected shortfall as an “almost coherent” risk measure. Analogous to what we saw in case of VaR-EC, we can write

$$\sum_{i=1}^m w_i \zeta_i = \text{EC}_{\text{TCE}_\alpha},$$

such that shortfall-based EC can be obtained as a weighted sum of the corresponding contributions.

Remarks With expected shortfall we have identified a coherent (or close to coherent) risk measure, which overcomes the major drawbacks of classical VaR approaches. Furthermore, shortfall-based measures allow for a consistent definition of risk contributions. We continue with some further remarks:

- The results on shortfall contributions together with the findings on differentiability in [105] indicate that the proposed capital allocation ζ_i can be used as a performance measure, as pointed out in Theorem 4.4 in [122], for example. In particular, it shows that if one increases the exposure to a counterparty having a RAROC above portfolio RAROC, the portfolio RAROC will be improved. Here RAROC is defined as the return over (contributory) economic capital.
- We obtain $\zeta_i < \hat{L}_i$, i.e., by construction the capital is always less than the exposure, a feature that is not shared by risk contributions defined in terms of covariances.
- Shortfall contributions provide a simple “first-order” statistics of the distribution of L_i conditional on $L > c$. Other statistics like conditional variance could be useful. (We do not know if conditional variance is coherent under all circumstances.)
- The definition of shortfall contributions reflects a causality relation. If counterparty i contributes higher to the overall loss than counterparty j in extreme loss scenarios, then, as a consequence, business with i should be more costly (assuming stand-alone risk characteristics are the same).
- Since $L, L_i \geq 0$, capital allocation rules according to shortfall contributions can easily be extended to the space of all coherent risk measures as defined in this chapter

5.2.4 A Simulation Study

In the simulation study we want to compare the two different allocation techniques, namely allocation based on VaR and allocation based on expected shortfall. We first tested it on a transaction base. In a subsequent test case we considered the allocation of capital to business units. There are at least two reasons justifying the efforts for the second test. First it might not be reasonable to allocate economic capital

that is based on extreme loss situations to a single transaction, since the risk in a single transaction might be driven by short-term volatility and not by the long-term view of extreme risks. The second reason is more driven by the computational feasibility of expected shortfall. In the “binary world” of default simulations, too many simulations are necessary in order to obtain a positive contribution conditional on extreme default events for all counterparties.

The basic result of the simulation study is that analytic contributions produce a steeper gradient between risky and less risky loans than tail risk contributions. In particular, loans with a high default probability but moderate exposure concentration require more capital in the analytic contribution method, whereas loans with high concentration require relatively more capital in the shortfall contribution method.

Transaction View The first simulation study is based on a credit portfolio considered in detail in [105]. The particular portfolio consists of 40 counterparties.

As capital definition, the 99% quantile of the loss distribution is used. Within the Monte-Carlo simulation it is straightforward to evaluate risk contributions based on expected shortfall. The resulting risk contributions and its comparison to the analytically calculated risk contributions based on the volatility decomposition are shown in [Figure 5.2](#).

In the present portfolio example the difference between the contributory capital of two different types, namely analytic risk contributions and contributions to shortfall, should be noticed, since even the order of the assets according to their risk contributions changed. The asset with the largest shortfall contribution is the one with the second largest var/covar risk contribution, and the largest var/vovar risk contribution goes with the second largest shortfall contribution. A review of the portfolio shows that the shortfall contributions are more driven by the relative asset size. However, it is always important to bear in mind that these results are still tied to the given portfolio.

It should also be noticed that the gradient of the EC is steeper for the analytic approach. Bad loans might be able to breach the hurdle rate in a RAROC-Pricing tool if one uses the expected shortfall approach, but might fail to earn above the hurdle rate if EC is based on var/covar.

Business Unit View The calculation of expected shortfall contributions requires a lot more computational power, which makes it less




FIGURE 5.2

The bar chart depicts the different risk contributions for every counterparty in the portfolio. The dark bars belong to the counterparty contribution measured by the short-fall; the white ones correspond to the analytic Var/Covar-contribution.

feasible for large portfolios. However, the capital allocation on the business level can accurately be measured by means of expected shortfall contributions. [Figure 5.3](#) shows an example of a bank with 6 business units. Again we see that expected shortfall allocation differs from var/covar allocation.

Under var/covar, it sometimes can even happen that the capital allocated to a business unit is larger if considered consolidated with the bank than capitalized standalone. This again shows the non-coherency of VaR measures. Such effects are very unpleasant and can lead to significant misallocations of capital. Here, expected shortfall provides the superior way of capital allocation. We conclude this chapter by a simple remark how one can calculate EC on VaR-basis but allocate capital shortfall-based.

If a bank calculates its total EC by means of VaR, it still can allocate capital in a coherent way. For this purpose, one just has to determine some threshold $c < \text{VaR}_\alpha$ such that

$$EC_{\text{TCE}}(c) \approx EC_{\text{VaR}_\alpha} .$$

This *VaR-matched expected shortfall* is a coherent risk measure preserving the VaR-based overall economic capital. It can be viewed as an approximation to VaR-EC by considering the whole tail of the loss distribution, starting at some threshold below the quantile, such that the resulting mean value matches the quantile. Proceeding in this way, allocation of the total VaR-based EC to business units will reflect the coherency of shortfall-based risk measures.

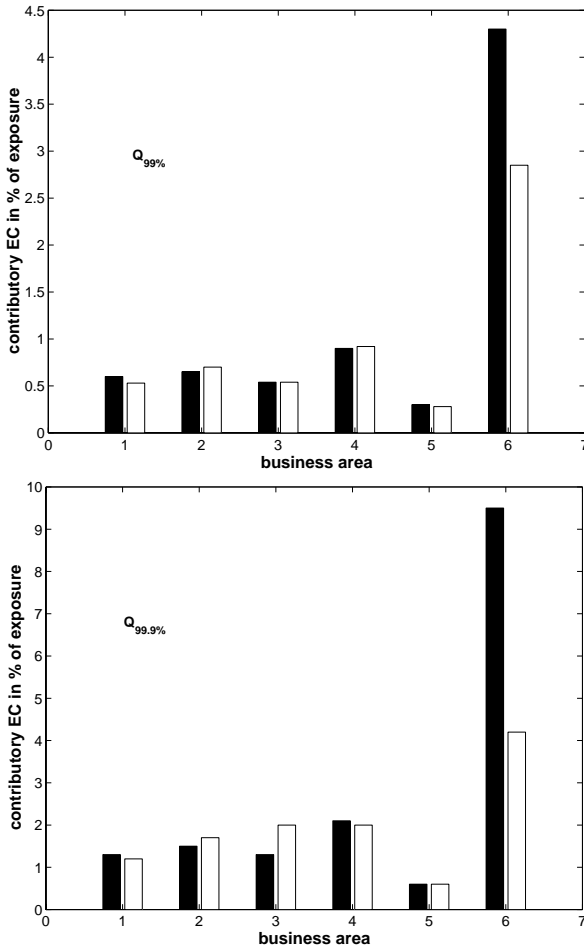


FIGURE 5.3

The bar charts depict the different risk contributions (top: 99% quantile, bottom: 99.9% quantile) of the business areas of a bank. The black bars are based on a Var/Covar approach; the white ones correspond to shortfall risk.

Chapter 6

Term Structure of Default Probability

So far, default has mostly been modeled as a binary event (except the intensity model), suited for single-period considerations within the regulatory framework of a fixed planning horizon. However, the choice of a specific period like one year is more or less arbitrary. Even more, default is an inherently time-dependent event. This chapter serves to introduce the idea of a term structure of default probability. This credit curve represents a necessary prerequisite for a time-dependent modeling as in [Chapters 7](#) and [8](#). In principle, there are three different methods to obtain a credit curve: from historical default information, as implied probabilities from market spreads of defaultable bonds, and through Merton's option theoretic approach. The latter has already been treated in a previous chapter, but before introducing the other two in more detail we first lay out some terminology used in survival analysis (see [15,16] for a more elaborated presentation).

6.1 Survival Function and Hazard Rate

For any model of default timing, let $S(t)$ denote the probability of surviving until t . With help of the “time-until-default” τ (or briefly “default time”), a continuous random variable, the *survival function* $S(t)$ can be written as

$$S(t) = \mathbb{P}[\tau > t], \quad t \geq 0.$$

That is, starting at time $t = 0$ and presuming no information is available about the future prospects for survival of a firm, $S(t)$ measures the likelihood that it will survive until time t . The probability of default between time s and $t \geq s$ is simply $S(s) - S(t)$. In particular, if $s = 0$,

and because $S(0) = 1$, then the probability of default $F(t)$ is

$$F(t) = 1 - S(t) = \mathbb{P}[\tau \leq t], \quad t \geq 0. \tag{6. 1}$$

$F(t)$ is the distribution function of the random default time τ . The corresponding probability density function is defined by

$$f(t) = F'(t) = -S'(t) = \lim_{\Delta \rightarrow 0^+} \frac{\mathbb{P}[t \leq \tau < t + \Delta]}{\Delta},$$

if the limit exists. Furthermore, we introduce the *conditional* or *forward* default probability

$$p(t|s) = \mathbb{P}[\tau \leq t | \tau > s], \quad t \geq s \geq 0,$$

i.e., the probability of default of a certain obligor between t and s conditional on its survival up to time s , and

$$q(t|s) = 1 - p(t|s) = \mathbb{P}[\tau > t | \tau > s] = S(t)/S(s), \quad t \geq s \geq 0,$$

the forward survival probability. An alternative way of characterizing the distribution of the default time τ is the *hazard function*, which gives the instantaneous probability of default at time t conditional on the survival up to t . The hazard function is defined via

$$\mathbb{P}[t < \tau \leq t + \Delta t | \tau > t] = \frac{F(t + \delta t) - F(t)}{1 - F(t)} \approx \frac{f(t)\Delta t}{1 - F(t)}$$

as

$$h(t) = \frac{f(t)}{1 - F(t)}.$$

Equation(6. 1) yields

$$h(t) = \frac{f(t)}{1 - F(t)} = -\frac{S'(t)}{S(t)},$$

and solving this differential equation in $S(t)$ results in

$$S(t) = e^{-\int_0^t h(s)ds}. \tag{6. 2}$$

This allows us to express $q(t|s)$ and $p(t|s)$ as

$$q(t|s) = e^{-\int_s^t h(u)du}, \tag{6. 3}$$

$$p(t|s) = 1 - e^{-\int_s^t h(u)du}. \tag{6. 4}$$

Additionally, we obtain

$$F(t) = 1 - S(t) = 1 - e^{-\int_0^t h(s)ds},$$

and

$$f(t) = S(t)h(t).$$

One could assume the hazard rate to be piecewise constant, i.e., $h(t) = h_i$ for $t_i \leq t < t_{i+1}$. In this case, it follows that the density function of τ is

$$f(t) = h_i e^{-h_i t} \mathbf{1}_{[t_i, t_{i+1}[}(t),$$

showing that the survival time is exponentially distributed with parameter h_i . Furthermore, this assumption entails over the time interval $[t_i, t_{i+1}[$ for $0 < t_i \leq t < t_{i+1}$

$$q(t|t_i) = e^{-\int_{t_i}^t h(u)du} = e^{-h_i(t-t_i)}.$$

Remark The “forward default rate” $h(t)$ as a basis of a default risk term structure is in close analogy to a forward interest rate, with zero-coupon bond prices corresponding to survival probabilities. The hazard rate function used to characterize the distribution of survival time can also be called a “credit curve” due to its similarity to a yield curve. If h is continuous then $h(t)\Delta t$ is approximately equal to the probability of default between t and $t + \Delta t$, conditional on survival to t . Understanding the first arrival time τ as associated with a Poisson arrival process, the constant mean arrival rate h is then called intensity and often denoted by λ^1 . Changing from a deterministically varying intensity to random variation, and thus closing the link to the stochastic intensity models [32], turns Equation (6. 3) into

$$q(t|s) = \mathbb{E}_s \left[e^{-\int_s^t h(u)du} \right],$$

where \mathbb{E}_s denotes expectation given all information available at time s .

¹Note that some authors explicitly distinguish between the intensity $\lambda(t)$ as the arrival rate of default at t conditional on *all* information available at t , and the forward default rate $h(t)$ as arrival rate of default at t , conditional *only* on survival until t .

6.2 Risk-neutral vs. Actual Default Probabilities

When estimating the risk and the value of credit-related securities we are faced with the question of the appropriate probability measure, *risk-neutral* or *objective probabilities*. But in fact the answer depends on the objective we have. If one is interested in estimating the economic capital and risk charges, one adopts an actuarial-like approach by choosing historical probabilities as underlying probability measure. In this case we assume that actual default rates from historical information allow us to estimate a capital quota protecting us against losses in worst case default scenarios. The objective is different when it comes to pricing and hedging of credit-related securities. Here we have to model under the risk-neutral probability measure. In a risk-neutral world all individuals are indifferent to risk. They require no compensation for risk, and the expected return on all securities is the risk-free interest rate. This general principle in option pricing theory is known as risk-neutral valuation and states that it is valid to assume the world is risk-neutral when pricing options. The resulting option prices are correct not only in the risk-neutral world, but in the real world as well. In the credit risk context, risk-neutrality is achieved by calibrating the default probabilities of individual credits with the market-implied probabilities drawn from bond or credit default swap spreads. The difference between actual and risk-neutral probabilities reflects risk-premiums required by market participants to take risks. To illustrate this difference suppose we are pricing a one-year par bond that promises its face value 100 and a 7% coupon at maturity. The one-year risk-free interest rate is 5%. The actual survival probability for one year is $1 - DP = 0.99$; so, if the issuer survives, the investor receives 107. On the other hand, if the issuer defaults, with actual probability $DP = 0.01$, the investor recovers 50% of the par value. Simply discounting the expected payoff computed with the actual default probability leads to

$$\frac{(107 \times 0.99 + 50\% \times 100 \times 0.01)}{1 + 5\%} = 101.36,$$

which clearly overstates the price of this security. In the above example we have implicitly adopted an actuarial approach by assuming that the price the investor is to pay should exactly off set the expected loss due to a possible default. Instead, it is natural to assume that investors

are concerned about default risk and have an aversion to bearing more risk. Hence, they demand an additional risk premium and the pricing should somehow account for this risk aversion. We therefore turn the above pricing formula around and ask which probability results in the quoted price, given the coupons, the risk-free rate, and the recovery value. According to the risk-neutral valuation paradigm, the fact that the security is priced at par implies that

$$100 = \frac{(107 \times (1 - DP^*) + 50\% \times 100 \times DP^*)}{1 + 5\%}.$$

Solving for the market-implied risk-neutral default probability yields $DP^* = 0.0351$. Note that the actual default probability $DP = 0.01$ is less than DP^* . Equivalently, we can say that the bond is priced as though it were a break-even trade for a “stand-in” investor who is not risk adverse but assumes a default probability of 0.0351. The difference between DP and DP^* reflects the risk premium for default timing risk. Most credit market participants think in terms of spreads rather than in terms of default probabilities, and analyze the shape and movements of the spread curve rather than the change in default probabilities. And, indeed, the link between credit spread and probability of default is a fundamental one, and is analogous to the link between interest rates and discount factors in fixed income markets. If s represents a multiplicative spread over the risk-free rate one gets

$$DP^* = \frac{1 - \frac{1}{1+s}}{1 - REC} \approx \frac{s}{1 - REC},$$

where the approximation is also valid for additive spreads.

“Actuarial credit spreads” are those implied by assuming that investors are neutral to risk, and use historical data to estimate default probabilities and expected recoveries. Data from Fons [43] suggest that corporate yield spreads are much larger than the spreads suggested by actuarial default losses alone. For example, actuarially implied credit spreads on a A-rated 5-year US corporate debt were estimated by Fons to be six basis points. The corresponding market spreads have been in the order of 100 basis points. Clearly, there is more than default risk behind the difference between “actuarial credit spreads” and actual yield spreads, like liquidity risk, tax-related issues, etc. But even after measuring spreads relative to AAA yields (thereby stripping out treasury effects), actuarial credit spreads are smaller than actual market spreads, especially for high-quality bonds.

6.3 Term Structure Based on Historical Default Information

Multi-year default probabilities can be extracted from historical data on corporate defaults similarly to the one-year default probabilities. But before going into details we first show a “quick and dirty” way to produce a whole term structure if only one-year default probabilities are at hand.

6.3.1 Exponential Term Structure

The derivation of exponential default probability term structure is based on the idea that credit dynamics can be viewed as a two-state time-homogeneous Markov-chain, the two states being survival and default, and the unit time between two time steps being Δ . Suppose a default probability DP_T for a time interval T (e.g., one year) has been calibrated from data; then the survival probability for the time unit Δ (e.g., one day) is given by

$$\mathbb{P}[\tau > t + \Delta | \tau \geq t] = (1 - DP_T)^{\Delta/T}, \quad (6.5)$$

and the default probability for the time t , in units of Δ , is then

$$DP_t = 1 - (1 - DP_T)^{t/T}. \quad (6.6)$$

In the language of survival analysis we can write for the probability of survival until T

$$1 - DP_T = q(T|0) = e^{-\int_0^T h(u)du} = e^{-\bar{h}T},$$

where the last equation defines the average default rate \bar{h} ,

$$\bar{h} = -\log(1 - DP_T) / T.$$

Assuming a constant default rate over the whole life time of the debt, Equation (6.6) reads

$$F(t) = q(t|0) = 1 - p(t|0) = 1 - e^{-\bar{h}t}.$$

6.3.2 Direct Calibration of Multi-Year Default Probabilities

Rating agencies also provide data on multi-year default rates in their reports. For example, Moody's [95] trailing $T + 1$ -month default rates for month t and rating universe k are defined as

$$D_{k,t} = \frac{\sum_{i=t-T}^t Y_{k,i}}{I_{k,t-11}}. \quad (6.7)$$

k , for example, could be all corporate issuers, US speculative grade issuers, or Ba-rated issuers in the telecom sector. The numerator is the sum of defaulters, Y , in month t that were in the rating universe k as of $t - T$. The denominator, $I_{k,t}$, is the number of issuers left in the rating universe k in month $t - T$, adjusted to reflect the withdrawal from the market of some of those issuers for noncredit-related reasons (e.g., maturity of debt). The adjustment for withdrawal is important because the denominator is intended to represent the number of issuers who could potentially have defaulted in the subsequent $T + 1$ -month period. Underlying Equation (6.7) is the assumption that defaults in a given rating universe are independent and identically distributed Bernoulli random variables, i.e., the number of defaults w.r.t. a certain pool, rating, and year follow a binomial distribution. Note that this assumption is certainly not correct in a strict sense; in fact, correlated defaults are the core issue of credit portfolio models.

Moody's employs a dynamic cohort approach to calculating multi-year default rates. A cohort consists of all issuers holding a given estimated senior rating at the start of a given year. These issuers are then followed through time, keeping track of when they default or leave the universe for noncredit-related reasons. For each cumulation period, default rates based on dynamic cohorts express the ratio of issuers who did default to issuers who were in the position to default over that time period. In terms of Equation (6.7) above, this constitutes lengthening the time horizon T ($T = 11$ in the case of one-year default rates). Since more and more companies become rated over the years, Moody's and S&P use an issuer weighted average to averaged cumulative default rates. To estimate the average risk of default over time horizons longer than one year, Moody's calculates the risk of default in each year since a cohort was formed. The issuer-weighted average of each cohort's one-year default rate forms the average cumulative one-year default rate. The issuer-weighted average of the second-year (marginal) default rates (default in exactly the second year) cumulated with that of the first year




FIGURE 6.1
Cumulative default rate for A-rated issuer. Weighted average default rates smooth economic cycles effects on individual cohorts, from [95].

yields the two-year average cumulative default rate, and so on. [Figure 6.1](#) shows how different cohorts produce a different credit history in response to different economic and market conditions.

[Table 6.1](#) gives the average cumulative default rates as reported by Moody's [95]. A closer look reveals some unpleasant features in this table. For example, one would expect column monotony for each year, i.e., high credit quality should never show a higher default rate than low credit quality, which is violated at various entries. Furthermore, some marginal default rates are zero, even for non-triple A rated corporates, which is unrealistic. Clearly, these problems stem from a lack of sufficient data for a reliable statistical analysis, and, obviously, pooling on a more aggregated level produces more satisfactory results w.r.t. these plausibility conditions; see [Table 6.2](#). In the next section we show a way to avoid these deficiencies by use of migration analysis.

One can argue that the issuer-weighted arithmetic mean is perhaps not the right thing to do. Since more and more corporates are rated during the last years, issuer-weighted averaging means that recent years

TABLE 6.1: Average cumulative default from 1 to 10 years – 1983-2000, from [95].

	1	2	3	4	5	6	7	8	9	10
Aaa	0,00%	0,00%	0,00%	0,06%	0,18%	0,25%	0,34%	0,43%	0,43%	0,43%
Aa1	0,00%	0,00%	0,00%	0,21%	0,21%	0,35%	0,35%	0,35%	0,35%	0,35%
Aa2	0,00%	0,00%	0,06%	0,18%	0,41%	0,49%	0,59%	0,71%	0,85%	1,01%
Aa3	0,06%	0,09%	0,17%	0,26%	0,37%	0,49%	0,49%	0,49%	0,49%	0,49%
A1	0,00%	0,03%	0,30%	0,47%	0,59%	0,73%	0,79%	0,86%	0,86%	0,96%
A2	0,00%	0,02%	0,16%	0,41%	0,62%	0,84%	0,99%	1,35%	1,63%	1,71%
A3	0,00%	0,12%	0,22%	0,30%	0,35%	0,47%	0,68%	0,77%	0,97%	1,09%
Baa1	0,07%	0,30%	0,53%	0,86%	1,19%	1,43%	1,82%	2,05%	2,20%	2,20%
Baa2	0,06%	0,29%	0,61%	1,22%	1,89%	2,54%	2,93%	3,17%	3,46%	3,81%
Baa3	0,39%	1,05%	1,62%	2,47%	3,15%	4,09%	4,99%	5,95%	6,54%	7,03%
Ba1	0,64%	2,10%	3,81%	6,15%	8,12%	10,09%	11,43%	12,75%	13,35%	14,08%
Ba2	0,54%	2,44%	4,95%	7,32%	9,27%	10,88%	12,59%	13,60%	14,27%	14,71%
Ba3	2,47%	6,82%	11,68%	16,18%	20,63%	24,74%	28,39%	32,28%	35,83%	38,22%
B1	3,48%	9,71%	15,59%	20,56%	25,62%	30,78%	36,15%	40,30%	44,16%	48,01%
B2	6,23%	13,70%	20,03%	24,63%	28,24%	31,14%	32,73%	34,33%	35,03%	35,90%
B3	11,88%	20,18%	26,71%	31,95%	36,68%	39,89%	42,81%	46,80%	51,42%	53,53%
Caa1-C	18,85%	28,29%	34,51%	40,23%	43,42%	46,48%	46,48%	49,73%	53,92%	59,04%
Investment-Grade	0,05%	0,17%	0,35%	0,60%	0,84%	1,08%	1,28%	1,47%	1,62%	1,73%
Speculative-Grade	3,69%	8,39%	12,87%	16,80%	20,39%	23,61%	26,44%	29,04%	31,22%	32,89%
All Corporates	1,21%	2,72%	4,12%	5,34%	6,39%	7,30%	8,05%	8,71%	9,23%	9,61%

TABLE 6.2: Average cumulative default by letter rating from 1 to 10 years –
1970-2000, from [95].

Rating	1	2	3	4	5	6	7	8	9	10
Aaa	0.00%	0.00%	0.00%	0.04%	0.12%	0.21%	0.31%	0.42%	0.54%	0.67%
Aa	0.02%	0.04%	0.08%	0.20%	0.31%	0.43%	0.55%	0.67%	0.76%	0.83%
A	0.01%	0.05%	0.18%	0.31%	0.45%	0.61%	0.78%	0.96%	1.18%	1.43%
Baa	0.14%	0.44%	0.83%	1.34%	1.82%	2.33%	2.86%	3.39%	3.97%	4.56%
Ba	1.27%	3.57%	6.11%	8.65%	11.23%	13.50%	15.32%	17.21%	19.00%	20.76%
B	6.16%	12.90%	18.76%	23.50%	27.92%	31.89%	35.55%	38.69%	41.51%	44.57%
Investment-Grade	0.05%	0.17%	0.35%	0.59%	0.82%	1.07%	1.34%	1.61%	1.91%	2.21%
Speculative-Grade	4.15%	8.39%	12.19%	15.48%	18.56%	21.26%	23.48%	25.60%	27.54%	29.46%
All Corporates	1.30%	2.61%	3.76%	4.77%	5.67%	6.46%	7.13%	7.76%	8.37%	8.96%

have much more impact than years further back in history and the result does not reflect a typical year as averaged over economic cycles.

Having now extracted cumulative default probabilities at discrete points in time, DP_i , we might be interested in a continuous version of this term structure. The simplest answer is an adroit linear interpolation of the multi-year default probability table (the interpolation between rating classes ought to be done on a logarithmic scale).

A slightly more sophisticated method can be formulated with the help of the forward default rate h . The forward default probability between t_i and t_{i+1} is given by

$$p(t_{i+1}|t_i) = \frac{DP_{i+1} - DP_i}{1 - DP_i} = 1 - \exp\left(-\int_{t_i}^{t_{i+1}} h(u)du\right).$$

Note that $DP_i = F(t_i)$. Define for the time interval $[t_i, t_{i+1}]$ an average forward default rate by

$$\bar{h}_i = \frac{1}{t_{i+1} - t_i} \int_{t_i}^{t_{i+1}} h(u)du, \quad \text{for } i = 0, \dots, n.$$

In terms of the multi-year default probabilities the forward default rate for period i is

$$\bar{h}_i = -\frac{1}{t_{i+1} - t_i} \log\left(\frac{1 - DP_{i+1}}{1 - DP_i}\right).$$

Two hazard rate functions obtained from the multi-year default probabilities in Table 6.2 are depicted in Figure 6.2 and show a typical feature: investment grade securities tend to have an upward sloping hazard rate term structure, whereas speculative grades tend to have a downward sloping term structure.

The cumulative default probability until time t , $t_i \leq t < t_{i+1}$ boils down to

$$\begin{aligned} DP_t &= F(t) = 1 - q(t_i|0)q(t|t_i) \\ &= 1 - (1 - DP_i) \left(\frac{1 - DP_{i+1}}{1 - DP_i}\right)^{(t-t_i)/(t_{i+1}-t_i)}. \end{aligned}$$

For $0 < t < 1$ we obtain again the exponential term structure. For $t > t_n$ the term structure can be extrapolated by assuming a constant forward default rate \bar{h}_{n-1} beyond t_{n-1} ,

$$DP_t = 1 - (1 - DP_{n-1}) \left(\frac{1 - DP_n}{1 - DP_{n-1}}\right)^{(t-t_n)/(t_{i+1}-t_i)}.$$




FIGURE 6.2
Hazard rate functions for ratings A and B, calculated from Moody’s cumulative default rates.

6.3.3 Migration Technique and Q-Matrices

The reliability of the default frequencies reported in [Table 6.1](#) strongly depends on the quality of the underlying data. In the previous section we pointed out that the amount of data used for the calibration decreases with increasing time horizon. As a consequence, the quality of the calibration suffers from a lack of sufficient data for a reliable statistical analysis, especially at larger time horizons. For this reason, we now discuss a different approach to multi-year default probabilities, whose key idea is the use of migration analysis.

According to migration frequency tables reported by Moody’s, S&P, and other rating agencies it is quite likely that corporates experience changes in credit quality over the years. This phenomenon is called *credit migration*, and the likelihoods of transitions from a given rating category to another are collected in *migration matrices*.

The migration technique can best be treated within the mathematical framework of Markov chains, i.e., we assume the existence of a credit migration process controlled solely by the transition probabilities given in the one-year migration matrix. More precisely, we define the finite state space of the chain covering possible bond ratings, e.g.,

$\Omega = \{AAA, AA, A, BBB, BB, B, CCC, Default\}$ and assign to every pair (i, j) of states a transition or migration probability

$$m_{ij} = P[i \rightarrow j] \quad (i = 1, \dots, 7; j = 1, \dots, 8),$$

where $P[i \rightarrow j]$ denotes the probability of change from rating class i at the beginning of a year to rating class j at year's end. In the present context the Markov property represents the assumption that the evolution of credit migration is independent of the past credit migration history. Through the homogeneity we assume the migration rates to be independent of time, i.e., the probability of a one-year migration $i \rightarrow j$ does not depend on the considered year.

Remark Clearly, both the assumptions of time-homogeneity and of the Markov property are very likely not precisely true for real credit ratings and in fact open to vivid debate. We do not consider such issues here, but rather rely on the simplified assumptions.

Now we collect the migration probabilities into a one-year migration (or transition) matrix $M = (m_{ij})_{i,j=1,\dots,8}$ where the 8-th row is given by the vector $(0, 0, 0, 0, 0, 0, 0, 1)$. The following properties of M follow immediately:

- (i) M has only nonnegative entries: $m_{ij} \geq 0$ for $i, j = 1, \dots, 8$.
- (ii) All row sums of M are equal to 1: $\sum_{j=1}^8 m_{ij} = 1$ for $i = 1, \dots, 8$.
- (iii) The last column contains the 1-year default probabilities: $m_{i,8} = PD(1, i)$ for $i = 1, \dots, 7$.
- (iv) The default state is absorbing: $m_{8,j} = 0$ for $j = 1, \dots, 7$, and $m_{8,8} = 1$.

This means that there is no escape from the default state.

Rating agencies publish transition matrices for one, two, or more years, usually with an additional column representing the no-longer rated debts. Since this primarily occurs when a company's outstanding debt issue expires, this portion is typically distributed along the rows proportionally to the probability weights in the rated states. Moody's average one-year migration matrix [95], for example, is

	<i>Aaa</i>	<i>Aa</i>	<i>A</i>	<i>Baa</i>	<i>Ba</i>	<i>B</i>	<i>C</i>	<i>Default</i>	<i>WR</i>
<i>Aaa</i>	0.8617	0.0945	0.0102	0.0000	0.0003	0.0000	0.0000	0.0000	0.0333
<i>Aa</i>	0.0110	0.8605	0.0893	0.0031	0.0011	0.0001	0.0000	0.0003	0.0346
<i>A</i>	0.0006	0.0285	0.8675	0.0558	0.0066	0.0017	0.0001	0.0001	0.0391
<i>Baa</i>	0.0006	0.0034	0.0664	0.8100	0.0552	0.0097	0.0008	0.0016	0.0523
<i>Ba</i>	0.0003	0.0006	0.0054	0.0546	0.7550	0.0818	0.0053	0.0132	0.0838
<i>B</i>	0.0001	0.0004	0.0020	0.0056	0.0592	0.7593	0.0303	0.0641	0.0790
<i>C</i>	0.0000	0.0000	0.0000	0.0087	0.0261	0.0562	0.5701	0.2531	0.0858

and reads after adjustment for rating withdrawal

$$M_{\text{Moody's}} = \begin{pmatrix} \begin{array}{c|cccccccc} \textit{Aaa} & \textit{Aaa} & \textit{Aa} & \textit{A} & \textit{Baa} & \textit{Ba} & \textit{B} & \textit{C} & \textit{D} \\ \hline \textit{Aaa} & 0.8914 & 0.0978 & 0.0106 & 0.0000 & 0.0003 & 0.0000 & 0.0000 & 0.0000 \\ \textit{Aa} & 0.0114 & 0.8913 & 0.0925 & 0.0032 & 0.0011 & 0.0001 & 0.0000 & 0.0003 \\ \textit{A} & 0.0006 & 0.0297 & 0.9028 & 0.0581 & 0.0069 & 0.0018 & 0.0001 & 0.0001 \\ \textit{Baa} & 0.0006 & 0.0036 & 0.0701 & 0.8547 & 0.0582 & 0.0102 & 0.0008 & 0.0017 \\ \textit{Ba} & 0.0003 & 0.0007 & 0.0059 & 0.0596 & 0.8241 & 0.0893 & 0.0058 & 0.0144 \\ \textit{B} & 0.0001 & 0.0004 & 0.0022 & 0.0061 & 0.0643 & 0.8244 & 0.0329 & 0.0696 \\ \textit{C} & 0.0000 & 0.0000 & 0.0000 & 0.0095 & 0.0285 & 0.0615 & 0.6236 & 0.2769 \\ \textit{D} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array} \end{pmatrix}.$$

A useful consequence of the Markovian property and the time-homogeneity is the fact that the n -year transition matrix is simply given by the n th power of the one-year transition matrix,

$$M_n = M_1^n,$$

where again the cumulative n th year default probabilities for a rating class is given by the last column of M_n .

Properties (i), (ii) make M a *stochastic matrix*. Furthermore, one might want to impose the following plausibility constraints to reflect our intuition.

- (v) Low-risk states should never show a higher default probability than high-risk states, i.e., $M_{i8} \leq M_{i+18}$, $i = 1, \dots, 7$.
- (vi) It should be more likely to migrate to closer states than to more distant states (row monotony towards the diagonal),

$$M_{ii+1} \geq M_{ii+2} \geq M_{ii+3} \dots$$

$$M_{ii-1} \geq M_{ii-2} \geq M_{ii-3} \dots$$

- (vii) The chance of migration into a certain rating class should be greater for more closely adjacent rating categories (column monotony towards the diagonal).

$$M_{i+1i} \geq M_{i+2i} \geq M_{i+3i} \dots$$

$$M_{i-1i} \geq M_{i-2i} \geq M_{i-3i} \dots$$

Insofar as a lower rating presents a higher credit risk, Jarrow et al. [64] formulated the condition:

- (viii) $\sum_{j \geq k} m_{ij}$ is a nondecreasing function of i for every fixed k ,

which is equivalent to requiring that the underlying Markov chain be *stochastically monotonic*. Note that row and column monotony towards the diagonal (properties (vi) and (vii)) implies stochastic monotony but not vice versa.

The problem with this wish list is that one cannot expect these properties to be satisfied by transition matrices sampled from historical data; so, the question remains how to best match a transition matrix to sampled data but still fulfill the required properties. Ong [104] proposes to solve this optimization problem, with the plausibility constraints stated as “soft conditions”, through a “simulated annealing” approach, where perturbed matrices are produced through additional random terms and tested to find an optimal solution. At this point we do not want to dive into the vast world of multidimensional optimization algorithms, but rather turn to another approach for obtaining a suitable migration matrix, namely via generators.

Generator Matrix The shortest time interval from which a transition matrix is estimated is typically one year. Data quality of rating transition observations within a shorter period is too poor to allow for a reliable estimate of a migration matrix. Nevertheless, for valuation purposes or loans that allow for the possibility of under-year intervention we are interested in transition matrices for shorter time periods. One might be tempted to approach this problem by fractional powers of M , but unfortunately the roots of transition matrices are not stochastic in general nor is it clear which root to choose when more of them exist. The idea is now to try to embed the time-discrete Markov

chain in a time-continuous Markov process, the latter being totally controlled by its generator. Collecting the rates of migration in a matrix Q , a time-dependent transition matrix $M(t)$, $t \geq 0$ then satisfies the matrix-valued (backward) differential equation

$$dM(t) = QM(t)dt.$$

Under the boundary condition $M(0) = I$, where I is the identity matrix, the formal solution of the differential equation is the matrix exponential

$$M(t) = e^{tQ} = \sum_{k=0}^{\infty} \frac{(tQ)^k}{k!}. \tag{6. 8}$$

Furthermore, the following theorem holds [102]:

6.3.1 Theorem $M(t)$ defined by (6. 8) is a stochastic matrix for all $t \geq 0$ if and only if $Q = (q_{ij})$ satisfies the following properties:

- (i) $0 \leq -q_{ii} < \infty$ for all $i = 1, \dots, 8$;
- (ii) $q_{ij} \geq 0$ for all $i \neq j$;
- (iii) $\sum_{j=1}^8 q_{ij} = 0$ for all $i = 1, \dots, 8$.

In Markov chain theory such matrices are called Q -matrices or *generators*. Unfortunately, this theorem is not that much of a help. Since we only have a single-period transition matrix available the existence of the true generator is not necessarily guaranteed. The problem of finding generators for empirical transition matrices has been comprehensively treated by Israel et al. [62]. They rediscovered some previously found results and derived some new findings. In the following we freely cite the most useful ones for our purposes and refer for proofs to the literature.

Equation (6. 8), nevertheless, can give us some guidance on how to find a valid, or at least construct an approximate, generator, i.e., the matrix logarithm.

6.3.2 Theorem ([62]) Let $M = (m_{ij})$ be an $n \times n$ strictly diagonally dominant Markov transition matrix, i.e., $m_{ii} > 1/2$ for all i . Then the series

$$\tilde{Q} = \sum_{k=1}^l (-1)^{k+1} \frac{(M - I)^k}{k}$$

converges geometrically quickly for $l \rightarrow \infty$, and gives rise to an $n \times n$ matrix \tilde{Q} having row-sums zero, such that $\exp(\tilde{Q}) = M$ exactly.

Note that the condition of strictly diagonal dominance is only a sufficient one. It is usually satisfied by credit migration matrices. For theorems on the (non-)existence of true generators see [62] and the references therein. The main problem of the log-expansion is that the matrix \tilde{Q} is not guaranteed to have nonnegative off-diagonal entries, which we need by the first theorem. However, any negative off-diagonal entries of \tilde{Q} will usually be quite small. Therefore, we try to correct the matrix simply by replacing these negative entries by zero, and redistribute the values by some appropriate ad hoc rules to the other entries to preserve the property of having vanishing row sum, in the hope that the thus obtained Q-matrix yields an, in some sense close, embedding.

One version is to define a Q-matrix Q from \tilde{Q} as (see also Stromquist [119])

$$q_{ij} = \max(\tilde{q}_{ij}, 0), \quad i \neq j; \quad q_{ii} = \tilde{q}_{ii} + \sum_{i \neq j} \min(\tilde{q}_{ij}, 0), \quad (6.9)$$

i.e., the sum of the negative off-diagonal entries is allotted in full to the diagonal element of the respective row.

A different Q-matrix is obtained by adding the negative values back to all entries of the same row that have the correct sign, proportional to their absolute values (see also Araten [3] for a closely related algorithm), i.e., let

$$g_i = |\tilde{q}_{ii}| + \sum_{i \neq j} \max(\tilde{q}_{ij}, 0); \quad b_i = \sum_{i \neq j} \max(-\tilde{q}_{ij}, 0)$$

and set

$$\hat{q}_{ij} = \begin{cases} 0, & i \neq j \text{ and } \tilde{q}_{ij} < 0 \\ \tilde{q}_{ij} - b_i |\tilde{q}_{ij}| / g_i, & \text{otherwise if } g_i > 0 \\ \tilde{q}_{ij}, & \text{otherwise if } g_i = 0. \end{cases} \quad (6.10)$$

In both cases the new matrix will still have by construction row sum zero, but now with nonnegative off-diagonals. Clearly, it will no longer satisfy $\exp(Q) = M$. Other, more refined “redistribution” choices are conceivable; however, they would rarely induce substantial difference to the distance of $\exp(Q)$ and M . Note that it is possible that a valid generator exists even if the \tilde{Q} computed by the series expansion is not a valid one. Furthermore, it may be possible that there exist more than one valid generator for a given transition matrix M .

Assuming that there is at most one migration per year Jarrow et al. [64] derived the following formula for a suitable generator:

$$\check{q}_{ii} = \log(m_{ii}), \quad \check{q}_{ij} = m_{ij} \log(m_{ii}) / (m_{ii} - 1) \quad \text{for } i \neq j. \quad (6. 11)$$

Let us now consider some numerical examples. From the Moody's one-year transition matrix $M_{Moody's}$ we arrive with Equation (6. 11) at the Q-matrix

$$\check{Q}_{Moody's} = \begin{pmatrix} -0.1150 & 0.1035 & 0.0112 & 0.0000 & 0.0003 & 0.0000 & 0.0000 & 0.0000 \\ 0.0121 & -0.1150 & 0.0979 & 0.0034 & 0.0012 & 0.0001 & 0.0000 & 0.0003 \\ 0.0007 & 0.0312 & -0.1023 & 0.0611 & 0.0072 & 0.0019 & 0.0001 & 0.0001 \\ 0.0007 & 0.0039 & 0.0757 & -0.1570 & 0.0629 & 0.0111 & 0.0009 & 0.0018 \\ 0.0004 & 0.0007 & 0.0065 & 0.0655 & -0.1935 & 0.0982 & 0.0064 & 0.0158 \\ 0.0001 & 0.0005 & 0.0024 & 0.0067 & 0.0707 & -0.1930 & 0.0362 & 0.0765 \\ 0.0000 & 0.0000 & 0.0000 & 0.0119 & 0.0358 & 0.0771 & -0.4722 & 0.3473 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

with the matrix exponential

$$\exp(\check{Q}_{Moody's}) = \begin{pmatrix} 0.8919 & 0.0925 & 0.0146 & 0.0006 & 0.0004 & 0.0000 & 0.0000 & 0.0000 \\ 0.0108 & 0.8933 & 0.0882 & 0.0057 & 0.0015 & 0.0003 & 0.0000 & 0.0003 \\ 0.0008 & 0.0282 & 0.9063 & 0.0540 & 0.0080 & 0.0023 & 0.0002 & 0.0003 \\ 0.0007 & 0.0045 & 0.0670 & 0.8585 & 0.0535 & 0.0120 & 0.0010 & 0.0028 \\ 0.0003 & 0.0009 & 0.0079 & 0.0556 & 0.8288 & 0.0816 & 0.0060 & 0.0189 \\ 0.0001 & 0.0005 & 0.0026 & 0.0078 & 0.0591 & 0.8284 & 0.0262 & 0.0753 \\ 0.0000 & 0.0001 & 0.0006 & 0.0099 & 0.0282 & 0.0570 & 0.6247 & 0.2797 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 1.0000 \end{pmatrix}.$$

Using instead Equation (6. 10) we obtain

$$\hat{Q}_{Moody's} = \begin{pmatrix} -0.1159 & 0.1095 & 0.0061 & 0.0000 & 0.0003 & 0.0000 & 0.0000 & 0.0000 \\ 0.0128 & -0.1175 & 0.1032 & 0.0002 & 0.0009 & 0.0000 & 0.0000 & 0.0003 \\ 0.0005 & 0.0330 & -0.1066 & 0.0660 & 0.0056 & 0.0014 & 0.0001 & 0.0000 \\ 0.0007 & 0.0027 & 0.0797 & -0.1621 & 0.0690 & 0.0084 & 0.0007 & 0.0009 \\ 0.0003 & 0.0005 & 0.0039 & 0.0708 & -0.2003 & 0.1082 & 0.0054 & 0.0111 \\ 0.0001 & 0.0004 & 0.0021 & 0.0042 & 0.0773 & -0.1990 & 0.0456 & 0.0692 \\ 0.0000 & 0.0000 & 0.0000 & 0.0114 & 0.0359 & 0.0834 & -0.4748 & 0.3440 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \end{pmatrix}$$

with

$$\exp(\hat{Q}_{Moody's}) = \begin{pmatrix} 0.8912 & 0.0976 & 0.0105 & 0.0003 & 0.0003 & 0.0000 & 0.0000 & 0.0000 \\ 0.0114 & 0.8913 & 0.0925 & 0.0032 & 0.0011 & 0.0001 & 0.0000 & 0.0003 \\ 0.0006 & 0.0297 & 0.9028 & 0.0581 & 0.0069 & 0.0018 & 0.0001 & 0.0001 \\ 0.0006 & 0.0036 & 0.0701 & 0.8547 & 0.0582 & 0.0102 & 0.0008 & 0.0017 \\ 0.0003 & 0.0007 & 0.0059 & 0.0596 & 0.8241 & 0.0893 & 0.0058 & 0.0144 \\ 0.0001 & 0.0004 & 0.0022 & 0.0061 & 0.0643 & 0.8244 & 0.0329 & 0.0696 \\ 0.0000 & 0.0000 & 0.0005 & 0.0095 & 0.0285 & 0.0614 & 0.6234 & 0.2766 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 1.0000 \end{pmatrix}.$$

Computing the L^1 -norm² yields $\|M_{Moody's} - \exp(\check{Q}_{Moody's})\|_1 = 0.10373$ and $\|M_{Moody's} - \exp(\hat{Q}_{Moody's})\|_1 = 0.00206$, i.e., the generator obtained from the log series expansion seems to be a better approximation in this case. Note that some, but not all entries violating the monotony conditions have been smoothed out.

Kealhofer et al. [73] question that rating changes are a good indicator for credit quality changes. In particular, they claim that rating agencies are too slow in changing ratings and therefore the probability of staying in a given grade overstates the true probability of keeping approximately the same credit quality. Suppose firms are classified according to KMV's respective expected default frequencies (EDF), based upon non-overlapping ranges of default probabilities. Each of these ranges corresponds then to a rating class, i.e., firms with default rates less than or equal to 0.002% are mapped to AAA, 0.002% to 0.04% corresponds to AA, etc. The historical frequencies of changes from one range to another are estimated from the history of changes in default rates as measured by EDFs. This yields the following KMV one-year transition matrix.

$$M_{KMV} = \begin{pmatrix} 0.6626 & 0.2222 & 0.0737 & 0.0245 & 0.0086 & 0.0067 & 0.0015 & 0.0002 \\ 0.2166 & 0.4304 & 0.2583 & 0.0656 & 0.0199 & 0.0068 & 0.0020 & 0.0004 \\ 0.0276 & 0.2034 & 0.4419 & 0.2294 & 0.0742 & 0.0197 & 0.0028 & 0.0010 \\ 0.0030 & 0.0280 & 0.2263 & 0.4254 & 0.2352 & 0.0695 & 0.0100 & 0.0026 \\ 0.0008 & 0.0024 & 0.0369 & 0.2293 & 0.4441 & 0.2453 & 0.0341 & 0.0071 \\ 0.0001 & 0.0005 & 0.0039 & 0.0348 & 0.2047 & 0.5300 & 0.2059 & 0.0201 \\ 0.0000 & 0.0001 & 0.0009 & 0.0026 & 0.0179 & 0.1777 & 0.6995 & 0.1013 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 1.0000 \end{pmatrix}.$$

²The L^1 -norm of a matrix M is defined as $\|M\|_1 = \sum_{i,j} |m_{ij}|$.

Estimating an appropriate generator according to Equation (6. 10) leads to

$$\hat{Q}_{KMV} = \begin{pmatrix} -0.4941 & 0.4512 & 0.0043 & 0.0290 & 0.0006 & 0.0087 & 0.0002 & 0.0001 \\ 0.4503 & -1.1663 & 0.6945 & 0.0000 & 0.0180 & 0.0009 & 0.0025 & 0.0001 \\ 0.0000 & 0.5813 & -1.2204 & 0.6237 & 0.0025 & 0.0122 & 0.0000 & 0.0007 \\ 0.0217 & 0.0000 & 0.6583 & -1.3037 & 0.6119 & 0.0000 & 0.0109 & 0.0009 \\ 0.0000 & 0.0314 & 0.0000 & 0.6355 & -1.2157 & 0.5418 & 0.0000 & 0.0070 \\ 0.0013 & 0.0000 & 0.0235 & 0.0000 & 0.4805 & -0.8571 & 0.3463 & 0.0054 \\ 0.0000 & 0.0010 & 0.0000 & 0.0118 & 0.0000 & 0.3004 & -0.4276 & 0.1145 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \end{pmatrix}$$

with $\|M_{KMV} - \exp(\hat{Q}_{KMV})\|_1 = 0.5717$, whereas Equation (6. 11) results in

$$\check{Q}_{KMV} = \begin{pmatrix} -0.4116 & 0.2711 & 0.0899 & 0.0299 & 0.0105 & 0.0082 & 0.0018 & 0.0002 \\ 0.3206 & -0.8430 & 0.3823 & 0.0971 & 0.0295 & 0.0101 & 0.0030 & 0.0006 \\ 0.0404 & 0.2976 & -0.8167 & 0.3357 & 0.1086 & 0.0288 & 0.0041 & 0.0015 \\ 0.0045 & 0.0417 & 0.3366 & -0.8547 & 0.3499 & 0.1034 & 0.0149 & 0.0039 \\ 0.0012 & 0.0035 & 0.0539 & 0.3348 & -0.8117 & 0.3582 & 0.0498 & 0.0104 \\ 0.0001 & 0.0007 & 0.0053 & 0.0470 & 0.2765 & -0.6349 & 0.2781 & 0.0272 \\ 0.0000 & 0.0001 & 0.0011 & 0.0031 & 0.0213 & 0.2113 & -0.3574 & 0.1205 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \end{pmatrix}$$

with $\|M_{KMV} - \exp(\check{Q}_{KMV})\|_1 = 1.289$, again showing that the log-expansion approximation is superior in this case. We therefore conclude that, despite the awkward ad hoc transformation from \tilde{Q} to \check{Q} resp. \hat{Q} , the L^1 -error based on a log-expansion is smaller than the error based on the method proposed by Jarrow et al. for the given examples.

Adjustment of the default column Clearly, any approximate generator results in a more or less modified default column. Additionally, one might have some exogenously given default master scale and still want to represent the dynamics of rating migration by a given Markov generator. The following property allows adjusting the generator appropriately.

6.3.3 Proposition *Let Q be a 8×8 generator matrix and define $\Lambda \in \mathbb{R}^{8 \times 8}$ as*

$$(\Lambda)_{i,j} = \begin{cases} 0 & \text{if } i \neq j \\ \lambda_i > 0 & \text{if } i = j \end{cases}$$

Then ΛQ is again a generator matrix, i.e., row scaling by constant positive factors is a closed operation in the space of Q -matrices.

The proof is obvious by the properties of Q -matrices. We can now use this proposition to successively adjust the generator to reproduce a given default column according to the following algorithm:

1. Choose $\Lambda^{(1)}$ with $\lambda_1 > 0$ and $\lambda_{i \neq 1} = 1$ such that $(\exp(\Lambda^{(1)}Q))_{1,8} = m_{1,8}$.
2. Choose $\Lambda^{(2)}$ with $\lambda_2 > 0$ and $\lambda_{i \neq 2} = 1$ such that $(\exp(\Lambda^{(2)}\Lambda^{(1)}Q))_{2,8} = m_{2,8}$.
- \vdots \vdots \vdots \vdots \vdots
7. Choose $\Lambda^{(7)}$ with $\lambda_7 > 0$ and $\lambda_{i \neq 7} = 1$ such that $(\exp(\Lambda^{(7)} \dots \Lambda^{(1)}Q))_{7,8} = m_{7,8}$.
8. Scaling a row of a Q -matrix has some impact on every single entry in the exponential of that matrix. This means, for example, that after we achieved the right AAA-DP (step 1), step 2 produces the right AA-DP but slightly changes the just calibrated AAA-DP. Therefore one has to repeat steps 1-7 until the default column of the respective matrix exponential agrees with the default column of M within some error bars.

The factors λ_i as described in the algorithm can be found by a simple trial-and-error method, e.g., using dyadic approximation. The above algorithm converges due to the fact that the mappings $\lambda \mapsto \exp(\Lambda^{(i)}Q)$, $i = 1, \dots, 7$, are continuous. This follows from the power series representation (6. 8). Furthermore, from (6. 8) follows that $\exp(D_{i,\lambda}Q) \approx I + \Lambda^{(i)}Q$ indicating that multiplication of the i -th row of Q by a factor λ_i mainly affects the i -th row of the corresponding matrix exponential. For example, starting from \hat{Q}_{KMV} the modified adjusted generator \bar{Q} reads

$$\bar{Q}_{KMV} = \begin{pmatrix} -0.4965 & 0.4535 & 0.0043 & 0.0291 & 0.0006 & 0.0087 & 0.0002 & 0.0001 \\ 0.4151 & -1.0752 & 0.6403 & 0.0000 & 0.0166 & 0.0008 & 0.0023 & 0.0001 \\ 0.0000 & 0.5764 & -1.2100 & 0.6184 & 0.0025 & 0.0121 & 0.0000 & 0.0007 \\ 0.0239 & 0.0000 & 0.7253 & -1.4364 & 0.6742 & 0.0000 & 0.0120 & 0.0010 \\ 0.0000 & 0.0273 & 0.0000 & 0.5526 & -1.0571 & 0.4711 & 0.0000 & 0.0061 \\ 0.0014 & 0.0000 & 0.0253 & 0.0000 & 0.5180 & -0.9239 & 0.3733 & 0.0058 \\ 0.0000 & 0.0010 & 0.0000 & 0.0127 & 0.0000 & 0.3243 & -0.4616 & 0.1235 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \end{pmatrix}$$

with corresponding matrix exponential

$$\exp(\bar{Q}_{KMV}) = \begin{pmatrix} 0.6587 & 0.2290 & 0.0693 & 0.0256 & 0.0093 & 0.0063 & 0.0016 & 0.0002 \\ 0.2090 & 0.4482 & 0.2420 & 0.0688 & 0.0230 & 0.0064 & 0.0023 & 0.0004 \\ 0.0548 & 0.2177 & 0.4301 & 0.2025 & 0.0727 & 0.0171 & 0.0041 & 0.0010 \\ 0.0224 & 0.0736 & 0.2378 & 0.3576 & 0.2333 & 0.0589 & 0.0138 & 0.0026 \\ 0.0070 & 0.0249 & 0.0716 & 0.1915 & 0.4575 & 0.1974 & 0.0430 & 0.0071 \\ 0.0023 & 0.0077 & 0.0232 & 0.0546 & 0.2173 & 0.4754 & 0.1993 & 0.0201 \\ 0.0005 & 0.0017 & 0.0050 & 0.0125 & 0.0415 & 0.1732 & 0.6642 & 0.1013 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 1.0000 \end{pmatrix}$$

and $\|M_{KMV} - \exp(\bar{Q}_{KMV})\|_1 = 0.6855$.

Remark Before closing this section we briefly mention a theorem on the non-existence of a valid generator.

6.3.4 Theorem ([62]) *Let M be a transition matrix and suppose that either*

- (i) $\det(M) \leq 0$; or
- (ii) $\det(M) > \prod_i m_{ii}$; or
- (iii) *there exist states i and j such that j is accessible from i , i.e., there is a sequence of states $k_0 = i, k_1, k_2, \dots, k_m = j$ such that $m_{k_l k_{l+1}} > 0$ for each l , but $m_{ij} = 0$.*

Then there does not exist an exact generator.

Strictly diagonal dominance of M implies $\det(M) > 0$; so, part (i) does usually not apply for credit migration matrices (for a proof, see references). But case (iii) is quite often observed with empirical matrices. For example, $M_{Moody's}$ has zero *Aaa* default probability, but a transition sequence from *Aaa* to *D* is possible. Note that if we adjust a generator to a default column with some vanishing entries the respective states become trapped states due to the above theorem ($\exp(\bar{Q}_{Moody's})$ and $\exp(\hat{Q}_{Moody's})$ are only accurate to four decimals), i.e., states with zero default probability and an underlying Markov process dynamics are irreconcilable with the general ideas of credit migration with default as the only trapped state.

Remark Strictly diagonal dominance is a necessary prerequisite for the logarithmic power series of the transition matrix to converge [62]. Now, the default state being the only absorbing state, any transition matrix M risen to the power of some $t > 1$, M^t , loses the property of diagonal dominance, since in the limit $t \rightarrow \infty$ only the default state is populated, i.e.,

$$M(t) = M^t \rightarrow \begin{pmatrix} 0 & \dots & 0 & 1 \\ \dots & \dots & \dots & \dots \\ 0 & \dots & 0 & 1 \end{pmatrix} \quad \text{as } t \rightarrow \infty,$$

which is clearly not strictly diagonally dominant. Kreinin and Sidelnikova [75] proposed regularization algorithms for matrix roots and generators that do not rely on the property of diagonal dominance. These algorithms are robust and computationally efficient, but in the time-continuous case are only slightly advantageous when compared to the weighted adjustment. In the time-discrete case, i.e., transition matrices as matrix-roots, their method seems to be superior for the given examples to other known regularization algorithms.

6.4 Term Structure Based on Market Spreads

Alternatively, we can construct an implied default term structure by using market observable information, such as asset swap spreads or defaultable bond prices. This approach is commonly used in credit derivative pricing. The extracted default probabilities reflect the market agreed perception today about the future default tendency of the underlying credit; they are by construction risk-neutral probabilities. Yet, in some sense, market spread data presents a classic example of a joint observation problem. Credit spreads imply loss severity given default, but this can only be derived if one is prepared to make an assumption as to what they are simultaneously implying about default likelihoods (or vice versa). In practice one usually makes exogenous assumptions on the recovery rate, based on the security's seniority. In any credit-linked product the primary risk lies in the potential default of the reference entity: absent any default in the reference entity, the expected cash flow will be received in full, whereas if a default event occurs the investor will receive some recovery amount. It is therefore

natural to model a risky cash flow as a portfolio of contingent cash flows corresponding to the different scenarios weighted by the probability of these scenarios.

The time origin, $t = 0$, is chosen to be the current date and our time frame is $[0, T]$, i.e., we have market observables for comparison up to time T . Furthermore, assume that the event of default and the default-free discount factor are statistically independent. Then the present value for a risky payment X promised for time t (assuming no recovery) equals

$$B(0, t)S(t)X,$$

where $B(0, t)$ is the risk-free discount factor (zero bond prices) and $S(t)$ as usual the cumulative survival probability as of today. Consider a credit bond from an issuer with notional V , fixed coupon c , and maturity T_n , and let the accrual dates for the promised payments be $0 \leq T_1 < T_2 < \dots < T_n$. We assume that the coupon of the bond to be paid at time T_i is $c\Delta_i$ where Δ_i is the day count fraction for period $[T_{i-1}, T_i]$ according to the given day count convention. When the recovery rate REC is nonzero, it is necessary to make an assumption about the claim made by the bond holders in the event of default. Jarrow and Turnbull [65] and Hull and White [59] assume that the claim equals the no-default value of the bond. In this case value additivity is given, i.e., the value of the coupon-bearing bond is the sum of the values of the underlying zero bonds. Duffie and Singleton [30] assume that the claim is equal to the value of the bond immediately prior to default. In [60], Hull and White advocate that the best assumption is that the claim made in the event of default equals the face value of the bond plus accrued interests. Whilst this is more consistent with the observed clustering of asset prices during default it makes splitting a bond into a portfolio of risky zeros much harder, and value additivity is no longer satisfied. Here, we define recovery as a fraction of par and suppose that recovery rate is exogenously given (a refinement of this definition is made in [Chapter 7](#)), based on the seniority and rating of the bond, and the industry of the corporation. Obviously, in case of default all future coupons are lost.

The net present value of the payments of the risky bond, i.e., the

dirty price, is then given as

$$\begin{aligned} \text{dirty price} = & \sum_{T_i > 0} B(0, T_i) \Delta_i S(T_i) + \\ & + V \left[B(0, T_n) S(T_n) + REC \int_0^{T_n} B(0, t) F(dt) \right]. \end{aligned} \quad (6. 12)$$

The interpretation of the integral is just the recovery payment times the discount factor for time t times the probability to default “around” t summed up from time zero to maturity.

Similarly, for a classic default swap we have spread payments $\Delta_i s$ at time T_i where s is the spread, provided that there is no default until time T_i . If the market quotes the fair default spread s the present value of the spread payments and the event premium $V(1 - REC)$ cancel each other:

$$0 = \sum_{i=1}^n B(0, T_i) s \Delta_i S(T_i) - V(1 - REC) \int_0^{T_n} B(0, t) F(dt). \quad (6. 13)$$

Given a set of fair default spreads or bond prices (but the bonds have to be from the same credit quality) with different maturities and a given recovery rate one now has to back out the credit curve. To this end we have to specify also a riskless discount curve $B(0, t)$ and an interpolation method for the curve, since it is usually not easy to get a smooth default curve out of market prices. In the following we briefly sketch one method:

Fitting a credit curve Assuming that default is modeled as the first arrival time of a Poisson process we begin by supposing that the respective hazard rate is constant over time. Equations (6. 12) and (6. 13), together with Equation (6. 2) $S(t) = e^{-\int_0^t h(s) ds} = e^{-ht}$, allow us then to back out the hazard rate from market observed bond prices or default spreads. If there are several bond prices or default spreads available for a single name one could in principle extract a term structure of a piece-wise constant hazard rate. In practice, this might lead to inconsistencies due to data and model errors. So, a slightly more sophisticated but still parsimonious model is obtained by assuming a time-varying, but deterministic default intensity $h(t)$. Suppose, for example, that $\int_0^t h(s) ds = \Phi(t) \cdot t$, where the function $\Phi(t)$ captures term structure effects. An interesting candidate for the

fit function Φ is the Nelson-Siegel [100] yield curve function:

$$\Phi(t) = a_0 + (a_1 + a_2) \left(\frac{1 - \exp(-t/a_3)}{t/a_3} \right) - a_2 \exp(-t/a_3). \quad (6. 14)$$

This function is able to generate smooth upward sloping, humped and downward sloping default intensity curves with a small number of parameters, and, indeed, we have seen in [Figure 6.2](#) that investment grade bonds tend to have a slowly upward sloping term structure whereas those of speculative grade bonds tend to be downward sloping. Equation (6. 14) implies that the default intensity of a given issuer tends towards a long-term mean. Other functions like cubic or exponential spline may also be used in Equation (6. 14), although they might lead to fitting problems due to their greater flexibility and the frequency of data errors. The parameter a_0 denotes the long-term mean of the default intensity, whereas a_1 represents its current deviation from the mean. Specifically, a positive value of a_1 implies a downward sloping intensity and a negative value implies an upward sloping term structure. The reversion rate towards the long-term mean is negatively related to $a_3 > 0$. Any hump in the term structure is generated by a nonzero a_2 . However, in practice, allowing for a hump may yield implausible term structures due to overfitting. Thus, it is assumed that $a_2 = 0$, and the remaining parameters $\{a_0, a_1, a_3\}$ are estimated from data. The Nelson-Siegel function can yield negative default intensities if the bonds are more liquid or less risky than the “default-free” benchmark, or if there are data errors.

Using Equations (6. 2) and (6. 14) the survival function $S(t)$ can then be written as

$$S(t) = \exp \left[- \left(a_0 + a_1 \left(\frac{1 - \exp(-t/a_3)}{t/a_3} \right) \right) \cdot t \right]. \quad (6. 15)$$

Now, we construct default curves from reference bond and default swap prices as follows: Consider a sample of N constituents which can be either bonds or swaps or both. To obtain the values of the parameters of the default intensity curve, $\{a_0, a_1, a_3\}$, we fit equations (6. 12, 6. 13), and with the help of Equation (6. 15), to the market observed prices by use of a nonlinear optimization algorithm under the constraints $a_3 > 0$, $S(0) = 1$, and $S(t) - S(t + 1) \geq 0$. Mean-Absolute-Deviation regression seems to be more suitable than Least-Square regression since the former is less sensitive to outliers.

KMV's risk-neutral approach (See Crouhy et al. [21]) Under the Merton-style approach the actual cumulative default probability from time 0 to time t has been derived in a real, risk averse world as (cf. Chapter 3)

$$DP_t^{real} = N \left(-\frac{\log(A_0/C) + (\mu - \sigma^2/2)t}{\sigma\sqrt{t}} \right), \quad (6. 16)$$

where A_0 is the market value of the firm's asset at time 0, C is the firm's default point, σ the asset volatility, and μ the expected return of the firm's assets. In a world where investors are neutral to risk, all assets should yield the same risk-free return r . So, the risk-neutral default probabilities are given as

$$DP_t^{rn} = N \left(-\frac{\log(A_0/C) + (r - \sigma^2/2)t}{\sigma\sqrt{t}} \right), \quad (6. 17)$$

where the expected return μ has been replaced by the risk-free interest rate r . Because investors refuse to hold risky assets with expected return less than the risk-free base rate, μ must be larger than r . It follows that

$$DP_t^{rn} \geq DP_t^{real}.$$

Substituting Equation (6. 16) into Equation (6. 17) and rearranging, we can write the risk-neutral probability as:

$$DP_t^{rn} = N \left(N^{-1}(DP_t^{real}) + \frac{\mu - r}{\sigma} \sqrt{t} \right). \quad (6. 18)$$

From the continuous time CAPM we have

$$\mu - r = \beta\pi \quad \text{with} \quad \beta = \frac{\text{Cov}(r_a, r_m)}{\mathbb{V}(r_m)} = \rho_{a,m} \frac{\sigma}{\sigma_m}$$

as beta of the asset with the market. r_a and r_m denote the continuous time rate of return on the firm's asset and the market portfolio, σ and σ_m are the respective volatilities, and $\rho_{a,m}$ denotes the correlation between the asset and the market return. The market risk premium is given by

$$\pi = \mu_m - r$$

where μ_m denotes the expected return on the market portfolio. Putting all together leads to

$$DP_t^{rn} = N \left(N^{-1}(DP_t^{real}) + \rho_{a,m} \frac{\pi}{\sigma_m} \sqrt{t} \right). \quad (6. 19)$$

The correlation $\rho_{a,m}$ is estimated from a linear regression of the asset return against the market return. The market risk premium π is time varying, and is much more difficult to estimate statistically. KMV uses a slightly different mapping from distance-to-default to default probability than the normal distribution. Therefore, the risk-neutral default probability is estimated by calibrating the market Sharpe ratio, $SR = \pi/\sigma_m$, and θ , in the following relation, using bond data:

$$DP_t^{rn} = N[N^{-1}(DP_t^{real}) + \rho_{a,m}SRt^\theta]. \quad (6. 20)$$

From Equation (6. 12) we obtain for the credit spread s of a risky zero bond

$$e^{-(r+s)t} = [(1 - DP_t^{rn}) + (1 - LGD)DP_t^{rn}] e^{-rt}. \quad (6. 21)$$

Combining Equation (6. 20) and Equation (6. 21) yields

$$s = -\frac{1}{t} \log \left[1 - N(N^{-1}(DP_t^{real}) + \rho_{a,m}SRt^\theta)LGD \right],$$

which then serves to calibrate SR and θ in the least-square sense from market data.

Chapter 7

Credit Derivatives

Credit derivatives are instruments that help banks, financial institutions, and debt security investors to manage their credit-sensitive investments. Credit derivatives *insure* and *protect* against adverse movements in the credit quality of the counterparty or borrower. For example, if a borrower defaults, the investor will suffer losses on the investment, but the losses can be offset by gains from the credit derivative transaction. One might ask why both banks and investors do not utilize the well-established insurance market for their protection. The major reasons are that credit derivatives offer lower transaction cost, quicker payment, and more liquidity. Credit default swaps, for instance, often pay out very soon after the event of default¹; in contrast, insurances take much longer to pay out, and the value of the protection bought may be hard to determine. Finally, as with most financial derivatives initially invented for hedging, credit derivatives can now be traded speculatively. Like other over-the-counter derivative securities, credit derivatives are privately negotiated financial contracts. These contracts expose the user to operational, counterparty, liquidity, and legal risk. From the viewpoint of quantitative modeling we here are only concerned with *counterparty risk*. One can think of credit derivatives being placed somewhere between traditional credit insurance products and financial derivatives. Each of these areas has its own valuation methodology, but neither is wholly satisfactory for pricing credit derivatives. The insurance techniques make use of historical data, as, e.g., provided by rating agencies, as a basis for valuation (see [Chapter 6](#)). This approach assumes that the future will be like the past, and does not take into account market information about credit quality. In contrast, derivative technology employs market information as a basis for valuation. Derivative securities pricing is based on the assumption of *risk-neutrality* which assumes arbitrage-free and com-

¹Especially under the ISDA master agreement, cf. [61].

plete markets, but it is not clear whether these conditions hold for the credit market or not. If a credit event is based on a freely observable property of market prices, such as credit spreads, then we believe that conventional derivative pricing methodology may be applicable.

Credit derivatives are bilateral financial contracts that isolate specific aspects of credit risk from an underlying instrument and *transfer* that *risk* between two counterparties. By allowing credit risk to be freely traded, risk management becomes far more flexible. There are lots of different types of credit derivatives, but we shall only treat the most commonly used ones. They could be classified into two main categories according to valuation, namely the *replication* products, and the *default* products. The former are priced off the capacity to replicate the transaction in the money market, such as credit spread options. The latter are priced as a function of the exposure underlying the security, the default probability of the reference asset, and the expected recovery rate, such as credit default swaps. Another classification could be along their performance as protection-like products, such as credit default options and exchange-like products, such as total return swaps. In the next sections we describe the most commonly used credit derivatives and illustrate simple examples. For a more elaborate introduction to the different types of credit derivatives and their use for risk management see [68,107]; for documentation and guidelines we refer to [61].

7.1 Total Return Swaps

A total return swap (TRS) [63,97] is a mean of duplicating the cash flows of either selling or buying a reference asset, without necessarily possessing the asset itself. The TRS seller pays to the TRS buyer the total return of a specified asset and receives a floating rate payment plus a margin. The total return includes the sum of interest, fees, and any change in the value with respect to the reference asset, the latter being equal to any appreciation (positive) or depreciation (negative) in the market value of the reference security. Any net depreciation in value results in a payment to the TRS seller. The margin, paid by the TRS buyer, reflects the cost to the TRS seller of financing and servicing the reference asset on its own balance sheet. Such a transaction *transfers the entire economic benefit and risk* as well as the reference security to




FIGURE 7.1
Total return swap.

another counterparty.

A company may wish to sell an asset that it holds, but for tax or political reasons may be unable to do so. Likewise, it might hold a view that a specific asset is likely to depreciate in value in the near future, and wish to short it. However, not all assets in the market are easy to short in this way. Whatever the reason, the company would like to receive the cash flows which would result from selling the asset and investing the proceeds. This can be achieved exactly with a total return swap. Let us give an example: Bank A decides to get the economic effect of selling securities (bonds) issued by a German corporation, X. However, selling the bonds would have undesirable consequences, e.g., for tax reasons. Therefore, it agrees to swap with bank B the total return on one million 7.25% bonds maturing in December 2005 in return for a six-month payment of LIBOR plus 1.2% margin plus any decrease in the value of the bonds. [Figure 7.1](#) illustrates the total return swap of this transaction.

Total return swaps are popular for many reasons and attractive to different market segments [63,68,107]. One of the most important features is the facility to obtain an almost unlimited amount of leverage. If there is no transfer of physical asset at all, then the notional amount on which the TRS is paid is unconstrained. Employing TRS, banks can diversify credit risk while maintaining confidentiality of their client's financial records. Moreover, total return swaps can also give investors access to previously unavailable market assets. For instance, if an investor can not be exposed to Latin America market for various reasons, he or she is able to do so by doing a total return swap with a counterparty that has easy access to this market. Investors can also receive cash flows

that duplicate the effect of holding an asset while keeping the actual assets away from their balance sheet. Furthermore, an institution can take advantage of another institution's back-office and documentation experience, and get cash flows that would otherwise require infrastructure, which it does not possess.

7.2 Credit Default Products

Credit default swaps [84] are bilateral contracts in which one counterparty pays a fee periodically, typically expressed in basis points on the notional amount, in return for a contingent payment by the protection seller following a credit event of a reference security. The credit event could be either default or downgrade; the credit event and the settlement mechanism used to determine the payment are flexible and negotiated between the counterparties. A TRS is importantly distinct from a CDS in that it exchanges the total economic performance of a specific asset for another cash flow. On the other hand, a credit default swap is triggered by a credit event. Another similar product is a credit default option. This is a binary put option that pays a fixed sum if and when a predetermined credit event (default/downgrade) happens in a given time.

Let us assume that bank A holds securities (swaps) of a low-graded firm X, say BB, and is worried about the possibility of the firm defaulting. Bank A pays to firm X floating rate (Libor) and receives fixed (5.5%). For protection bank A therefore purchases a credit default swap from bank B which promises to make a payment in the event of default. The fee reflects the probability of default of the reference asset, here the low-graded firm. [Figure 7.2](#) illustrates the above transaction.

Credit default swaps are almost exclusively inter-professional transactions, and range in nominal size of reference assets from a few millions to billions of euros. Maturities usually run from one to ten years. The only true limitation is the willingness of the counterparties to act on a credit view. Credit default swaps allow users to reduce credit exposure without physically removing an asset from their balance sheet. Purchasing default protection via a CDS can hedge the credit exposure of such a position without selling for either tax or accounting purposes.




FIGURE 7.2
Credit default swap.

When an investor holds a credit-risky security, the return for assuming that risk is only the net spread earned after deducting the cost of funding. Since there is no up-front principal outlay required for most protection sellers when assuming a CDS position, they take on credit exposure in off-balance sheet positions that do not need to be funded. On the other hand, financial institutions with low funding costs may fund risky assets on their balance sheets and buy default protection on those assets. The premium for buying protection on such securities may be less than the net spread earned over their funding costs.

Modeling For modeling purposes let us reiterate some basic terminology; see [55,56]. We consider a frictionless economy with finite horizon $[0, T]$. We assume that there exists a unique martingale measure Q making all the default-free and risky security prices martingales, after renormalization by the money market account. This assumption is equivalent to the statement that the markets for the riskless and credit-sensitive debt are complete and arbitrage-free [55]. A filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{(t \geq 0)}, Q)$ is given and all processes are assumed to be defined on this space and adapted to the filtration \mathcal{F}_t (\mathcal{F}_t describes the information observable until time t). We denote the conditional expectation and the probability with respect to the equivalent martingale

measure by $\mathbb{E}_t(\cdot)$ and $Q_t(\cdot)$, respectively, given information at time t . Let $B(t, T)$ be the time t price of a default-free zero-coupon bond paying a sure currency unit at time T . We assume that forward rates of all maturities exist; they are defined in the continuous time by

$$f(t, T) = -\frac{\partial}{\partial T} \log B(t, T).$$

The default free spot rate is defined by

$$r(t) = \lim_{T \rightarrow t} f(t, T).$$

Spot rates can be modeled directly as by Cox et al. [17] or via forward rates as in Heath et al. [56]. The money market account that accumulates return at the spot rate is defined as

$$A(t) = e^{\int_0^t r(s) ds}.$$

Under the above assumptions, we can write default-free bond prices as the expected discount value of a sure currency unit received at time T , that is,

$$B(t, T) = \mathbb{E}_t \left[\frac{A(t)}{A(T)} \right] = \mathbb{E}_t \left[e^{-\int_t^T r(s) ds} \right].$$

Now, let $\tilde{B}(t, T)$ be the time t price of a credit risky zero-coupon bond promising to pay a currency unit at time T . This promised payment may not be made in full if the firm is bankrupt at time T , i.e., only a fraction of the outstanding will be recovered in the event of default. Here we assume that the event premium is the difference of par and the value of a specified reference asset after default. Let again τ represent the random time at which default occurs, with a distribution function $F(t) = \mathbb{P}[\tau \leq t]$ and $\mathbf{1}_{\{\tau < T\}}$ as the indicator function of the event. Then the price of the risky zero-coupon can be written in two ways:

$$\tilde{B}(t, T) = \mathbb{E}_t \left[e^{-\int_t^T r(s) ds} (\mathbf{1}_{\{\tau > T\}} + REC(T) \mathbf{1}_{\{\tau < T\}}) \right] \quad (7.1)$$

$$= \mathbb{E}_t \left[e^{-\int_t^T r(s) ds} \mathbf{1}_{\{\tau > T\}} + e^{-\int_t^\tau r(s) ds} REC(\tau) \mathbf{1}_{\{\tau < T\}} \right]. \quad (7.2)$$

In the first expression the recovery rate $REC(T)$ is thought of as a payout received at maturity, whereas in the second expression, we think of $REC(\tau)$ as the payment made at the time of default. Given the

existence of the money market account, we can easily translate from one representation of the recovery to the other by

$$REC(T) = REC(\tau)e^{\int_{\tau}^T r(s)ds}.$$

A credit default swap now has a default leg and a premium leg. The present value of the contingent payment $1 - REC(\tau)$ is then

$$A_{def,t} = \mathbb{E}_t \left[e^{-\int_t^T r(u)du} (1 - REC(\tau)) \mathbf{1}_{\{\tau < T\}} \right].$$

The present value of the spread payments s is given by:

$$A_{fee,t} = s \mathbb{E}_t \left[e^{-\int_t^T r(u)du} \mathbf{1}_{\{\tau > T\}} \right].$$

From arbitrage-free arguments the value of the swap should be zero when it is initially negotiated. In the course of time its present value from the protection buyer’s point of view is $A_{def,t} - A_{fee,t}$. In order to calculate the value of the CDS, it is required to estimate the survival probability, $S(t) = 1 - F(t)$, and the recovery rates $REC(t)$.

Swap premiums are typically due at prespecified dates and the amount is accrued over the respective time interval. Let $0 \leq T_0 \leq T_1 \leq \dots T_n$ denote the accrual periods of the default swap, i.e., at time $T_i, i \geq 1$ the protection buyer pays $s\Delta_i$, where Δ_i is the day count fraction for period $[T_{i-1}, T_i]$, provided that there is no default until time T_i . Assuming furthermore a deterministic recovery rate at default, $REC(\tau) = REC$, and no correlation between default and interest rates we arrive at

$$A_{def,t} = (1 - REC) \int_{T_0}^{T_n} B(T_0, u) F(du) \tag{7. 3}$$

$$A_{fee,t} = \sum_{i=1}^n s\Delta_i B(T_0, T_i) (1 - F(T_i)). \tag{7. 4}$$

The integral describes the present value of the payment $(1 - REC)$ at the time of default. For a default “at” time u , we have to discount with $B(T_0, u)$ and multiply with the probability $F(du)$ that default happens “around” u .

In some markets a plain default swap includes the features of paying the accrued premium at default, i.e., if default happens in the period (T_{i-1}, T_i) the protection buyer is obliged to pay the already accrued

part of the premium payment. In this case the value of the premium leg changes to

$$A_{fee,t} = \sum_{i=1}^n s \left[\Delta B(T_0, T_i)(1 - F(T_i)) + \int_{T_{i-1}}^{T_i} (u - T_{i-1})B(T_0, u)F(du) \right], \tag{7. 5}$$

where the difference $u - T_{i-1}$ is according to the given day count convention.

Both reduced-form models (intensity models) and structural models can in principle be applied to price default swaps. In the reduced-form model framework the relation between the intensity process h_t and the random survival probabilities at future times t provided $\tau > t$ is given by

$$q(t, T) = \mathbb{P}[\tau > T | \mathcal{F}_t] = \mathbb{E}_t \left(e^{-\int_t^T h(s)ds} \right).$$

If we assume a deterministic recovery rate REC and understand the recovery as a fraction of a corresponding riskless zero with the same maturity, we can write the price for a risky zero bond (7. 1) as (on $\{\tau > t\}$):

$$\begin{aligned} \tilde{B}(t, T) = & REC \mathbb{E}_t \left(e^{-\int_t^T r(s)ds} \right) \\ & + (1 - REC) \mathbb{E}_t \left(e^{-\int_t^T (r(s)+h(s))ds} \right). \end{aligned} \tag{7. 6}$$

In the case of zero correlation between the short rate and the intensity process both processes in the exponent would factorize when taking the expectation value. But a really sophisticated default swap model would call for correlated default and interest rates, which leads us beyond the scope of this presentation. Instead, we turn in the following section back to correlated defaults and their application to basket swaps.

7.3 Basket Credit Derivatives

Basket default swaps are more sophisticated credit derivatives that are linked to several underlying credits. The standard product is an insurance contract that offers protection against the event of the k th

default on a basket of n , $n \geq k$, underlying names. It is similar to a plain default swap but now the credit event to insure against is the event of the k th default and not specified to a particular name in the basket. Again, a premium, or spread, s is paid as an insurance fee until maturity or the event of k th default. We denote by s^{kth} the fair spread in a k th-to-default swap, i.e., the spread making the value of this swap equal to zero at inception.

If the n underlying credits in the basket default swap are independent, the fair spread s^{1st} is expected to be close to the sum of the fair default swap spreads s_i over all underlyings $i = 1, \dots, n$. If the underlying credits are in some sense “totally” dependent the first default will be the one with the worst spread; therefore $s^{1st} = \max_i(s_i)$.

The question is now how to introduce dependencies between the underlying credits to our model. Again, the concept of copulas as introduced in Section 2.6 can be used, and, to our knowledge, Li [78,79] was the first to apply copulas to valuing basket swaps by generating correlated default times as random variables via a correlation model and a credit curve. For more on copulas we refer to Section 2.6 and the literature referenced there, but see also Embrechts et al. [34] for possible pitfalls.

Modeling Dependencies via Copulas Denote by τ_i , $i = 1, \dots, n$ the random default times for the n credits in the basket, and let furthermore $(F_i(t))_{t \geq 0}$ be the curve of cumulative (risk-neutral) default probabilities for credit i :

$$F_i(t) = \mathbb{P}[\tau_i \leq t], \quad t \geq 0,$$

with $S_i(t) = \mathbb{P}[\tau_i > t] = 1 - F_i(t)$. $F(t)$ is assumed to be a strictly increasing function of t with $F(0) = 0$ and $\lim_{t \rightarrow \infty} F(t) = 1$. This implies the existence of the quantile function $F^{-1}(x)$ for all $0 \leq x \leq 1$. From elementary probability theory we know that for any standard uniformly distributed U ,

$$U \sim U(0, 1) \quad \Rightarrow \quad F^{-1}(U) \sim F. \tag{7. 7}$$

This gives a simple method for simulating random variates with distribution F , i.e., random default times in our case. The cash flows in a basket default swap are functions of the whole random vector (τ_1, \dots, τ_n) , but in order to model and evaluate this basket swap we

need the joint distribution of the τ_i 's:

$$\mathbf{F}(t_1, \dots, t_n) = \mathbb{P}[\tau_1 \leq t_1, \dots, \tau_n \leq t_n].$$

Similarly, we define the multivariate survival function \mathbf{S} by

$$\mathbf{S}(t_1, \dots, t_n) = \mathbb{P}[\tau_1 > t_1, \dots, \tau_n > t_n].$$

Note that

$$\begin{aligned} S_i(t_i) &= \mathbf{S}(0, \dots, 0, t_i, 0, \dots, 0), \\ \mathbf{S}(t_1, \dots, t_n) &\neq 1 - \mathbf{F}(t_1, \dots, t_n), \quad \text{in general,} \\ \text{but } \mathbf{S}(t, \dots, t) &= 1 - \mathbf{F}(t, \dots, t). \end{aligned}$$

We exploit again the concept of copula function where, for uniform random variables, U_1, U_2, \dots, U_n ,

$$C(u_1, u_2, \dots, u_n) = \mathbb{P}[U_1 \leq u_1, U_2 \leq u_2, \dots, U_n \leq u_n]$$

defines a joint distribution with uniform marginals. The function $C(u_1, u_2, \dots, u_n)$ is called a Copula function. Remember that $U_i = F_i(\tau_i)$ admits a uniform distribution on the interval $[0, 1]$; so, the joint distribution of (τ_1, \dots, τ_n) can be written as:

$$\mathbf{F}(t_1, \dots, t_n) = C(F_1(t_1), \dots, F_n(t_n)). \tag{7. 8}$$

Hence, the Copula function introduces a mutual correlation by linking univariate marginals to their full multivariate distribution thereby separating the dependency structure C , i.e., the ingredients are some credit curve for each credit as marginal distribution functions for the default times and a suitable chosen copula function. Observe that by Sklar's theorem (Section 2.6) any joint distribution can be reduced to a copula and the marginal distributions, although it may be difficult to write down the copula explicitly.

One of the most elementary copula functions is the multivariate normal distribution

$$C(u_1, u_2, \dots, u_n) = N_n [N^{-1}(u_1), N^{-1}(u_2), \dots, N^{-1}(u_n); \Sigma] \tag{7. 9}$$

where N_n is as before the cumulative multivariate normal distribution with correlation matrix Σ and N^{-1} is the inverse of a univariate normal distribution. Clearly, there are various different copulas generating all

kinds of dependencies, and the choice of the copula entails a significant amount of model risk [45,47]. The advantage of the normal copula, however, is that, as we have seen in Chapter 2, it relates to the latent variable approach to model dependent default. Assume that the default event of credit i up to time T is driven by a single random variable r_i (ability-to-pay variable) being below a certain threshold $c_i(T)$:

$$\tau_i < T \Leftrightarrow r_i < c_i(T).$$

If the Z_i 's admit a multivariate standard normal distribution with correlation matrix $\tilde{\Sigma}$, then to be consistent with our given default curve, we set $c_i(T) = N^{-1}(F_i(T))$. The pairwise joint default probabilities are now given in both representations by:

$$\begin{aligned} \mathbb{P}[\tau_i \leq T, \tau_j \leq T] &= \mathbb{P}[r_i \leq c_i(T), r_j \leq c_j(T)] \\ &= N_2[N^{-1}(F_i(T)), N^{-1}(F_j(T)); \tilde{\Sigma}_{ij}] \end{aligned} \quad (7. 10)$$

We see that these probabilities (7. 10) only coincide with those from the normal copula approach (7. 8), (7. 9), if the asset correlation matrix $\tilde{\Sigma}$ and the correlation matrix Σ in the normal copula are the same. But note that since the asset value approach can only model defaults up to a single time horizon T , the calibration between the two models can only be done for one fixed horizon. So, we see again that the factor model approach to generate correlated defaults based on standard normal asset returns is tantamount to a normal copula approach.

Remark Analogously to the default distribution we can apply Sklar's theorem to the survival function, i.e., when \mathbf{S} is a multivariate survival function with margins S_1, \dots, S_n , then there exists a copula representation

$$\mathbf{S}(t_1, \dots, t_n) = \check{C}(S_1(t_1), \dots, S_n(t_n)). \quad (7. 11)$$

There is an explicit, although rather complex relation between the survival copula \check{C} and the distribution copula C [50]; in the two-dimensional case we obtain

$$\begin{aligned} \check{C}(u_1, u_2) &= \mathbf{S}(S_1^{-1}(u_1), S_2^{-1}(u_2)) = \mathbf{S}(t_1, t_2) \\ &= 1 - F_1(t_1) - F_2(t_2) + \mathbf{F}(t_1, t_2) \\ &= S_1(t_1) + S_2(t_2) - 1 + C(1 - S_1(t_1), 1 - S_2(t_2)) \\ &= u_1 + u_2 - 1 + C(1 - u_1, 1 - u_2), \end{aligned}$$

where it can easily be shown that \check{C} is indeed a copula function. At this point let us state that a copula is radially symmetric if and only if $C = \check{C}$ (proof [50]). The normal copula is radial symmetric; so, e.g., in two dimensions we find indeed that

$$\begin{aligned} \check{C}(u_1, u_2) &= u_1 + u_2 - 1 + C(1 - u_1, 1 - u_2) \\ &= u_1 + u_2 - 1 + N_2 [N^{-1}(1 - u_1), N^{-1}(1 - u_2); \Sigma] \\ &= N_2 [N^{-1}(u_1), +\infty; \Sigma] + N_2 [+ \infty, N^{-1}(u_2); \Sigma] \\ &\quad - N_2 [+ \infty, + \infty; \Sigma] + N_2 [-N^{-1}(u_1), -N^{-1}(u_2); \Sigma] \\ &= N_2 [N^{-1}(u_1), N^{-1}(u_2); \Sigma] \\ &= C(u_1, u_2). \end{aligned}$$

This property is very interesting for computational purposes, since in the radially symmetric case it is thus equivalent to work with the distribution copula or with the survival copula.

Summarizing, the normal copula function approach for modeling correlated default times is as follows (Figure 7.3):

- Specify the cumulative default time distribution F_i (credit curve), such that $F_i(t)$ gives the probability that a given asset i defaults prior to t .
- Assign a standard normal random variable r_i to each asset, where the correlation between distinct r_i and r_j is ρ_{ij} .
- Obtain the default time τ_i for asset i through

$$\tau_i = F_i^{-1} (N(r_i)) .$$

Note that since $F_i(t)$ is a strictly increasing continuous function with $\lim_{t \rightarrow \infty} F_i(t) = 1$ there is always a default time, though it may be very large.

In the one-period case, positively correlated defaults mean that if one asset defaulted it is more likely that the second defaults as well, compared two independent defaults. For default times, a positive correlation means that the time between the two default events is smaller, on average, than if they were uncorrelated. Figure 7.4 depicts the average standard deviation of default times τ_i , $1 \leq i \leq 5$, $\langle \text{std}_i[\tau_i] \rangle$ (the average is taken over numerous scenarios), in units of the average default




FIGURE 7.3
 Generating correlated default times via the copula approach.

time, $\langle \text{mean}_i[\tau_i] \rangle$, the average first-to-default-time $\langle \min_i[\tau_i] \rangle$, and the average last-to-default-time $\langle \max_i[\tau_i] \rangle$, for a uniform basket of five loans in dependence of the asset correlation with cumulative multi-year default probabilities as in Table 7.1.

TABLE 7.1: Term structure of cumulative default probability.

year	1	2	3	4	5	6	7
DP	0.0071	0.0180	0.0320	0.0484	0.0666	0.0859	0.1060

year	8	9	10
DP	0.1264	0.1469	0.1672

Pricing In order to price basket default swaps, we need the distribution of the time τ^{kth} of the k th default. The k th default time is in fact the order statistic $\tau_{(k:n)}$, $k \leq n$, and in general, we have for the distribution functions

$$S_{(k:n)}(t) = 1 - F_{(k:n)}(t).$$

The distribution of the first order statistic $\tau_{(1:n)}$ is

$$F_{(1:n)}(t) = \mathbb{P}[\tau_{(1:n)} \leq t] = 1 - \mathbb{P}[\tau_1 > t, \dots, \tau_n > t] = 1 - \mathbf{S}(t, \dots, t),$$

and the one of the last order statistic (the time of the last default) is obviously

$$F_{(n:n)}(t) = \mathbb{P}[\tau_1 \leq t, \dots, \tau_n \leq t] = \mathbf{F}(t, \dots, t). \tag{7. 12}$$

The corresponding formulas for the other distribution function $F_{(k:n)}$ in terms of the copula function are much more involved (see [50]); we only state the special cases $n = 2, 3$:

n=2:

$$\begin{aligned} \mathbf{F}_{(2:2)}(t) &= C(F_1(t), F_2(t)) \\ \mathbf{F}_{(1:2)}(t) &= F_1(t) + F_2(t) - C(F_1(t), F_2(t)) \end{aligned}$$




FIGURE 7.4

The average standard deviation of the default times $\langle \text{std}_i[\tau_i] \rangle$ (\diamond), the average first-to-default time $\langle \min_i[\tau_i] \rangle$ ($+$), and the average last-to-default-time $\langle \max_i[\tau_i] \rangle$ (\circ) in units of the average default time $\langle \text{mean}_i[\tau_i] \rangle$ for a uniform basket of five loans in dependence of the asset correlation for normal distributed (solid) and t-distributed (dashed) latent variables.

$n=3$:

$$\begin{aligned}
 \mathbf{F}_{(3:3)}(t) &= C(F_1(t), F_2(t), F_3(t)) \\
 \mathbf{F}_{(2:3)}(t) &= C(F_1(t), F_2(t)) + C(F_1(t), F_3(t)) + C(F_2(t), F_3(t)) \\
 &\quad - 2C(F_1(t), F_2(t), F_3(t)) \\
 \mathbf{F}_{(1:3)}(t) &= F_1(t) + F_2(t) + F_3(t) - C(F_1(t), F_2(t)) - C(F_1(t), F_3(t)) \\
 &\quad - C(F_2(t), F_3(t)) + C(F_1(t), F_2(t), F_3(t))
 \end{aligned}$$

The fair spread s^{kth} for maturity T_m is then given by (compare Equations (7. 3), (7. 4))

$$\begin{aligned}
 0 &= s^{kth} \sum_{i=1}^m \Delta_i B(T_0, T_i) S_{(k:n)}(T_i) \\
 &\quad - \sum_{i=1}^n (1 - REC_i) \int_{T_0}^{T_m} B(T_0, u) F_{(k:n)}^{kth=i}(du). \quad (7. 13)
 \end{aligned}$$

The first part is the present value of the spread payments, which stops at τ^{kth} . The second part is the present value of the payment at the time of the k th default. Since the recovery rates might be different for the n underlying names, we have to sum up over all names and weights with the probability that the k th default happens around u and that the k th name is just i . (We assume that there are no joint defaults at exactly the same time.) So $F_{(k:n)}^{kth=i}$ is the probability distribution of the k th order statistic of the default times and that $kth = i$. Figure 7.5 show the k th-to-default spreads for a basket of three underlyings with fair spreads $s_1 = 0.009$, $s_2 = 0.010$, and $s_3 = 0.011$, and pair-wise equal correlation. Schmidt and Ward [110] already observed that the sum of the k th-to-default swap spreads is greater than the sum of the individual spreads, i.e., $\sum_{k=1}^n s^{kth} > \sum_{i=1}^n s_i$. Both sides insure exactly the same risk; so, this discrepancy is due to a windfall effect of the first-to-default swap. At the time of the first default one stops paying the huge spread s^{1st} on the one side but on the plain-vanilla side one stops just paying the spread s_i of the first default i . Of course this mismatch is only a superficial one, since the sums of the present values of the spreads on both sides are equal. Note also the two extreme cases. For fully correlated underlyings, $\rho = 1$, the first-to-default spread is the worst of all underlyings. Of course in the normal copula framework

perfect linear correlation means that the state variables are identical and that the name with the largest default probability dominates all others (assuming the same recovery rates for all underlyings). On the other hand, for $\rho = 0$, from an arbitrage-free argument one can show that the first-to-default spread is close to the sum of the individual spreads. If the correlation is greater than zero the underlying names are dependent, which entails a spread widening of the remaining names as a consequence of the default of credit i . Schmidt and Ward [110] investigated how this implied spread widening is reflected in the copula approach and found that given a flat correlation structure the size of the spread widenings depends on the quality of the credit first defaulting, i.e., the less riskier the defaulting name the larger the impact. Also the implied spread widening admits a pronounced term structure: the earlier the first default, the larger the impact on the remaining spreads.

Counterparty Risk So far, we have tacitly ignored the counterparty risk of the protection seller to default. This feature could also be dealt with in the context of the copula approach (but see also Hull and White [58] for another approach). For simplicity we reduce the problem to a single obligor CDS, but the generalization to baskets is straightforward. We now have the additional risk that the protection seller, i.e., the swap counterparty, can default, together with the reference security. So, instead of making the promised payments $1 - REC$ in the event of the reference default, only a fraction δ of that payment is recovered by the protection buyer. The formulas for the default leg (7. 3) and the premium leg (7. 4) change then, in informal notation, to

$$\begin{aligned}
 A_{def,t} &= \int_{T_0}^{T_n} B(T_0, u)(1 - REC(u))(-\mathbf{S}(du, u)) \\
 &\quad + \int_{T_0}^{T_n} B(T_0, u)(1 - REC(u))\delta(u)\mathbf{F}(du, u) \\
 &= \int_{T_0}^{T_n} B(T_0, u)(1 - REC(u)) [(1 - \delta(u))(-\mathbf{S}(du, u)) \\
 &\quad \quad \quad + \delta(u)F_{ra}(du)] \\
 A_{fee,t} &= \sum_{i=1}^n s\Delta_i B(T_0, T_i)\mathbf{S}(T_i, T_i),
 \end{aligned}$$

where F_{ra} denotes the default curve of the reference asset. $-\mathbf{S}(du, u) = -\partial_1\mathbf{S}(u, u)du$ is the probability that the reference asset defaults be-




FIGURE 7.5

*k*th-to-default spread versus correlation for a basket with three underlyings: (solid) s^{1st} , (dashed) s^{2nd} , (dashed-dotted) s^{3rd} .




FIGURE 7.6

Default spread versus correlation between reference asset and swap counterparty: (solid) $\delta = 0.2$, (dashed) $\delta = 0$ as fraction of recovery payment made.

tween u and $u + du$ while the swap counterparty is still alive, whereas $F(du, u) = \partial_1 F(u, u)du$ is the probability that the reference asset defaults between u and $u + du$ and the counterparty has already defaulted. The bivariate survival function S can then again be represented by a copula function \check{C} . Figure 7.6 shows the spread of a single-asset default swap as a function of the correlation of the reference asset to the swap counterparty. The risk-free rate is $r = 4\%$, the hazard rates are supposed to be constant at $\lambda_{ra} = 0.011$ and $\lambda_{cp} = 0.01$, and the recovery rate is at $REC = 0.2$.

Remark Obviously, the normal distribution is only one choice for a possible copula function. See, for example, [34,45] for possible pit-

falls in modeling dependencies via copula functions. For comparison, we also compute default times with t-distributed latent variables for the uniform basket of five loans (see also [83] for a use of t-copulas in modeling default baskets). Choosing identical linear correlations, ρ_n , ρ_t in the normal- and t-copula function ($\nu = 5$ degrees of freedom), produces the results in Figure 7.4. The standard deviations and the maxima (last-to-default) of the default times in the t-distributed case are slightly lower over the full range of linear correlation than the ones based on normal distributed variables. The first-to-default times are slightly higher in the t-distributed case than in the normal distributed one. Next, we calibrated for some cases the linear correlation parameters in the t- and normal distributed case to match the one-year default correlation, based on a one-year default probability of $DP_1 = 0.0071$. Note that this is not generally possible. The following table shows that now the differences are much larger.

$DP_1 = 0.0071$	$\langle \text{std}_i[\tau_i] \rangle$	$\langle \min_i[\tau_i] \rangle$	$\langle \max_i[\tau_i] \rangle$
$\rho_t = 0.1, \nu = 5$ $\rho_n = 0.455$	0.75 0.61	0.27 0.38	2.07 1.85
$\rho_t = 0.2, \nu = 5$ $\rho_n = 0.52$	0.71 0.56	0.31 0.42	2.00 1.79
$\rho_t = 0.1, \nu = 10$ $\rho_n = 0.32$	0.76 0.67	0.27 0.33	2.11 1.97
$\rho_t = 0.2, \nu = 10$ $\rho = 0.393$	0.72 0.64	0.30 0.36	2.03 1.9

7.4 Credit Spread Products

Credit spread is the difference between the yield on a particular debt security and a benchmark yield, usually on a government bond. A credit spread option (CSO) [84], say call, is an option that pays an amount equal to $\max(S(t) - K, 0)$ at maturity, where $S(t)$ is the credit spread and K is the strike price. Likewise, a put CSO pays $\max(K - S(t), 0)$. A call (put) CSO buyer benefits from a increase (decrease) in the credit spread. One of the main characteristics of these products is that the return is not dependent on a specific credit event. It merely depends on the value of one reference credit spread against another.

If the reference asset owner's credit rating goes down, and therefore the default probability increases, the credit spread goes up and vice versa. A debt issuer can make use of credit spread call options to hedge against a rise in the average credit spread. On the other hand, a financial institution that holds debt securities can purchase CSO puts to hedge against a fall in the credit spread.

Credit spread derivatives are priced by means of a variety of models. One can value them by modeling the spread itself as an asset price. The advantage of this approach is its relative simplicity. Longstaff and Schwartz [81] developed a simple framework for pricing credit spread derivatives, which we will summarize in the following. It captures the major empirical properties of observed credit spreads. They use this framework to derive closed-form solutions for call and put CSOs.

Let x denote the logarithm of the credit spreads, that is $x_t = \log(S(t))$. We assume that x is given by the SDE

$$dx = (a - bx)dt + sdB_1,$$

where a, b, s are parameters and B_1 is a Wiener process. This implies that changes in x are mean-reverting and homoscedastic, which is consistent with the empirical data. We assume that the default-free term structure is determined by a one-factor-model [124], that is

$$dr = (\alpha - \beta r)dt + \sigma dB_2.$$

Again α, β, σ_2 are parameters and B_2 is a Wiener process. The correlation coefficient between dB_1 and dB_2 is $\hat{\rho}$. Let us assume market prices of the risk premium are incorporated into a and α . Thus, both a and α are risk-adjusted parameters rather than empirical ones. This assumption is consistent with Vasicek [124] and Longstaff and Schwartz [80]. The risk-adjusted process for x is given by [81]

$$dx = \left[a - bx - \frac{\hat{\rho}\sigma s}{\beta} \left(1 - e^{-\beta(T-t)} \right) \right] dt + sdB_1. \quad (7. 14)$$

This SDE in (7. 14) can be solved by making a change of variables and then integrating. The resulting solution implies that x_T is conditionally normally distributed with respect to (7. 14) with mean μ and variance η^2 , where

$$\mu = e^{-bT}x + \frac{1}{b} \left(\alpha - \frac{\hat{\rho}\sigma s}{\beta} \right) [1 - e^{-bT}] + \frac{\hat{\rho}\sigma s}{\beta(b + \beta)} [1 - e^{-(b+\beta)T}]$$

$$\eta^2 = \frac{s^2[1 - e^{-2bT}]}{2b}.$$

Note that as $T \rightarrow \infty$ the values of μ and η^2 converge to fixed values, and the distribution of x_T converges to a steady-state stationary distribution. With this framework we can find the price of a European call CSO. Let $C(x, r, T)$ denote the value of the option. The payoff function for this option is simply $H(x) = \max(e^x - K, 0)$. The closed-form solution for the call CSO is given by

$$C(x, r, T) = p(r, T) \left[e^{\mu + \eta^2/2} N(d_1) - KN(d_2) \right].$$

Here, $N(\cdot)$ is the cumulative standard normal distribution, $p(r, T)$ is a riskless discount bond, and

$$d_1 = \frac{-\log(K) + \mu + \eta^2}{\eta}, \quad d_2 = d_1 - \eta.$$

The value of a European put CSO is

$$P(x, r, T) = C(x, r, T) + p(r, T) \left[K - e^{\mu + \eta^2/2} \right].$$

The option formula has some similarities with the Black-Scholes option pricing formula. However, the value of a call option can be less than its intrinsic value even when the call is only slightly in the money. This surprising result is due to the mean reversion of the credit spreads. When the spread is above the long-run mean, it is expected to decline over time. This can not happen in the B-S model because the underlying asset must appreciate like the riskless rate in the risk-neutral framework. The delta for a call is always positive, as in the B-S framework, but the delta of a CSO call decreases to zero as the time until the expiration increases. A change in the current credit spread is heavily outweighed by the effects of mean reversion if the expiration date of the call is far in the future.

A credit spread collar combines a credit spread put and a credit spread call. An investor that wishes to insure against rising credit spreads by buying a credit spread call can reduce the cost by selling a credit spread put. In a credit spread forward (CSF), counterparty A pays at time T a pre-agreed fixed payment and receives the credit spread of the reference asset at time T . Conversely, counterparty B receives the fee and pays the credit spread. The fixed payment is chosen at time $t < T$ to set the initial value of the credit spread forward to zero. The credit spread forward can also be structured around the relative credit spread between two different defaultable bonds. Credit




FIGURE 7.7
Credit spread swap.

spread forwards can be combined to a credit spread swap (Figure 7.7) in which one counterparty pays periodically the relative credit spread, $(S_1(t) - S_2(t))$, to the other.

7.5 Credit-linked Notes

Credit-linked notes exist in various forms in the credit derivatives market; see [23,98,68,26]. In its most common form, a credit-linked note (CLN) is a *synthetic bond* with an *embedded default swap* as illustrated in Figure 7.8.

CLNs are initiated in several ways. In the following we outline four examples of typical CLN structures.

The first case we present is the situation of an (institutional) investors who wants to have access to a credit exposure (the *reference asset*) for which by policy, regulation, or other reasons he has no direct access. In such cases, a CLN issued by another institution (the *issuer*) which has access to this particular credit exposure offers a way to evade the problems hindering the investor to purchase the exposure he is interested in. The issuer sells a note to the investor with underlying exposure equal to the face value of the reference asset. He receives the face value of the reference asset as cash proceeds at the beginning of the transaction and in turn pays interest, including some premium for the default risk, to the investor. In case the reference asset experiences a *credit event*, the issuer pays to the investor the recovery proceeds of the reference asset. The spread between the face value and the recovery value of the reference asset is the investor's *exposure at risk*. In case no credit

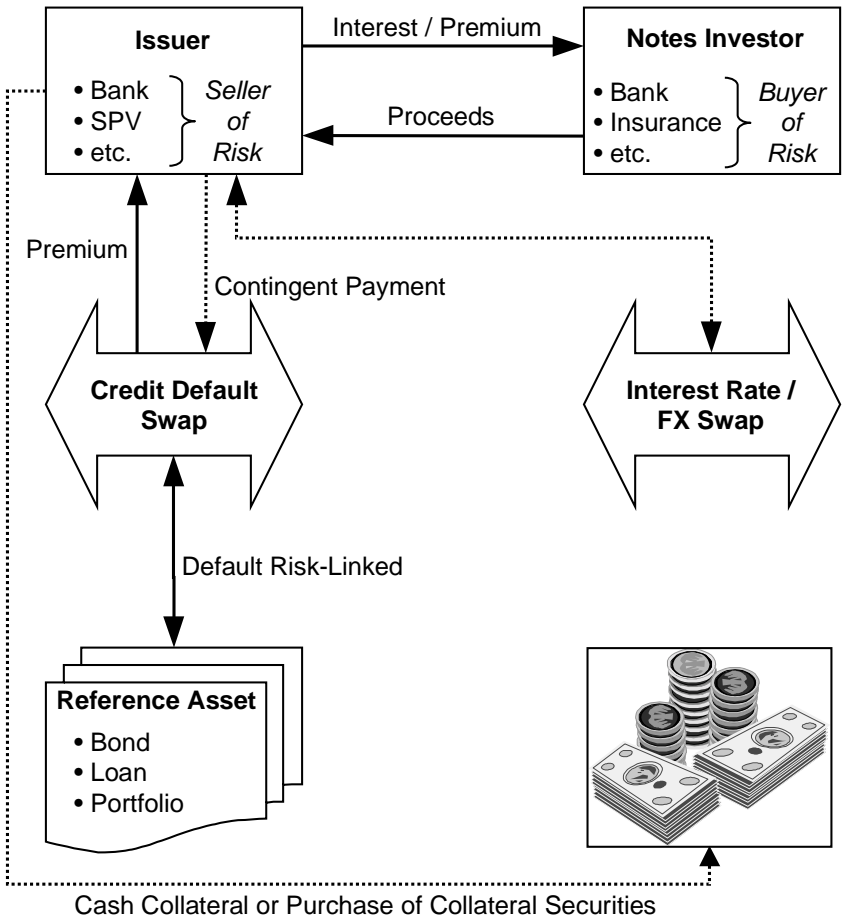


FIGURE 7.8
Example of a Credit-linked Note.

event occurred during the lifetime of the reference note, the issuer pays the full principal back to the investor. So in this example one could summarize a CLN as a *synthetic bond* with an *embedded default swap*.

In our second example, an investor, who has no access to the credit derivatives market or is not allowed to do off-balance sheet transactions, wants to invest in a credit default swap, selling protection to the owner of some reference asset. This can be achieved by investing in a CLN in the same way as described in our first example. Note that from the investor's point of view the CLN deal differs from a default swap agreement by the cash payment made upfront. In a default swap, no principal payments are exchanged at the beginning.

Another common way to set-up a CLN is *protection* buying. Assume that a bank is exposed to the default risk of some reference asset. This could be the case by means of an asset on the balance sheet of the bank or by means of a situation where the bank is the protection seller in a credit default swap. In both cases the bank has to carry the reference asset's default risk; see [Figure 7.8](#). The bank can now issue a CLN to some investor who pays the exposure of the reference asset upfront in cash to the bank and receives interest, including some premium reflecting the riskiness of the reference asset, during the lifetime of the note. If the reference asset defaults, the bank suffers a loss for its balance sheet asset (funded case) or has to make a contingent payment for the default swap (unfunded case). The CLN then compensates the bank for the loss, such that the CLN functions as an *insurance*.

In this example, the difference between a CLN and just another default swap arises from the *cash proceeds* the bank receives *upfront* from the CLN investor. As a consequence, the bank is *not exposed to the counterparty risk* of the protection selling investor. Therefore, the credit quality of the investor is of no relevance². The proceeds from the CLN can be kept as a cash collateral or be invested in high-quality collateral securities, so that losses on the reference asset will be covered *with certainty*.

Our last example refers to [Chapter 8](#), where CLNs will be discussed as notes issued by a *special purpose vehicle* (SPV) in order to set-up a *synthetic CDO*. In this case, CLNs are used for the exploitation of *regulatory arbitrage* opportunities and for *synthetic risk transfer*.

²Of course, for a short time at the start of the CLN their could be a *settlement risk*.

Besides the already mentioned reasons, there are certainly more advantages of CLNs worthwhile to be mentioned. For example, CLNs do not require an ISDA master agreement, but rather can contractually rely on the term sheet of the notes. Another advantage of CLNs is that not only the investor's credit quality but also his *correlation* with the reference asset is of no relevance to the CLN, because the money for the protection payment is delivered upfront. This concludes our discussion of credit derivatives.

Chapter 8

Collateralized Debt Obligations

Collateralized debt obligations constitute an important class of so-called *asset backed securities* (ABS), which are securities backed by a pool of assets. Depending on the underlying asset class, ABS include various subclasses, for example *residential* or *commercial mortgage backed securities* (RMBS, CMBS), *trade receivables* ABS, *credit card* ABS (often in the form of so-called *CC Master Trusts*), *consumer loan* ABS, and so on. Not long ago, it started that ABS were also structured based on pools of derivative instruments, like *credit default swaps*, resulting in a new ABS class, so-called *collateralized swap obligations* (CSO). In general, one could say that ABS can be based on *any* pool of assets generating a cash flow suitable for being structured in order to meet investor's risk preferences.

As an important disclaimer at the beginning of this chapter, we must say that a deeper treatment of ABS would easily justify the beginning of a new book, due to the many different structures and asset classes involved in the ABS market. Moreover, it is almost impossible to capture the “full” range of products in the ABS market, because in most cases a new transaction will also carry new innovative cash flow elements or special features. Therefore, ABS transactions have to be considered on a careful case-by-case analysis.

This chapter presents an introduction to CDO modeling in a “storyline” style. Some references for further reading are given in the last section of this chapter.

8.1 Introduction to Collateralized Debt Obligations

Figure 8.1 shows a segmentation of the CDO market. There are basically two types of debt on which CDOs are based, namely *bonds* and *loans*, constituting

- *Collateralized bond obligations* (CBO):
In this case, the collateral pool contains credit risky bonds. Many of the CBOs we currently find in the market are motivated by *arbitrage spread* opportunities, see [Section 8.2.1](#).
- *Collateralized loan obligations* (CLO):
Here the collateral pool consists of loans. *Regulatory capital relief*, *cheaper funding*, and, more general, *regulatory arbitrage* combined with *economic risk transfer* are the major reasons for the *origination* of CLOs by banks all over the world, see [Section 8.2.1](#).

Besides these two, CSOs (see the introductory remarks) are of increasing importance. Their advantage is the reduction of funding costs, because instead of *funded* instruments like loans or bonds, the cash flows from credit derivatives are structured in order to generate an attractive arbitrage spread. A second advantage of CSOs is the fact that credit derivatives are *actively traded* instruments, such that, based on the *fair market spread* of the collateral instruments, a *fair price* of the issued securities can be determined, for example, by means of a *risk-neutral valuation* approach.

Another class of CDOs gaining much attention are *multisector CDOs*. In this case, the collateral pool is a *mixture* of different ABS bonds, high-yield bonds or loans, CDO pieces, mortgage-backed securities, and other assets. Multisector CDOs are more difficult to analyze, mainly due to *cross-collateralization* effects, essentially meaning that bonds issued by a distressed company could be contained in more than one instrument in the collateral pool. For example, “fallen angels” (like Enron not too long ago) typically cause performance difficulties simultaneously to all CDOs containing this particular risk. Cross-collateralization can only be treated by looking at the union of all collateral pools of all instruments in the multisector pool simultaneously in order to get an idea about the *aggregated* risk of the combined portfolios. Then, based on every aggregated scenario in a Monte Carlo simulation, the cash flows of the different instruments have to be collected and combined in order to investigate the structured cash flows of the multisector CDO.

The credit risk modeling techniques explained in this book can be used for modeling (multisector) CDOs. Of course, a sound factor model, like the one explained in [Section 1.2.3](#), is a necessary prerequisite for modeling CDOs by taking industry and country diversification

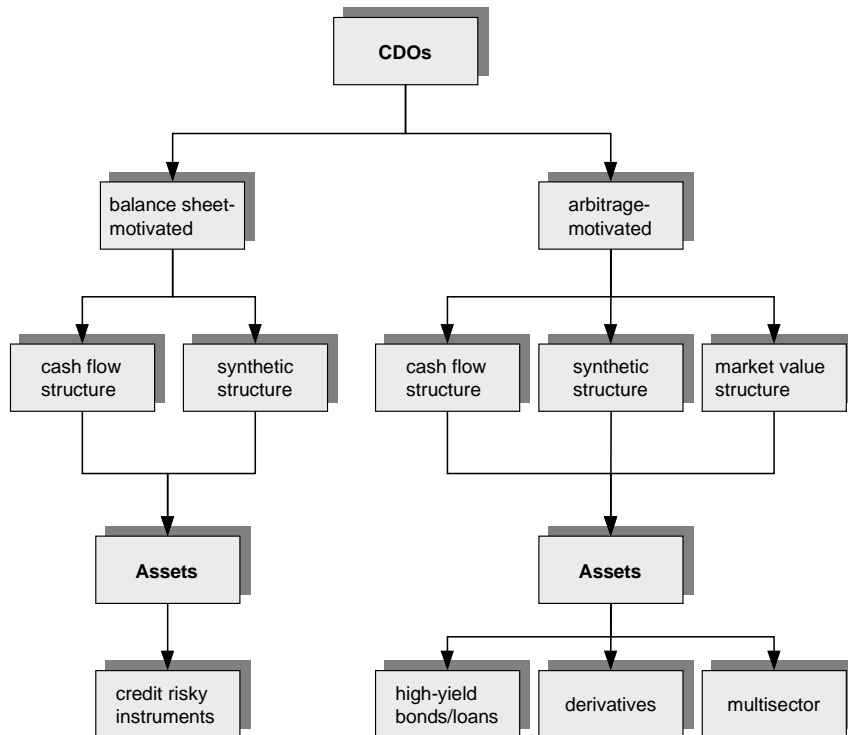


FIGURE 8.1
Classification of CDOs.

effects into account. Moreover, in many cases one additionally has to incorporate an *interest rate term structure model* in order to capture interest rate risk in case of floating rate notes.

In general, *Market value* CDOs are more difficult to treat from a modeling point of view. These structures are more comparable to *hedge funds* than to traditional ABS structures. In a market value CDO, the *portfolio manager* has the right to freely trade the collateral. As a consequence, a market value CDO portfolio today sometimes has very few in common with the portfolio of the structure a few months later. The performance of market value CDOs completely relies on the portfolio manager's expertise to trade the collateral in a way meeting the principal and interest obligations of the structure. Therefore, investors will mainly focus on the manager's *deal track record* and experience when deciding about an investment in the structure. The difficulties on the modeling side arise from the unknown *trading strategy* of the portfolio manager and the need of the manager to react to a *volatile economic environment*. Such subjective aspects are difficult if not impossible to model and will not be treated here.

8.1.1 Typical Cash Flow CDO Structure

In this section, we explain a typical cash flow CDO transaction; see [Figure 8.2](#). For this purpose we focus on some of the main aspects without going too much into details.

- At the beginning there will always be some pool of credit risky assets. Admittedly, it will not always be the case that the pool intended to be securitized was existent at the originating bank's balance sheet for a long time; instead, there are many cases where banks purchased parts of the pool just a few months before launching the transaction. Such purchases are typically done in a so-called *ramp-up* period.
- In a next step, the assets are transferred to an SPV, which is a company set-up especially for the purpose of the transaction. This explains the notion *special purpose vehicle*. An important condition hereby is the *bankruptcy remoteness* of the SPV, essentially meaning that the SPV's own bankruptcy risk is minimized and that the SPV will not default on its obligations because of bankruptcy or insolvency of the originator. This is achieved by a

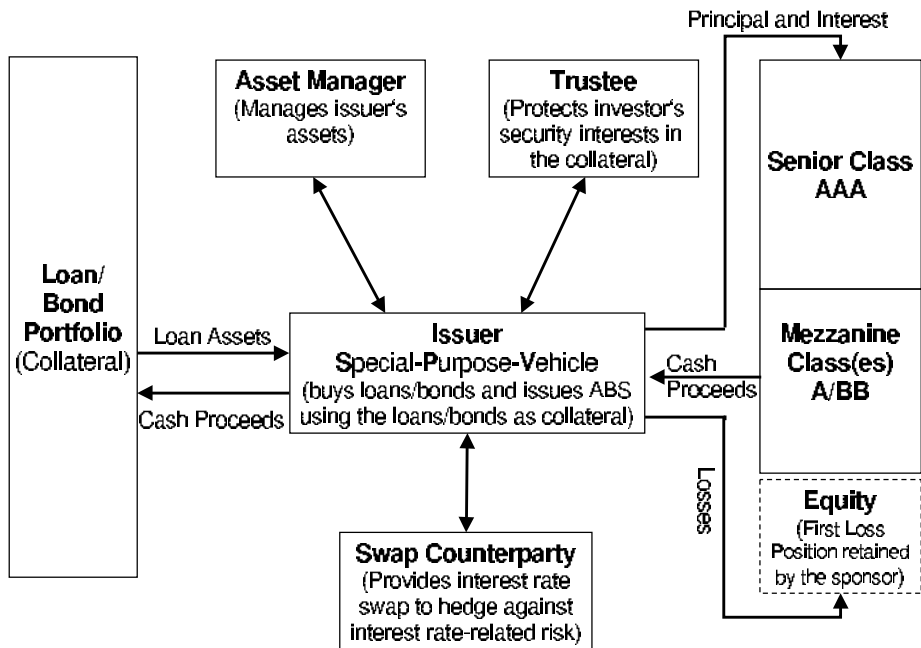


FIGURE 8.2
Example of a cash flow CDO transaction.

strict legal separation between the SPV and the originator, implying a legal and economic independence. Additionally, an SPV's obligations typically involve various structural features supporting the bankruptcy remoteness of the SPV.

In case of cash flow structures, a "true-sale" of the assets from the originator to the SPV completely removes the securitized assets from the originator's balance sheet. However, most often the administration of the asset pool remains the originator's responsibility. The originator receives the principal balance of the pool as cash proceeds, such that from the originator's point of view the funding of the asset pool is completed.

- After the true sale, the assets are property of the SPV. Therefore, the SPV is the owner of all of the cash flows arising from the asset pool. This can be used to establish a funding source for the SPV's purchase of assets from the originator. Note that as a special purpose company, the SPV itself has no money for paying the principal balance of the asset pool to the originating institution. A way out is the issuance of securities or *structured notes* backed¹ by the cash flow of the asset pool. In other words, the SPV now issues notes to investors, such that the total notional of notes reflects the principal balance of the pool. Interest and principal for the notes are paid from interest and principal proceeds from the asset pool. This mechanism changes the meaning of the asset pool towards a *collateral* pool. From the issuance of notes, the SPV receives cash proceeds from the investors, refinancing the original purchase of assets from the originating institution.

Because investor's proceeds (principal and interest) are paid from cash flows generated by the collateral pool, investors are taking the performance risk of the collateral pool. Because different investors have different risk appetite, the notes issued by an SPV are typically *tranche*d into different *risk classes*. The *first loss piece* (FLP), often also called the *equity piece*² is the most *subordinated* tranche, receiving interest and principal payments only if all other notes investors received their contractually promised payments.

¹This perfectly explains the name *asset backed* securities.

²The equity tranche is sometimes kept by the originating institution, therefore constituting equity.

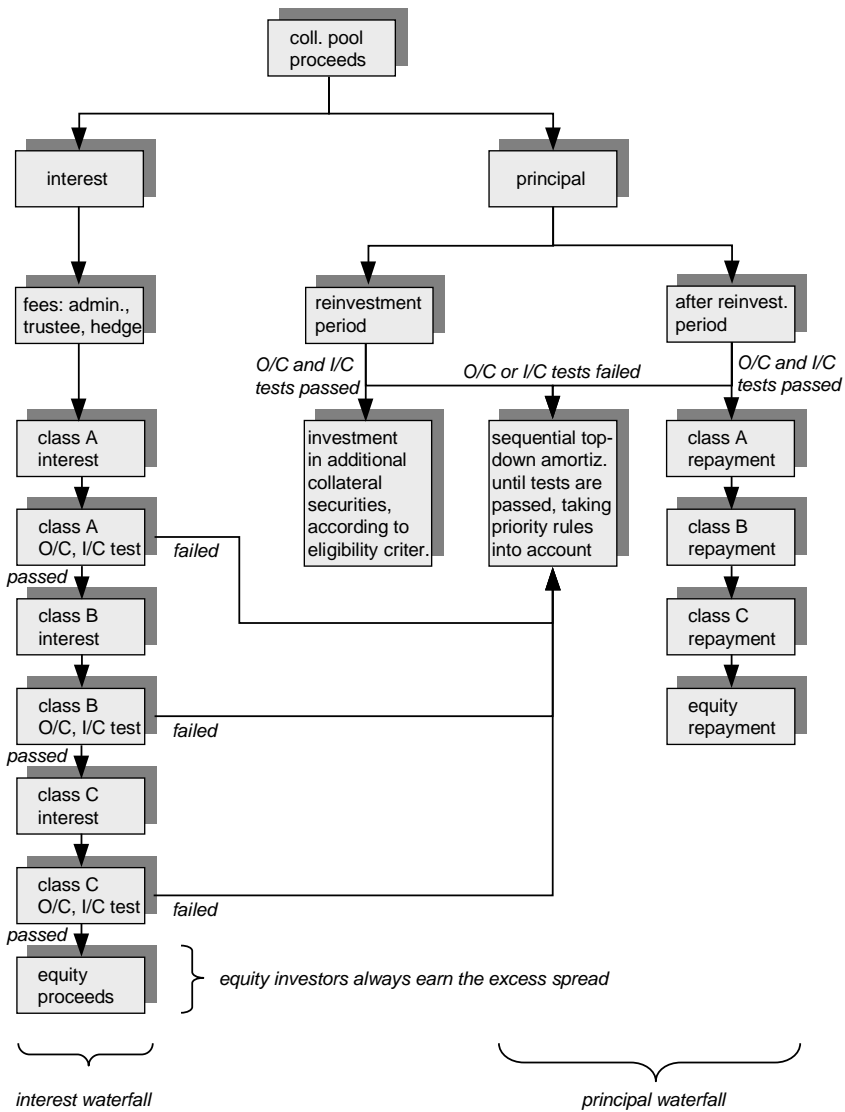


FIGURE 8.3
Example of cash flow waterfalls in a cash flow CDO.

TABLE 8.1: CDO example (illustrative only).

Class	Notional in USD	Tranche Size	Moody's Rating	Spread ov. LIBOR (bps)	minimum O/C ratio	minimum I/C ratio
A	225,000,000	75%	Aa2	100	120%	140%
B	30,000,000	10%	Baa3	250	110%	125%
C	15,000,000	5%	Ba1	550	105%	110%
Equity	30,000,000	10%	Not Rated	-	-	-

The FLP is followed by *junior*, *mezzanine*, and *senior* tranches, receiving interest and principal proceeds in the order of their seniority: Most senior note holders receive their payments first, more junior note investors receive payments only if more prioritized payments are in line with the documentation of the structure. Therefore, the most senior tranche always is the safest investment, carrying the lowest coupon. The more junior a tranche is, the higher the promised coupon, compensating investors for the taken risk.

An exception is the equity tranche, which typically carries no promised coupon. Instead, equity investors receive the *excess spread* of the structure in every payment period, where the excess spread is the left-over cash after paying all fees of the structure and all payments to notes investors senior to the equity piece.

From the discussion above it follows, that *subordination* is kind of a *structural credit enhancement*. For example, in a structure with only one equity, mezzanine, and senior tranche, the senior note holders are protected by a cushion consisting of the equity and mezzanine capital, and the mezzanine tranche is protected by the equity tranche.

Figure 8.3 describes the interest and principal proceeds “waterfalls” in a typical cash flow CDO. The figure also indicates the *deleveraging* mechanism inherent in CDO structures, realized by *overcollateralization* (O/C) and *interest coverage* (I/C) tests, which brings us to our last topic in this section.

But before continuing we should mention that there are additional parties involved in a CDO transaction, including

- rating agencies, which assign ratings to the issued notes,
- a trustee, which takes care that the legal documentation is honored and receives monthly trustee reports regarding the current

performance of the structure,

- some swap counterparties in case interest or currency risk has to be hedged, and
- lawyers, structuring experts, and underwriters at the beginning of the transaction, where the latter mentioned are hired from another investment bank or from inhouse business units.

Now, in order to explain the O/C and I/C mechanisms in a cash flow CDO, let us consider a simple illustrative example. Let us assume we are given a structure like the one outlined in [Table³ 8.1](#). Further we assume that

- the collateral pool contains 100 corporate bonds with an average default probability

$$DP = 3\% ,$$

and a weighted⁴ average coupon (WAC) of

$$WAC = 10.4\% ,$$

reflecting the risk inherent in the collateral securities;

- spreads and default probabilities are annualized values. The following discussion is independent of the maturity of the structure.

Now we are ready for explaining the O/C respectively I/C mechanisms. Basically these *coverage tests* are intended as an early warning (automatically redirecting cash flows) that interest or principal proceeds are running short for covering the notes coupons and/or repayments. In case of a broken coverage test, principal and interest proceeds are used for paying back the outstandings on notes sequentially (senior tranches first, mezzanine and junior tranches later) until all tests are passed again. This *deleveraging* mechanism of the structure reduces the exposure at risk for tranches in order of their seniority. So one can think of coverage tests as some kind of *credit enhancement* for protecting notes investors (according to prioritization rules) from suffering a loss (a missed coupon⁵ or a repayment below par).

³In the table, LIBOR refers to the 3-month *London Interbank Offered Rate*, which is a widely used benchmark or reference rate for short term interest rates.

⁴For reasons of simplicity assuming that the bonds trade at par, the weighting is done w.r.t. the principal values of the bonds.

⁵For mezzanine investors often a *deferred interest* is possible: If the cash flow from the collateral securities is not sufficient for passing the coverage tests, mezzanine investor's

8.1.1.1 Overcollateralization Tests

In these tests, which are done for every single tranche except equity, the principal coverage of collateral securities compared to the required amount for paying back the notional of the considered tranche *and* the tranches senior to the considered tranche is tested. In the structure according to [Table 8.1](#), three O/C tests have to be done:

O/C test for class A notes: Denote the par value of the pool by PV_{Pool} and the par value of class A notes by PV_A , where par values are taken w.r.t. the considered payment period in which the test is done. (Synonymously to “par value” we could also say “outstandings” on notes.) Define

$$(O/C)_A = \frac{PV_{Pool}}{PV_A} .$$

The O/C test for class A notes is passed if

$$(O/C)_A \geq (O/C)_A^{\min} = 120\%$$

reflecting the minimum O/C ratio for class A as given in [Table 8.1](#).

O/C test for class B notes: Define

$$(O/C)_B = \frac{PV_{Pool}}{PV_A + PV_B} .$$

The O/C test for class B is passed if

$$(O/C)_B \geq (O/C)_B^{\min} = 110\% .$$

Note that the O/C test for class B takes into account that class A notes have to be paid back before.

O/C test for class C notes: Set

$$(O/C)_C = \frac{PV_{Pool}}{PV_A + PV_B + PV_C} .$$

coupon payments are deferred to a later payment period, where all tests are in line again. Deferred interest is paid on an accrued basis.

TABLE 8.2: O/C ratios example (illustrative only).

Portfolio Par	O/C_A	O/C_B	O/C_C
300,000,000	133%	118%	111%
275,000,000	122%	108%	102%

The O/C test for class C investors is passed if

$$(O/C)_C \geq (O/C)_C^{\min} = 105\% .$$

Note that the O/C test for class C takes into account that the outstandings of classes A and B have to be paid back before class C investors get their invested money back; see also [Figure 8.3](#).

To give an example, assume we are in a payment period where the pool volume due to losses of 25,000,000 USD melted down to 275,000,000 USD from the previous to the current period. [Table 8.2](#) shows the O/C ratios for the previous and the current period. One can see that the coverage is still sufficient for class A to pass the test, but insufficient for classes B and C. Their O/C tests are broken. This will cause a deleveraging of the CDO until all tests are in line again.

8.1.1.2 Interest Coverage Tests

An I/C test for a tranche basically measures if the interest proceeds from the collateral pool are sufficient for paying the fees and coupons of the structure. In our particular example there are three tests:

I/C test for class A notes: For the considered payment period, denote the par value of the pool by PV_{Pool} , the par value of class A by PV_A , the amount of required annual fees by FEES, the weighted average coupon of the pool by WAC, and the coupon⁶ on class A notes by C_A . Define

$$(I/C)_A = \frac{PV_{Pool} \times WAC \times 0.5 - FEES \times 0.5}{PV_A \times C_A \times 0.5} .$$

⁶In our example we are dealing with floating-rate notes. Here, the coupon on notes is always defined as LIBOR+Spread.

Here, the factor 0.5 reflects that interest is calculated w.r.t. a semiannual horizon, covering two (quarterly) payment periods. Of course, the concrete calculation method for I/C and O/C ratios always has to be looked-up in the documentation of the structure. The I/C test for class A notes is passed if

$$(I/C)_A \geq (I/C)_A^{\min} = 140\%$$

reflecting the minimum required I/C ratio for class A notes according to [Table 8.1](#).

I/C test for class B notes: Define

$$(I/C)_B = \frac{PV_{Pool} \times WAC \times 0.5 - FEES \times 0.5}{(PV_A \times C_A + PV_B \times C_B) \times 0.5} .$$

The I/C test for class B is passed if

$$(I/C)_B \geq (I/C)_B^{\min} = 125\% .$$

Analogous to the O/C tests, the calculation reflects that class A notes have priority before class B notes regarding coupon payments.

I/C test for class C notes: Setting

$$(I/C)_C = \frac{PV_{Pool} \times WAC \times 0.5 - FEES \times 0.5}{(PV_A \times C_A + PV_B \times C_B + PV_C \times C_C) \times 0.5} ,$$

class C interest coverage requires classes A and B to be covered, before C-notes investors receive coupon payments. The test is passed if

$$(I/C)_C \geq (I/C)_C^{\min} = 110\% .$$

The interest waterfall as illustrated in [Figure 8.3](#) is clearly reflected by these calculations.

[Table 8.3](#) gives an example for the value of the three I/C ratios right at the beginning of the transaction. For calculating the I/C ratios we assumed the current 3-month LIBOR to be equal to 4%.

TABLE 8.3: I/C ratios example (illustrative only)

Portfolio Par	I/C_A	I/C_B	I/C_C
300,000,000	273%	233%	210%

Total Fees	450,000	USD p.a.
-------------------	---------	-----------------

8.1.1.3 Other Tests

Typically, there are some more tests whose outcome has to be reported to the trustee and to the investors. The collection of tests and criteria varies from deal to deal, and not all tests included in the monthly reports automatically have immediate consequences on the cash flow side of the structure. Some tests one frequently finds in deal documentations are

- an *average rating floor* test, reporting whether the weighted average rating of the collateral pool is above a critical threshold; a typical threshold for cash flow CDOs is Moody's B-rating;
- *industry* and *diversity score*⁷ tests, alarming investors in case the industry diversification of the collateral pool decreased more than expected; a common range for the highest admissible industry concentration is 8-12%;
- an *obligor concentration* test, measuring the highest exposure concentration in the collateral pool, often restricted to concentrations below 3%,

and possibly some more tests helping the investors to identify and quantify potential sources of risk and financial distress of the structure.

This concludes our description of cash flow CDOs for now.

8.1.2 Typical Synthetic CLO Structure

In contrast to cash flow CDOs, synthetic CDOs do not rely on the cash flows of the collateral pool. Instead, credit derivatives, e.g., credit-linked notes, are used to link the performance of securities issued by

⁷Diversity scores are a measure of the industry diversification of a portfolio; diversity scores are due to Moody's and will be explained in [Section 8.4](#).

an SPV to the performance, e.g., the losses, of some *reference pool*. In other words, synthetic CDOs do not include a true sale, as we just discussed it in case of cash flow deals, but rather leave the reference assets on the originator's balance sheet. [Figure 8.4](#) shows a typical synthetic CLO as we find it in the market. It could work as follows:

- The originator buys protection for *super senior* and *junior*⁸ pieces of the reference portfolio by entering into two credit default swaps with some swap counterparties (protection sellers, see [Chapter 7](#)). The volume referring to the two swaps is called the *unfunded part*, because there is no “sale” requiring certain sources of funding.
- An SPV, which has to be bankruptcy remote for regulatory reasons, enters into a swap with the originator for the volume of the reference portfolio which is not covered by the senior and junior swaps.
- In order to guarantee the contingent payments to the originator in case of credit events in the reference pool, the SPV has to invest some money in collateral securities. Then, in case of a credit event, the contingent payments from the SPV (protection seller) to the originator (protection buyer) can be funded by selling collateral securities in an amount matching the realized losses in the reference portfolio.
- For purchasing collateral securities, the SPV needs some source of funding. In the same way as we already saw it for cash flow deals, the SPV issues credit-linked notes in the capital market, linked to the performance of the reference pool. The outstandings of the issued notes match the volume of the reference pool reduced by the size of the junior and senior swap tranches. The SPV invests the cash proceeds from issuing the notes in low-risk (AAA-quality) or riskless (treasury) notes.
- The spreads on the notes the SPV has to pay to notes investors match the premium the SPV receives from the originator who bought protection from the SPV for the funded part of the reference portfolio.

⁸Selling the FLP of a synthetic transaction to an investor often involves a so-called *interest sub-participation*, essentially meaning that part of the reference pool's interest will be made available to the FLP-investor in case of losses.

- If a credit event in the reference pool occurs, it depends on the already cumulated losses who has to pay for it. Losses below the upper limit of the junior tranche are carried by the junior swap counterparty. Losses exceeding the junior piece but below the super senior tranche are (additionally to the contingent payment made by the junior swap protection seller) carried by the SPV from the originator's point of view and carried by the investor's from the SPV's point of view. Indeed, because collateral securities will be sold for funding the contingent payment the SPV has to make to the originator, investors will not get the complete face value of their invested money back at the final maturity of the structure. The more junior the notes, the more likely it is that investors will not be fully repaid. Super senior losses, which refer to loss events far out in the tail of the reference portfolio's loss distribution, are taken by the super senior swap counterparty. If super senior swap counterparties have to pay for losses, all subordinated investors already had to make their contingent payments on the swap agreements.

Please note that in the market one finds all kinds of variations of the illustrative synthetic CLO we just described. For example, instead of credit default swaps some form of *financial guarantee* could be used. In some cases there will be a *non-cash settlement* in that the protection buyer sells the defaulted loan to the protection seller for par right after a contractually specified credit event occurred.

Additionally, most synthetic structures involve *triggers* based on, e.g., rating distributions, diversity scores, collateral values, losses, defaults, etc. For example, a loss trigger could be defined by saying that in case losses exceed a critical threshold ("trigger event"), some structural features of the transaction change in a contractually pre-specified manner. In this way, triggers are structural elements providing protection to note holders, comparable to the coverage tests discussed above.

In the last years, many new innovative structures offered interesting investment opportunities. Due to inefficient markets and regulatory arbitrage (see our discussion in [Section 8.2](#)), this trend can be expected to continue.

Because in our example there is a funded *and* an unfunded part of the transaction, such a synthetic CLO is called *partially funded* accordingly. Again we should remark that all variations are possible and existent in the market: Fully funded, partially funded, and totally unfunded.

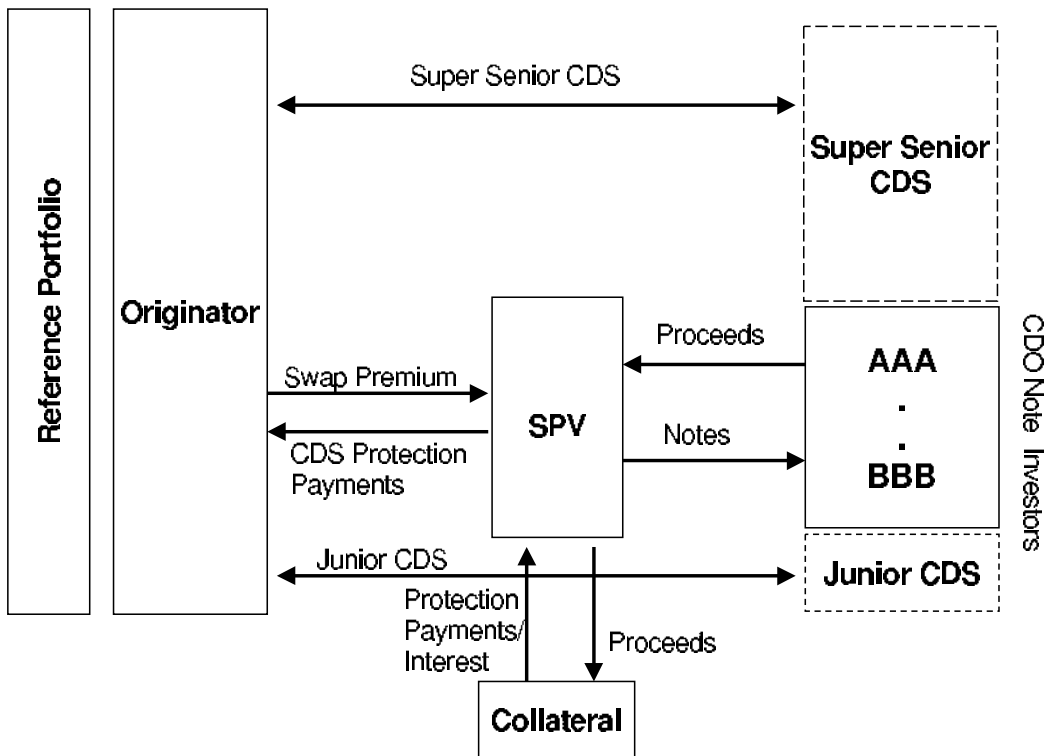


FIGURE 8.4
Example of a synthetic CDO transaction.

Agreements regarding the definition of credit events and the settlement after the occurrence of credit events can be made based on ISDA master agreements, see also [Chapter 7](#).

The *tranching* of the reference portfolio into junior, funded, and super senior parts follows analogous rules as we just saw it in case of cash flow CDOs. The more junior a note is, the higher the premium paid for buying protection for the considered tranche. The more senior a tranche is, the safer investors can invest, but the lower the premium they earn on the investment. Note that this is in line with the risk-adjusted pricing of swap contracts, see [Chapter 7](#).

8.2 Different Roles of Banks in the CDO Market

There are in general many roles of banks in the ABS market. In many cases, a bank will play the role of the *originator* or the role of the *investor*. But there are certainly other roles, which will not be discussed in this book. For example, banks also *provide liquidity*, *guarantee* for promised cash flows, offer different types of *credit enhancement*, and sell their services for *structuring* or *underwriting* ABS transactions. Of course, different roles require different models, so that in general one can say that parallel to the ABS market a whole range of models is needed to measure the different risks the bank is exposed to when participating in the ABS market. In the following section we discuss origination, and in a short subsequent section we make some remarks on ABS investments.

8.2.1 The Originator's Point of View

This section discusses *securitization*. The original meaning of the word “securitization” is funding by means of issuing (structured) securities. Today, banks mainly do securitizations for several reasons, including

- *transferring risk*,
- *arbitrage spread* opportunities.
- *funding* at better (cheaper) conditions, and

- exploitation of *regulatory* or *tax arbitrage*,

From a portfolio modeling point of view, there is a fundamental difference between the first three and the last securitization benefits: Risk transfer, arbitrage spread opportunities, and (to some extent) better funding are *correlation-driven* effects, whereas regulatory capital relief and tax arbitrage are *correlation-free* effects.

8.2.1.1 Regulatory Arbitrage and Capital Relief

The keyword *regulatory arbitrage* refers to opportunities in the markets due to inappropriate regulation by the regulatory authorities. For example, as indicated in [Section 1.3](#), regulatory capital for bank loans is still not risk-adjusted. This results in “pricing distortions”, due to the fact that the capital costs of a loan are independent of the credit quality of the borrower.

By means of a simple illustrative example we now want to show how regulatory arbitrage is exploited in securitization transactions. Because regulatory capital calculations are independent of the level of correlation in a portfolio, the following calculations focus on the *securitized portfolio only*, ignoring the (“context”) *source portfolio*, of which the securitized portfolio is a subportfolio.

Let us assume we are considering a subportfolio of a bank’s credit portfolio. The *risk-weighted assets* (RWA), see [Section 1.3](#), of the pool are assumed to be equal the total volume V of the pool. This essentially means that the collateral of the loans in the portfolio is not eligible for reduced *risk weights*. By the *8%-rule* mentioned in [Section 1.3](#), the *regulatory capital* (RC_{Pool}) of the portfolio is given by

$$RC_{Pool} = RWA \times 8\% = V \times 8\% .$$

Assume further that the portfolio’s expected loss (EL) is 50 bps and that the portfolio’s weighted average net⁹ margin (NM) equals 130 bps. This results in a *return on equity* (RoE_{Pool}) of the portfolio of

$$RoE_{Pool} = \frac{NM - EL}{RC_{Pool}} = \frac{130 - 50}{800} = 10\% .$$

The RoE measures the return, net of expected losses, in units of required equity. Note that although we are not proposing to steer a bank

⁹Net of liquidity, funding, and administration costs.

w.r.t. RoE benchmarks, for measuring just the regulatory arbitrage gain of a transaction the RoE gives an acceptable indication.

To continue the story, let us now assume that the bank decides to securitize the subportfolio by means of a synthetic CLO as illustrated in Figure 8.4, but without a junior swap. Let us say the costs of the structure sum up to 30 bps, including spreads on notes, swap fees, administrative costs, and upfront fees like rating agency payments, structuring and underwriting costs. We assume (illustrative only!) that

- the super senior swap refers to the upper 85% of the volume of the reference portfolio; swap counterparty is an OECD bank;
- the remaining 15% of the volume V are funded by issuing credit-linked notes,
- a first loss piece (FLP) of 1.5% is kept by the originating bank, whereas all tranches senior to the FLP are placed in the market.

The regulatory capital after Securitization is then given by

$$RC_{Sec} = V \times 1.5\% \times 100\% + V \times 85\% \times 20\% \times 8\% = V \times 2.86\% ,$$

because for capital covered by a credit default swap agreement with an OECD bank, regulatory authorities typically allow for a risk weight of 20%, and the funded part of the transaction is eligible for a zero risk weight, if the collateral securities purchased from the proceedings of issuing notes are attributed as “risk-free” according to national regulation standards. Equity pieces typically require a 100%-capital cushion from the regulator’s point of view. Ignoring for a moment the fact that the EL after securitization should be lower than before, because the bank now only is exposed to losses up to 1.5% of the portfolio’s total volume, the RoE after securitization can be calculated as

$$\begin{aligned} RoE_{Sec} &= \frac{NM - EL - COSTS}{1.5\% \times 100\% + 85\% \times 20\% \times 8\%} = & (8. 1) \\ &= \frac{130 - 50 - 30}{286} = 17.48\% . \end{aligned}$$

So the RoE has improved by more than 70%. A second advantage from the securitization is the capital relief. Before the securitization, the portfolio consumed $V \times 8\%$ of the bank’s equity. After the securitization, the portfolio consumes only $V \times 2.86\%$ of the bank’s equity,

although the loans are still on the bank's balance sheet (recall the discussion in the section on synthetic CLOs). In other words, after securitization the bank has the opportunity to use the relieved equity (in total: $V \times [8\% - 2.86\%]$) for entering into new business.

8.2.1.2 Economic Risk Transfer

In the discussion on regulatory capital above we already indicated that a securitization also should lower the EL of the portfolio. To understand this point, let us briefly recall the basic purpose of the EL as we introduced it at the very beginning of this book. The EL can be compared to an *insurance* premium, collected as a capital cushion against expected losses. Now, here is the point: After securitizing the portfolio, there is no longer the need to have an insurance against the full loss potential of the portfolio. Instead, in our example, the bank is only exposed to the first 1.5% of portfolio losses. All losses exceeding $V \times 1.5\%$ are taken by the notes investors and the super senior swap counterparty.

Moreover, the same argument conceptually also holds for the economic capital (EC; see [Section 1.2.1](#)) of the securitized portfolio. But because the EC involves correlations (and therefore incorporates *diversification* effects), a securitization not only impacts the securitized pool, but also influences the EC of the *source portfolio*, from which the securitized pool has been separated.

In more mathematic terms, we have the following situation:

Denote by $I = \{1, \dots, m\}$ an index set referring to the loans in the source portfolio, and let us assume that a subportfolio indexed by $S = \{i_1, \dots, i_q\} \subset I$ has been selected for securitization by means of a CLO. The bank now wants to quantify the securitization impact on the source portfolio. For this purpose, the portfolio's EL and EC have to be recalculated after the "portfolio shrinking" $I \rightarrow I \setminus S$.

Now, based on Monte Carlo simulation techniques, the securitization impact is not difficult to calculate. Let us assume the bank would manage to sell all tranches of the CLO except the equity piece, which is then hold by the bank. For reasons of simplicity we consider the one-year period from the launch of the deal until one year later. The size of the equity piece, FLP, is a random variable due to the uncertainty regarding the performance of the collateral securities. Denoting the loss statistics of the whole portfolio I by (L_1, \dots, L_m) , see [Chapter 2](#),

the gross portfolio loss *before* the securitization transaction equals

$$L = \sum_{i=1}^m L_i ,$$

hereby assuming an LGD of 100% and exposures equal to 1 for reasons of simplicity. The portfolio's gross loss *after* securitization obviously is given by

$$L_{Sec} = \sum_{i \in I \setminus S} L_i + \min \left(\sum_{k=1}^q L_{i_k}, \text{FLP} \right) , \quad (8. 2)$$

because the securitized portfolio S is protected against losses exceeding FLP. But the variables L and L_{Sec} can be easily simulated by use of the Monte Carlo engine of the bank. After simulation, we have a loss distribution of the portfolio *before* the transaction and a loss distribution of the portfolio *after* the securitization. The *expected loss gain* respectively *economic capital gain* of the transaction is given by

$$\Delta \text{EL} = \mathbb{E}[L] - \mathbb{E}[L_{Sec}]$$

respectively

$$\begin{aligned} \Delta \text{EC}_\alpha &= \text{EC}_\alpha(L) - \text{EC}_\alpha(L_{Sec}) = \\ &= (q_\alpha(L) - \mathbb{E}[L]) - (q_\alpha(L_{Sec}) - \mathbb{E}[L_{Sec}]) = \Delta q_\alpha - \Delta \text{EL} , \end{aligned}$$

where q_α denotes the α -quantile of the respective loss distribution (before and after securitization, respectively), and

$$\Delta q_\alpha = q_\alpha(L) - q_\alpha(L_{Sec}) .$$

These calculations are sufficient for capturing the securitization impact on the source portfolio.

“Risk transfer” refers to the possibility to reduce the required capital cushion against losses of a portfolio by means of a securitization. “Economic” risk transfer happens, if the risk transfer can be measured in terms of the EL and EC, such that ΔEL and Δq_α are positive.

Now assume that the securitized pool S belongs to some business unit of the bank with its own profit center. Then, the securitization impact additionally has to be measured from that profit center's point of view, so we additionally need to quantify EL and EC benefits for the

securitized pool only. Keeping the notation from above, the pool loss before and after securitization is

$$L_{Pool} = \sum_{k=1}^q L_{i_k} \quad \text{and} \quad L_{Pool}^{Sec} = \min \left(\sum_{k=1}^q L_{i_k}, \text{FLP} \right).$$

The pool's EL benefit of the securitization therefore is

$$\Delta EL_{Pool} = \mathbb{E}[L_{Pool}] - \mathbb{E}[L_{Pool}^{Sec}].$$

Obviously, ΔEL_{Pool} is positive if and only if there is at least one cumulative loss path¹⁰ for which the cap at FLP turns out to be effective.

Regarding EC, we now have to consider the *contributory economic capital* (CEC) of the securitized pool w.r.t. the source portfolio; see also [Section 5.2](#). The gain in CEC of the securitization is given by

$$\Delta CEC_{\alpha} = CEC_{Pool} - CEC_{Pool}^{Sec},$$

where CEC_{Pool} respectively CEC_{Pool}^{Sec} denotes the CEC of the securitized pool before respectively after securitization.

We conclude our discussion by briefly mentioning a common performance measure capturing the effects of economic risk transfer, namely *risk-adjusted return on capital* (RAROC). There are various definitions of RAROC measures in the literature, but here we use it just for illustrative purposes and therefore keep things as simple¹¹ as possible. More or less, RAROC always is defined as the risk-adjusted return of an instrument or portfolio divided by the corresponding EC. To illustrate the effect of securitization to RAROC benchmarks, let us assume that the CEC of the pool before securitization was $CEC_{Pool} = 5\%$. Let us further assume that after securitization the CEC of the pool melted down to $CEC_{Pool}^{Sec} = 150$ bps. The EL of the pool is assumed to be reduced from $\mathbb{E}[L_{Pool}] = 50$ bps to $\mathbb{E}[L_{Pool}^{Sec}] = 40$ bps, due to the securitization. This yields

$$\text{RAROC} = \frac{\text{NM} - \mathbb{E}[L_{Pool}]}{\text{CEC}_{Pool}} = \frac{130 - 50}{500} = 16\%$$

¹⁰Thinking in terms of a Monte Carlo simulation.

¹¹For example, we do not take, as often done, the *capital benefit* arising from riskfree interest earned on the EC into account.

for the portfolio before the securitization transaction, and

$$\text{RAROC}_{Sec} = \frac{\text{NM} - \mathbb{E}[L_{Pool}^{Sec}] - \text{COST}}{\text{CEC}_{Pool}^{Sec}} = \frac{130 - 40 - 30}{150} = 40\%$$

after securitizing the portfolio. So the securitization improves the RAROC of the portfolio by a factor of 2.5, just due to the protection limit of 1.5%.

Note that the discussion above was based on a one-period view, e.g., based on an average lifetime consideration. For measuring economic risk transfer and securitization effects on RAROC more accurately, much more work and modeling efforts are required, very often accompanied by strong assumptions, e.g., regarding the evolution of the reference pool.

8.2.1.3 Funding at Better Conditions

Funding is an important issue for banks. Because every loan needs to be backed by regulatory capital, the capital costs associated with a loan to a customer can be too high for making the lending business profitable. But if loans are pooled into portfolios for securitization, funding a loan can get significantly cheaper. The reasons why a securitization makes funding cheaper, are basically given in the two sections above: Because regulatory capital is relieved, equity costs of the securitized portfolio are much lower than they used to be. Moreover, if an economic risk transfer is achieved, EL costs and EC costs will be reduced to an extent reflecting the amount of risk transferred to the capital market. Both effects, and additional tax and other benefits can help a bank to refinance a loan portfolio at much better conditions than it was the case before the securitization.

8.2.1.4 Arbitrage Spread Opportunities

Arbitrage spread opportunities are created in the following way. The assets in the collateral pool (in our example of a cash flow CDO we are talking about high-yield bonds) are priced on a *single* asset basis, such that every bond coupon in the portfolio reflects the risk of the bond. Of course, in general the coupon of a bond *and* its price provide the “full” information about the risks inherent in the bond. However, assuming a bond trades at par allows for taking the coupon of the bond as a proxy for its riskiness. So the WAC of the collateral pool

really is a weighted sum of *single* asset risks, ignoring the potential for diversification effects typically inherent in a portfolio.

In contrast, on the CDO side, it is the *portfolio* risk which endangers the performance of the structure. Recalling our discussion on cash flow CDOs, we see that the tranching of notes really is a *tranching of the loss distribution* of the collateral pool, taking all possible diversification effects into account. But diversification decreases the risk of a portfolio, so that the price of the portfolio risk must be lower than the price obtained by just summing up exposure-weighted single risks. This is reflected by the spreads on notes as given in [Table 8.1](#): The spreads paid to notes investors are much lower than the spreads earned on the bonds in the collateral pool. Due to the risk tranching of notes, the spreads on senior notes is even lower, due to the credit enhancement by subordination provided from notes with lower seniority.

It is exactly the mismatch between the single asset based WAC of the portfolio and the much lower weighted average coupon on the notes of the CDO, which creates an arbitrage spread. This mismatch is in one part due to diversification effects, and in another part based on structural elements like subordination or other credit enhancement mechanisms. Calling special attention to the diversification point, one can say that *CDOs are “correlation products”*.

An example regarding arbitrage spread is given in the next section in the context of CDO investments. Conceptually, any originator of an arbitrage cash flow CDO keeping the CDO's first loss piece automatically takes on the role of the *equity investor*, earning the *excess spread* of the structure in its own pockets. Therefore, we can postpone the arbitrage spread example to the next section.

8.2.2 The Investor's Point of View

Very often banks are on the investment side of a CDO. In many cases, ABS bonds offer interesting and attractive investment opportunities, but require (due to their complexity) careful analytic valuation methods for calculating the risks and benefits coming with an ABS investment into the bank's portfolio. This will be made explicit by means of the following example.

Recall the sample cash flow CDO from [Table 8.1](#). In this example we assumed $WAC = 10.4\%$ and $DP = 3\%$. Assuming an LGD of

80% on the collateral securities, we obtain the portfolio's expected loss, $EL = 3\% \times 80\% = 2.4\%$.

Considering the CDO from an *expected return* point of view, what would an equity investor expect to earn on an investment in the equity tranche? A typical “back-of-the-envelope” calculation reads as follows:

From [Table 8.1](#) we obtain the weighted average coupon WAC_{Notes} of the structure as

$$WAC_{Notes} = 75\% \times 5\% + 10\% \times 6.5\% + 5\% \times 9.5\% = 4.875\% ,$$

again assuming the average 3-month LIBOR to be equal to 4%. Because cash flow CDOs completely rely on the cash flows from the collateral pool, the 10.4% of the pool's par value are the complete income of the structure. From this income, all expenses of the structure have to be paid. Paying¹² coupons to notes investors yields a *gross arbitrage spread* (gross excess spread) of

$$[\text{Pool Income}] - [\text{Notes Spreads}] = 10.4\% - 4.875\% = 5.525\% .$$

The *expected net excess spread* is then defined as

$$\begin{aligned} & [\text{Gross Arbitrage Spread}] - EL - \text{COSTS} = \\ & = 5.525\% - 2.4\% - \frac{450,000}{300,000,000} = 2.975\% . \end{aligned}$$

The *equity return* is then given by

$$[\text{Exp. Net Excess Spread}] \times \frac{\text{Pool Volume}}{\text{Equity Volume}} = 29.75\% .$$

So the “back-of-the-envelope” calculation promises a very attractive equity return of almost 30%.

Now let us look at this seemingly attractive investment from a portfolio modeling point of view. For this purpose we calculated the *equity return distribution* of the CDO by means of a *correlated default times* approach as outlined later on in this chapter; see also [Chapter 7](#). From a Monte Carlo simulation we obtained¹³ [Figure 8.5](#). Hereby we essentially followed the CDO modeling scheme as illustrated in [Figure 8.3](#), adapted to a default times approach according to [Figure 8.7](#).

¹²Referring to an average scenario.

¹³Under certain assumptions regarding the maturity of the bonds and the structure.




FIGURE 8.5
Equity return distribution of a CDO.

Looking at the equity return distribution in [Figure 8.5](#), it turns out that, in contrast to the above shown “back-of-the-envelope” calculation, the Monte Carlo simulation yields an *average equity return* of only 15.92%. Additionally, the *volatility of equity returns* turns out to be 9.05%, so by just one standard deviation move, the equity return can vary between 6.87% and 24.98%. This reflects the fact that equity investments are rather volatile and therefore very risky. Moreover, due to *tail events* of the collateral pool’s loss distribution, it can happen that the *downside risks* of equity investments dominate the *upside chances*.

We continue our example by looking at the return distribution for class-A notes investors. [Table 8.4](#) shows that in 94.17% of the cases the promised coupon of 5% has been paid to A-investors. However, in 5.83% of the cases, either not a full coupon payment or not a full repayment resulted in a loss. Here, loss means that at least one contractually promised dollar has not been paid. So the 5.83% are indeed the *default probability* of the senior tranche of the CDO. For a Aa2-rating, this is a very high chance for default. Additionally, the simulation yields an expected loss of the Aa2-tranche of 50bps, which again is very high compared to Aa2-rated bonds. Defining the loss given default of the

TABLE 8.4: Return statistics for class-A notes investors

Return Statistics for Class-A Notes	
Return Range	Relative Frequency
Return=5%	94.17%
4%≤Return<5%	2.78%
3%≤Return<4%	1.19%
2%≤Return<3%	0.59%
1%≤Return<2%	0.61%
0%≤Return<1%	0.39%
Return<0%	0.27%

TABLE 8.5: Weighted average life of tranches

Tranche	Weighted Average Life
A	5.03
B	9.84
C	10.00
Equity	10.00

tranche by

$$\text{LGD}(T_{Aa2}) = \frac{\text{EL}(T_{Aa2})}{\text{DP}(T_{Aa2})} = \frac{50}{583} = 8.6\% ,$$

shows that on the other side the LGD of the tranche is very low. This is also due to the large volume (thickness) of the tranche. In [Section 8.4](#) we will discuss rating agency models, and it will turn out that agency ratings of senior tranches typically underestimate the tranche’s “true” risk. This is due to the fact that rating agency models often neglect the *fat tail* of credit portfolio loss distributions. In our example we can clearly see that the Aa2-rating does not really reflect the “true” risk of the Aa2-tranche.

[Table 8.5](#) shows the weighted average life (WAL) of the four tranches. For the simulation, we assumed that the CDO matures in 10 years. The WAL for class-A notes is quite low, in part due to the amortization structure of the collateral pool, but to some extent also due to broken

coverage tests leading to a *deleveraging* of the outstandings of the notes. Because of the waterfall structure illustrated in Figure 8.3, the most senior class has to be repaid before lower classes receive repayments. This yields the low WAL for class A.

We conclude this section by a brief summary. In the discussion above, our calculations showed that it is very dangerous to rely on “average value” considerations like our “back-of-the-envelope” calculation. Only a full Monte Carlo simulation, based on portfolio models as introduced in this book, will unveil the downside risks and upside chances of an investment in a CDO.

8.3 CDOs from the Modeling Point of View

In this section, a general framework for CDO modeling is presented. Not all structures require all elements mentioned in the sequel. In some cases, shortcuts, approximations, or working assumptions (e.g. a fixed¹⁴, possibly stress-tested, LIBOR) can be used for evaluating a CDO quicker than by means of implementing a simulation model where all random elements are also drawn at random, hereby increasing the complexity of the model.

In our presentation, we will keep a somewhat abstract level, because going into modeling details or presenting a fully worked-out case study is beyond the introductory scope of this chapter. However, we want to encourage readers¹⁵ involved in ABS transactions to start modeling their deals by means of a *full Monte Carlo simulation* instead of just following the common practice to evaluate deals by stress tests and the assumption of fixed loss rates. The example in the previous section demonstrates how dangerous such a “shortcut model” can be.

The evaluation of CDO transactions involves three major steps:

¹⁴For example, if in the documentation of a structure one finds that fluctuations of LIBOR are limited by a predefined cap and floor, then one can think of stress testing the impact of LIBOR variations by just looking at the two extreme scenarios.

¹⁵As far as we know, most major banks use, additionally to the “classic” approaches and rating agency models, CDO models based on Monte Carlo simulation comparable to the approach we are going to describe.

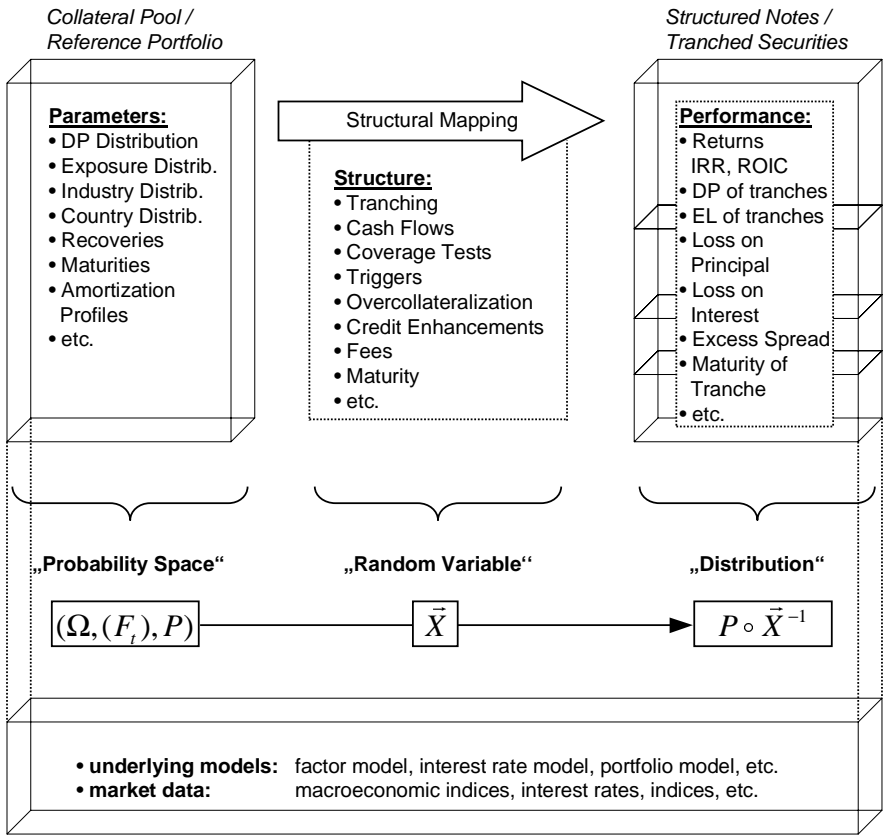


FIGURE 8.6
CDO modeling scheme.

1. Step: *Constructing a model for the underlying portfolio*

Underlying the structure is always an asset pool, for example a reference portfolio or a collateral pool. The structural elements of the considered deal are always linked to the performance of the underlying asset pool, so it is natural to start with a portfolio model similar to those presented in [Chapters 1-4](#). Additionally, such a model should include

- multi-year horizons due to maturities longer than one year,
- a sound factor model for measuring industry and country diversification in an appropriate manner, and
- a model for short term interest rates for capturing the interest rate risk of floating rate securities and notes.

This first step is the only part involving probability theory. The second and third step are much more elementary.

2. Step: *Modeling the cash flows of the structure*

Based on Step 1, the cash flows of the structure *conditioned* on the simulated scenario from the portfolio model representing the performance of the collateral securities should be modeled by taking all cash flow elements of the structure, including

- subordination structure,
- fees and hedge premiums,
- principal and interest waterfalls,
- coverage tests (O/C and I/C),
- credit enhancements (e.g. overcollateralization),
- triggers (e.g. early amortization, call options), etc.,

into account. From a programming point of view, Step 2 consists of implementing an algorithm for “distributing money” (e.g., in a cash flow CDO the cash income from the collateral securities) into “accounts” (some specified variables reflecting, e.g., principal and interest accounts) defined by the contract or documentation of the deal. Such an algorithm should exactly reflect the cash flow mechanisms specified in the documentation, because leaving out just a single element can already significantly distort the simulation results towards wrong impressions regarding the performance

of the structure. In addition to a cash flow model, a discounting method (e.g., a risk-neutral valuation model in case that the risks, e.g., the default probabilities, of the collateral securities are determined according to a risk-neutral approach) should be in place in order to calculate present values of future cash flows

3. Step: *Interpreting the outcome of the simulation engine*

After the simulation, the outcome has to be evaluated and interpreted. Because the performance of the structure is subject to *random fluctuations* based on the randomness of the behaviour of the collateral securities, the basic outcome of the simulation will always consist of *distributions* (e.g., return distributions, loss distributions, etc.); see Figure 8.5 and the discussion there.

Figure 8.6 illustrates and summarizes the three steps by means of a modeling scheme.

In [37], Finger compares four different approaches to CDO modeling, namely a discrete multi-step extension of the CreditMetrics™ portfolio model, a diffusion-based extension of CreditMetrics™, a copula function approach for correlated default times, and a stochastic default intensity approach. The first two mentioned approaches are both multi-step models, which will be briefly discussed in the next section. The basic methodology underlying the third and fourth approach will be outlined in two subsequent sections.

8.3.1 Multi-Step Models

Multi-step models are natural extensions of single-period portfolio models, like the models we discussed in previous chapters. Essentially, a multi-step model can be thought of as many “intertemporally connected” single-period models successively simulated. Considering the three major valuation steps discussed in the previous section, one could describe the three steps in a multi-step model context as follows:

Step 1 defines a *filtered probability space* $(\Omega, (\mathcal{F}_t), \mathbb{P})$, where:

- Ω consists of the whole universe of possible scenarios regarding the collateral pool and the interest rate model. More precisely, every scenario $\omega \in \Omega$ is a *vector* whose components are defined by the possible outcomes of the portfolio model, including a default/migration indicator realization for every collateral security, a realization of LIBOR, etc.

- $(\mathcal{F}_t)_{t=1,\dots,T}$ is a filtration of σ -algebras containing the measurable events up to the payment period t . Any σ -algebra \mathcal{F}_t can be interpreted as the collection of events reflecting informations known up to payment period t . For example, \mathcal{F}_t contains the event that up to time t the portfolio loss already crossed a certain limit, etc. Here, T represents the final maturity of the structure.
- The probability measure \mathbb{P} assigns probabilities to the events in the σ -algebras \mathcal{F}_t , $t = 1, \dots, T$. For example, the probability that up to time t more than 20% of the collateral securities defaulted is given by $\mathbb{P}(F)$, where $F \in \mathcal{F}_t$ is the corresponding measurable event.

Step 2 defines a *random variable* \vec{X} , because as soon as a scenario $\omega \in \Omega$ is fixed by the simulation engine, the distribution of cash flows conditional on ω follows a *deterministic* workflow defined by the documentation of the structure. The variable \vec{X} is a vector whose components contain the quantities relevant for the performance of the structure, e.g., realized returns for notes investors, the amount of realized repayments, the coupon payments made to notes investors, etc. The *distribution* $\mathbb{P} \circ \vec{X}^{-1}$ of the “performance vector” \vec{X} then is the final output, which has to be analyzed and interpreted in Step 3. For example, the *relative frequency* of scenarios in which at least one promised dollar to a mezzanine investor has not been paid, constitutes the *default probability* of that mezzanine tranche.

The filtration $(\mathcal{F}_t)_{t=1,\dots,T}$ defines a *dynamic information flow* during the simulated lifetime of the deal. For example, the simulation step from time t to time $t + 1$ will always be *conditioned* on the already realized *path* (the “history” up to time t). This very much reflects the approach an investor would follow during the term of a structure: At time t she or he will take all available information up to time t into account for making an analysis regarding the future performance of the structure.

8.3.2 Correlated Default Time Models

The multi-step model is a straightforward extension of the one-period models we discussed in previous chapters to a multi-period simulation model. Another “best practice” approach is to generate *correlated default times* of the collateral securities. We already discussed this approach in [Section 7.3](#). The correlated default times approach calibrates

default times compatible to a given one-year horizon asset value model by means of credit curves, assigned to the default probability of the collateral securities, and some copula function, generating a multivariate dependency structure for the single default times. It is not by chance that this approach already has been used for the valuation of *default baskets*: Focussing only on defaults and not on rating migrations, the collateral pool (or reference portfolio) of a CDO can be interpreted as a somewhat large default basket. The only difference is the cash flow model on top of the basket.

From a simulation point of view, the default times approach involves much less random draws than a multi-step approach. For example, a multi-step model w.r.t. a collateral pool consisting of 100 bonds, would for quarterly payments over 10 years require $100 \times 10 \times 4$ simulated random draws in every scenario. The same situation by means of a default times approach would only require to simulate 100 random draws in a scenario, namely realizations of 100 default times for 100 bonds. This saves computation time, but has the disadvantage that *rating distributions* (e.g., for modeling rating triggers) can not be incorporated in a straightforward manner as it is the case in multi-step models.

Time-consuming calculations in the default times approach could be expected in the part of the algorithm inverting the credit curve $F(t)$ in order to calculate default times according to the formula $\tau = F^{-1}(N[r])$; see [Section 7.3](#). Fortunately, for CDO models the exact time when a default occurs is not relevant. Instead, the only relevant information is if an instrument defaults between two consecutive payment dates. Therefore, the copula function approach for default times can be easily discretized by calculating thresholds at each payment date $t = 1, \dots, T$ according to

$$\alpha_t = N^{-1}[F(t)] ,$$

where F denotes the credit curve for some fixed rating, and $N[\cdot]$ denotes the cumulative standard normal distribution function. Clearly one has

$$\alpha_1 < \alpha_2 < \dots < \alpha_T .$$

Setting $\alpha_0 = -\infty$, asset i defaults in period t if and only if

$$\alpha_{t-1} < r_i \leq \alpha_t ,$$

where $(r_1, \dots, r_m) \sim N(0, \Gamma)$ denotes the random vector of standardized asset value log-returns with asset correlation matrix Γ . This reduces

the computational efforts substantially, since the thresholds have to be calculated only once and can then be stored in a look-up table *before* the actual random events are simulated. Figure 8.7 depicts the workflow of a CDO model based on default times.

8.3.3 Stochastic Default Intensity Models

The stochastic intensity approach [29,31] is a time continuous model and has already been presented in Section 2.4.4. Duffie and Gârleanu [29] studied a stochastic intensity approach to the valuation of CDOs by considering a basic affine process for the intensity λ , solving a stochastic differential equation of the form

$$d\lambda(t) = \kappa(\theta - \lambda(t))dt + \sigma\sqrt{\lambda(t)}dB(t) + \Delta J(t), \quad (8.3)$$

where B is a Wiener process and J is a pure-jump process, independent of B . In the course of their paper, they consider a simple subordinated structure, consisting of only three tranches: An equity piece, a mezzanine tranche, and a senior tranche. They experimented with different overcollateralization levels and different correlations and showed that *correlations significantly impact the market value of individual tranches*. For example, in cases where the senior tranche has only a small cushion of subordinated capital, the market value of the senior tranche decreases with decreasing correlation, whereas the market value of the equity piece increases with increasing correlation. Their calculations further show that this effect can be mitigated, but not removed, by assuming a higher level of overcollateralization. Regarding the behaviour of the mezzanine tranche in dependence on a changing correlation, they find that the net effect of the impact of correlation changes on the market value of the senior and equity tranche is *absorbed by the mezzanine tranche*. This interestingly results in an ambiguous behaviour of the mezzanine tranche: Increasing default correlation may raise or lower the mezzanine spreads.

For a practical implementation, the stochastic differential equation (2.47) has to be solved *numerically* by discretization methods, i.e., the intensity is integrated in appropriately small time steps. Unfortunately, this procedure can be quite *time-consuming* compared to other CDO modeling approaches.

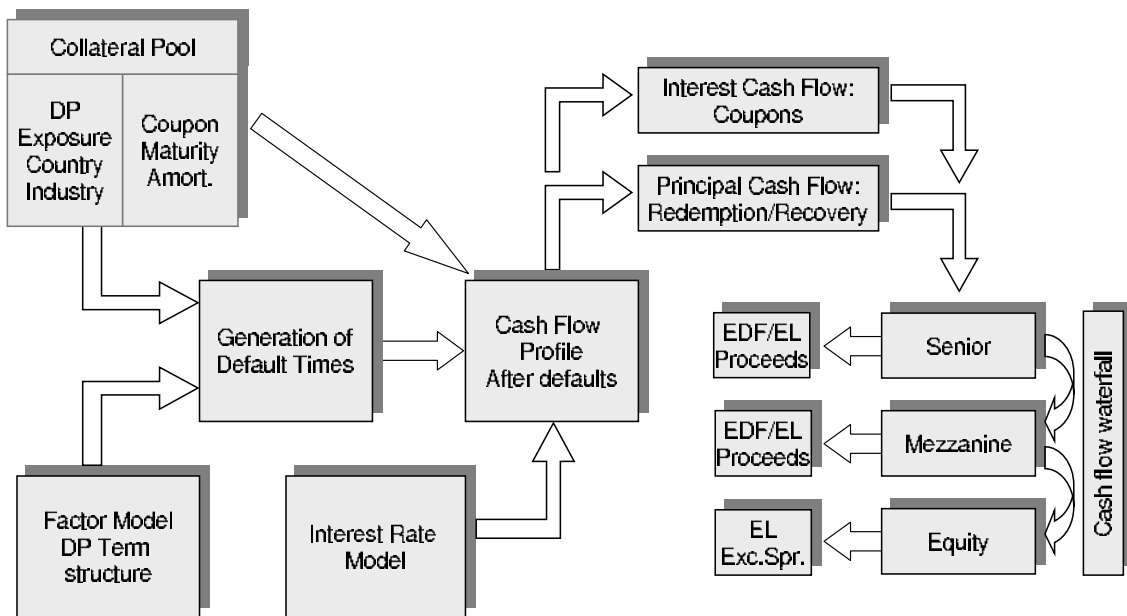


FIGURE 8.7
CDO modeling workflow based on default times.

8.4 Rating Agency Models: Moody's BET

At the beginning of this section, we should remark that all three major rating agencies (Moody's, S&P, and Fitch) have their own models for the valuation of ABS structures. Moreover, before a deal is launched into the market it typically requires at least two external ratings from two of the just mentioned agencies, some deals even admit a rating from all three of them. In this section we give an example for a rating agency model by means of discussing Moody's so-called *binomial expansion technique* (BET).

Moody's rating analysis of CDOs is based on the following idea: Instead of calculating the loss distribution of the original collateral portfolio of a CDO, Moody's constructs a *homogeneous comparison portfolio* satisfying the following conditions:

- All instruments have equal face values, summing up to the collateral pool's total par value.
- All instruments have equal default probability p , calibrated according to the *weighted average rating factor* (WARF), assigned to the portfolio by means of Moody's rating analysis.
- The instruments in the comparison portfolio are independent.

Moody's calibrates such a homogenous portfolio to any given pool of loans or bond, taking the rating distribution, exposure distribution, industry distribution, and the maturities of the assets into account.

Then, according to the assumptions made, the portfolio loss of the homogeneous comparison portfolio follows a *binomial distribution*; see also [Chapter 2](#).

The crucial parameter in this setting is the number n of instruments in the comparison portfolio. This parameter constitutes a key measure of diversification in the collateral pool developed by Moody's and is therefore called Moody's *diversity score* (DS) of the collateral portfolio. Regarding diversification, Moody's makes two additional assumptions:

- Every instrument in the comparison portfolio can be uniquely assigned to one industry group.

TABLE 8.6: Moody's Diversity Score; see[88]

Number of Firms in Same Industry	Diversity Score
1	1.00
2	1.50
3	2.00
4	2.33
5	2.67
6	3.00
7	3.25
8	3.50
9	3.75
10	4.00
<i>Source: Moody's Investors Services</i>	

- Two instruments in the comparison portfolio have positive correlation if and only if they belong to the same industry group.

Based on this assumption, the only driver of diversification is the industry distribution of the collateral pool. Table 8.6 reports the diversity score for different industry groupings¹⁶. The diversity score of a portfolio is then calculated by summing up the diversity scores for the single industries represented in the collateral pool. For illustration purposes, let us calculate two sample constellations.

1. Consider 10 bonds from 10 different firms, distributed over 3 industries:
 - 2 firms in industry no.1, yielding a diversity score of $DS_1 = 1.50$
 - 3 firms in industry no.2, yielding a diversity score of $DS_2 = 2.00$
 - 5 firms in industry no.3, yielding a diversity score of $DS_3 = 2.67$
 The portfolio's total diversity score equals
 $DS = DS_1 + DS_2 + DS_3 = 6.17$.
2. Consider 10 bonds from 10 different firms, distributed over 10 industries:
 - 10 times one firm in one single industry means
 - 10 times a diversity score of 1, such that the portfolio's total diversity score sums up to $DS = 10$.

¹⁶For more than 10 instruments in one industry group, the diversity score is determined by means of a case-by-case evaluation.

The industry distribution of Constellation 2 leads to an obviously better industry diversification, and therefore yields a higher diversity score. Altogether, Moody's distinguishes between 33 industry groups, yielding¹⁷ a best possible diversity score of $132 = 33 \times 4$.

The loss distribution of the homogeneous comparison portfolio is assumed to be binomially distributed with parameters DS and WARF, $L \sim B(\text{DS}, \text{WARF})$, such that the probability of k defaults in the comparison portfolio equals

$$\mathbb{P}[L = k] = \frac{(\text{DS})!}{k!(\text{DS} - k)!} (\text{WARF})^k (1 - \text{WARF})^{\text{DS} - k},$$

where k ranges from 0 to DS. Based on the so obtained loss distribution, *cash flow scenarios* are evaluated in order to determine the *rating of a tranche*. Dependent on how many losses in the collateral pool a tranche can bear without suffering a loss due to the credit enhancement mechanisms of the structure, the tranche gets assigned a rating reflecting its "default remoteness". For example, senior notes have to pass much stronger stress scenarios without suffering a loss than junior or mezzanine notes.

From time to time CDO tranches are *down- or upgraded* by the rating agencies, because their default remoteness decreased or increased. For example, last and this year we saw many downgrades of CDO tranches, sometimes downgraded by more than one notch on the respective rating scale, due to the heavy recession in the global economy.

In a next step, we now want to consider the BET from a more mathematical point of view. For this purpose we consider a sample portfolio of m bonds, all bonds having the same default probability p and equal face values. Additionally we assume that the pairwise default correlation¹⁸ of the bonds is uniform for the whole portfolio and given by r . Our modeling framework is a uniform Bernoulli mixture model, with asset values as latent variables, as introduced in [Section 2.5.1](#). According to Equation 2.10 and Proposition 2.5.1, the corresponding uniform asset correlation ϱ of the model can be calculated by solving

$$r = \frac{N_2[N^{-1}[p], N^{-1}[p]; \varrho] - p^2}{p(1 - p)}$$

¹⁷Ignoring deviations from [Table 8.6](#) due to special case-by-case evaluations.

¹⁸In contrast to the rest of this book we here denote the default correlation by r .

for ϱ . For example, for $r = 3\%$ and $p = 1\%$ we calculate $\varrho = 23.06\%$. Recall that the uniform Bernoulli mixture model is completely determined by specifying p and r (respectively ϱ).

In Proposition 2.5.7 we already discussed the two extreme cases regarding ϱ . In case of $\varrho = 0$, the distribution of the portfolio loss is binomial, $L \sim B(m, mp)$. In case of $\varrho = 1$, the loss distribution is of Bernoulli type, $L \sim B(1, p)$. Both extreme case distributions are *binomial* distributions with probability p . Looking at the respective first parameter of both distributions, we discover m bonds in the first case and 1 bond in the second case. The idea of the BET now is to introduce also the intermediate cases by establishing a relation between the assumed level of correlation and the number of bonds in a homogeneous comparison portfolio. More formally, for a given portfolio of m bonds, the BET establishes a functional relation

$$n : [0, 1] \rightarrow \{0, 1, \dots, m\}, \quad r \mapsto n(r),$$

between the default correlation and the number of bonds in a homogeneous portfolio of *independent* bonds with binomial loss distribution.

The function n can be determined by a matching of first and second moments. The first moments of both portfolios must be equal to p . The second moment of the original portfolio can be calculated as

$$\begin{aligned} \mathbb{V}[L^{(m)}] &= \frac{1}{m^2} \mathbb{V}\left[\sum_{i=1}^m L_i\right] = \frac{1}{m^2} \sum_{i,j=1}^m \text{Cov}[L_i, L_j] = \\ &= \frac{1}{m^2} \left(mp(1-p) + \sum_{i \neq j} rp(1-p) \right) = \\ &= \frac{mp(1-p) + m(m-1)rp(1-p)}{m^2} = \frac{p(1-p)(1 + (m-1)r)}{m}. \end{aligned}$$

The variance of the homogeneous comparison portfolio, consisting of $n(r)$ independent bonds, equals

$$\mathbb{V}[L^{(n(r))}] = \frac{p(1-p)}{n(r)}.$$

Matching both second moments finally yields

$$n(r) = \frac{m}{1 + r(m-1)}. \tag{8.4}$$




FIGURE 8.8
Diversification Score as a function of m for $r = 3\%$.

This number $n(r)$ is not necessarily an integer value, so that we have to round it to the closest integer. The so obtained number is comparable to Moody’s diversity score. In order to distinguish¹⁹ the two scores, we call $n(r)$ the *diversification score* of the original portfolio. Figure 8.8 shows the diversification score $n(r)$ in dependence of m for $r = 3\%$. Two facts illustrated by the plot follow immediately from Equation 8.4:

1. The diversification score is independent of the credit quality of the pool, captured by the default probability p .
2. The diversification score is bounded from above by $1/r$. For $r = 3\%$, the maximum diversification score is $DS = 33$, which is attained for $m \geq 1,261$. The reason for an upper bound of the diversification score w.r.t. a fixed default correlation comes from the fact that only *specific risk* can be eliminated by diversification. *Systematic risk* remains in the portfolio, no matter by how many obligors we enlarge the portfolio.

We now want to compare the loss distributions of a fictitious sample portfolio and the homogeneous portfolio of independent bonds fitted to

¹⁹Note that Moody’s diversity score purely relies on industry diversification.

the original portfolio by means of the BET. We assume that the original portfolio contains $m = 100$ bonds with uniform default probability $p = 1\%$ and uniform default correlation $r = 3\%$. As already mentioned, these assumptions yield an asset correlation of $\rho = 23.06\%$.

According to Equations 2. 8 and 2. 49, the probability for k defaults in the original bond portfolio is given by

$$\mathbb{P}[L^{(100)} = k] = \binom{100}{k} \int_0^1 p(y)^k (1 - p(y))^{100-k} dN(y) ,$$

$$\text{where } p(y) = N\left[\frac{N^{-1}(0.01) - \sqrt{0.2306}y}{\sqrt{1 - 0.2306}}\right] .$$

Therefore, we can easily calculate the loss distribution of the original portfolio. Next, we calculate the diversification score of the original portfolio. According to Equation 8. 4, we obtain

$$n(3\%) = \frac{100}{1 + 3\%(100 - 1)} = \frac{100}{1 + 2.97} = 25.19 ,$$

such that the diversification score after rounding equals 25. Therefore, the loss distribution of the homogeneous comparison portfolio follow a binomial distribution, $L^{(25)} \sim B(25, 2.5)$. So here the BET claims that *25 independent bonds carry the same risk as 100 bonds with default correlation $r = 3\%$.*

Figure 8.9 compares the original loss distribution with the BET-fitted binomial distribution. The plot clearly shows, that the BET-fit significantly underestimates the *tail probabilities* of the original loss distribution.

This does not come much as a surprise, because due to the *central limit theorem* binomial distributions tend to be approximately normal for a large number of bonds, whereas typical credit portfolio loss distributions are skewed with fat tails. Moreover, it is generally true that moment matching procedures do not automatically also fit the tails of the considered distributions in an accurate manner.

Now we come to an important conclusion: Because the BET significantly underestimates the tail probabilities of the original portfolio, *the risk of senior notes will typically be underestimated by the BET approach.*

To make this more explicite, we consider a situation like illustrated in Figure 8.10. Assume that the plot shows the probability density of




FIGURE 8.9
Fitting a loss distribution by means of the BET (original uniform portfolio: $p = 1\%$, $r = 3\%$, $m = 100$); note that the y-axes is logarithmically scaled.




FIGURE 8.10
Tranching a Loss Distribution.

the distribution of the cumulative net losses L of some collateral pool, calculated over the whole term of the structure. Let us further assume that the bank invested in an upper mezzanine respectively lower senior tranche $T_{[\alpha_1, \alpha_2]}$ with lower respectively upper bound α_1 respectively α_2 . Then, the default probability (DP) of this tranche and its expected loss (EL) can be calculated²⁰ as

$$\text{DP}(T_{[\alpha_1, \alpha_2]}) = \mathbb{P}[L > \alpha_1] ,$$

$$\text{EL}(T_{[\alpha_1, \alpha_2]}) = \frac{1}{\alpha_2 - \alpha_1} \int \min\left(\max[x - \alpha_1, 0], \alpha_2 - \alpha_1\right) d\mathbb{P}_L(x) ,$$

where \mathbb{P}_L denotes the probability density of L . If we now would replace the loss distribution \mathbb{P}_L by a binomial distribution fitted to \mathbb{P}_L by means of the BET, we can expect that $\text{DP}(T_{[\alpha_1, \alpha_2]})$ and $\text{EL}(T_{[\alpha_1, \alpha_2]})$ will be significantly lower; see [Figure 8.9](#). A moment-matched binomial distribution will not appropriately capture the risk of a tranche more outside in the tail, like it is the case for $T_{[\alpha_1, \alpha_2]}$.

Our discussion has far-reaching consequences. Whenever a bank intends to invest in a senior note, the model the bank uses for the evaluation of the investment should capture the *tail risk* of the collateral pool. But the tail risk of the collateral pool is driven by the *correlation* inherent in the collateral portfolio. The higher the overall correlation, the larger the tail probabilities and therefore the potential for losses in senior pieces of the structure. Because the bank wants to be compensated for taking this risk, it can not rely on the BET or other methods ignoring the *skewed fat-tailed* character of credit portfolio loss distributions. Only a full Monte Carlo simulation of an appropriate portfolio model, combined with a sound modeling of all relevant cash flow elements of the structure, will really show how much premium payment the bank needs to be compensated for the taken risk and to make some profit at the end.

²⁰Of course, certain cash flow elements in a structure can distort the “direct” effect of losses on a particular tranche, as we claimed it here, but for reasons of simplicity we ignore this greater complexity for the moment. However, in synthetic CLOs, where the performance of notes is linked to the performance of a reference pool (e.g. by means of credit-linked notes) this simplified view is very close to the truth.

8.5 Conclusion

In this chapter we introduced the motivations for participating in CDO transactions, considering both, the originator's and the investor's point of view. Additionally we explained some of the common cash flow elements common in these transactions. We made clear, that only a combination of a sound multi-period portfolio model *and* an exhaustive modeling of the cash flow elements of a structure will provide by means of a Monte Carlo simulation reliable information about the chances and risks of a CDO engagement. Hereby, the choice of the underlying portfolio model essentially drives the risk measured for different tranches. For example, models ignoring fat tails, like the BET or other rating agency approaches, tend to underestimate the risk potential of senior notes. To give another example, because the tails of Poisson mixture models typically are less fat than the tails of comparably calibrated Bernoulli mixture models (see [Chapter 2](#)), an analyst relying on a Poisson approach will find some senior tranche investment less risky than an analyst working with a Bernoulli approach.

We conclude this chapter by a last comment. In our opinion, mathematically rigorous *and* applicable models for complex structured products will provide a great challenge for the next years. Risk transfer by credit derivatives or ABS structures becomes more and more the standard daily business of banks all over the world. since the regulatory regimes will become more risk sensitive in the future, regulatory arbitrage opportunities accordingly require refined methods for measuring mismatches between regulatory and economic risk capital.

8.6 Some Remarks on the Literature

Regarding the literature on ABS structures, the rating agencies provide the richest source of information material. Among the many interesting papers, we especially recommend the following "collection":

- **Standard & Poor's:**
 - Global CBO/CLO criteria [116].

- Global synthetic securities criteria [117] (derivatives focussed).

- **Moody’s Investors Services:**

- Reference to the rating procedure for cash flow transactions [88] (includes reference to expected loss ratings and diversity scores).
- Papers on the BET for CDOs and multisector CDOs [89,90,93].
- Special report on downgrading CBO tranches [91] (case study style).
- Discussion of risk transfer through securitizations [92] (case study style).
- CDO market overview (Moody’-rated) [96] (report is regularly updated; very valuable overview for investors).
- Multisector CDO ratings: Moody’s approach [94].

- **Fitch:**

- Introduction to CLOs [38].
- Rating criteria for cash flow transactions [39].
- Brief introduction to synthetic CDOs [40].
- Discussion of Fitch’s evaluation of synthetic securitizations [41].

Besides the rating agency reports, which are regularly updated and from which we just mentioned a very small subset, most investment banks provide useful introductions to CDOs. Here we just list a small sample collection of papers.

- **Bank of America:**

- Introduction to arbitrage CDOs, cash flow and synthetic [6].

- **J.P.Morgan:**

- A “CDO Handbook”, providing a good overview and introduction to CDOs [69].
- The JPM guide to credit derivatives, also introducing synthetic CDOs [68].

- **Deutsche Bank:**

- CDO introduction [27]
- Introduction to the use of synthetic CLOs for balance sheet management [25].

Regarding the academic literature we do not see as many publications for CDO modeling as we see them for portfolio models. However, there are certainly more papers “out there” than we have seen so far. Again, a small collection is included.

- Duffie and Gârleanu [29] study the impact of overcollateralization and correlation changes on CDO tranches by means of a dynamic intensity approach.
- Boscher and Ward [14] discuss the pricing of synthetic CDOs and default baskets by a copula function approach (correlated default times).
- Skarabot [112] studies asset securitization in relation to the asset structure of a firm.
- Finger [37] studies differences arising from evaluating CDOs by different modeling approaches.

There are certainly more books and articles we could recommend to interested readers. For example, the fixed income literature also includes many books on ABS, e.g., the book [35] by Fabozzi. Additionally, many books on credit risk management contain passages about CDOs and related instruments.

References

- [1] C. Acerbi and D. Tasche. On the coherence of expected shortfall. <http://www-m4.mathematik.tu-muenchen.de/m4/pers/tasche/page.shtml>, 2000.
- [2] E. I. Altman, A. Resti, and A. Sironi. Analyzing and explaining default recovery rates. A report submitted to The International Swaps & Derivatives Association, December 2001.
- [3] M. Araten and L. Angbazo. Roots of transition matrices: Application to settlement risk. Chase Manhattan Bank, 1997. Practical Paper.
- [4] P. Artzner, F. Delbaen, J. Eber, and D. Heath. Thinking coherently. *Risk*, 10(11):68–71, 1997.
- [5] P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.
- [6] Banc of America Securities. *Arbitrage CDO Primer*, May 2001.
- [7] H. Bauer. *Probability Theory*. de Gruyter, 1996.
- [8] M. Baxter and A. Rennie. *Financial Calculus. An introduction to derivative pricing*. Cambridge University Press, 1996.
- [9] N. H. Bingham and R. Kiesel. *Risk-Neutral Valuation. Pricing and Hedging of Financial Derivatives*. Springer Finance. Springer-Verlag, 3rd edition, 2000.
- [10] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, (81):637–654, 1973.
- [11] C. Bluhm and L. Overbeck. Estimating systematic risk in uniform credit portfolios. Working paper, Deutsche Bank Group, December 1999.
- [12] C. Bluhm, L. Overbeck, and C. Wagner. Irreconcilable differences. *RISK*, 14(10):S33–S37, October 2001.
- [13] J. R. Bohn. A survey of contingent-claims approaches to risky debt valuation. Working paper, June 1999.
- [14] H. Boscher and Ian Ward. Long or short in cdos. *Risk*, 15(6):125–129, June 2002.
- [15] N. L. Bowers, H. U. Gerber, J. C. Hickman, D. A. Jones, and C. J. Nesbitt. *Actuarial Mathematics*. Schaumburg, Illinois, 2nd edition, 1997. Society of Actuaries.

- [16] D. R. Cox and D. Oakes. *Analysis of Survival Data*. Chapman and Hall, 1984.
- [17] J.C. Cox, J.E. Ingersoll and S.A. Ross. A theory of term structure of interest rates. *Econometrica*, 53:385–407, 1985.
- [18] Credit Suisse Financial Products. *CreditRisk⁺ – A Credit Risk Management Framework*, 1997.
- [19] P. Crosbie. Modeling default risk. KMV Corporation, <http://www.kmv.com>, 1999.
- [20] M. Crouhy, D. Galai, and R. Mark. A comparative analysis of current credit risk models. *Journal of Banking & Finance*, 24:59–117, 2000.
- [21] M. Crouhy, D. Galai, and R. Mark. *Risk Management*. McGraw-Hill, 2000.
- [22] M. Crouhy, D. Galai, and R. Mark. Prototype risk rating system. *Journal of Banking & Finance*, 25:41–95, 2001.
- [23] S.R. Das. Structured notes and derivatives embedded securities. Euromoney Publications PLC, 1996.
- [24] M. Denault. Coherent allocation of risk capital. <http://www.risklab.ch/Papers.html>, 1999.
- [25] Deutsche Bank Fixed Income Research. *Synthetic CLOs and their Role in Bank Balance Sheet Management*, March 1999.
- [26] Deutsche Bank Fixed Income Research. *Credit Derivatives and Structured Credit*, August 2000.
- [27] Deutsche Bank Securitization. *Structured Product CDOs*, March 2001.
- [28] D. Duffie. *Dynamic Asset Pricing Theory*. Princeton University Press, 1992.
- [29] D. Duffie and N. Gârleanu. Risk and valuation of collateralized debt obligations. Working paper, 1999. Graduate School of Business, Stanford University.
- [30] D. Duffie and K. Singleton. An econometric model of the term structure of interest-rate swap yields. *Journal of Finance*, 51(1):1287–1321, 1997.

- [31] D. Duffie and K. Singleton. Simulating correlated defaults. Working paper, 1998. Graduate School of Business, Stanford University.
- [32] D. Duffie and K. Singleton. Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12:687–720, 1999.
- [33] E. Eberlein and U. Keller. Hyperbolic distributions in finance. *Bernoulli*, 1:281–299, 1995.
- [34] P. Embrechts, A. McNeil, and D. Straumann. Correlation and dependence in risk management: Properties and pitfalls. Preprint, July 1999. ETH Zürich.
- [35] F. J. Fabozzi. *Investing in Asset-Backed Securities*. McGraw Hill, 2001.
- [36] C. C. Finger. Conditional approaches for creditmetrics portfolio distributions. *CreditMetricsTM Monitor*, April 1999.
- [37] C. C. Finger. A comparison of stochastic default rate models. *RiskMetrics Journal*, pages 49–73, November 2000.
- [38] Fitch IBCA, Duff & Phelps. *Bank Collateralized Loan Obligations: An Overview*, December 1997.
- [39] Fitch IBCA, Duff & Phelps. *Rating Criteria for Cash Flow Collateralized Debt Obligations*, November 2000.
- [40] Fitch IBCA, Duff & Phelps. *Synthetic CDOs: A Growing Market for Credit Derivatives*, February 2001.
- [41] Fitch IBCA, Duff & Phelps. *Synthetic Securitization: Unique Structures, Analytic Challenges*, June 2001.
- [42] H. Föllmer. Calcul d’Itô sans Probabilités. volume 850 of *Lecture Notes in Mathematics*, pages 143–150. Springer, New York, 1981. Seminaire Probabilités, Strasbourg XV.
- [43] J. Fons. Using default rates to model the term structure of credit risk. *Financial Analysts Journal*, pages 25–32, 1994.
- [44] E. W. Frees and E. A. Valdez. Understanding relationships using copulas. 1997. Paper presented at the 32nd Actuarial Research Conference, held August 6-8, at the University of Calgary, Alberta, Canada.

- [45] R. Frey and A. J. McNeil. Modelling dependent defaults. Draft, March 2001.
- [46] R. Frey and A. J. McNeil. Var and expected shortfall in portfolios of dependent credit risks: Conceptual and practical insights. <http://www.math.ethz.ch/~mcneil/ftp/rome.pdf>, 2002.
- [47] R. Frey, A.J. McNeil, and M. Nyfeler. Copulas and credit models. *RISK*, 14(10):111–114, 2001.
- [48] J. Frye. Collateral damage. *RISK*, 13(4):91–94, 2000.
- [49] J. Frye. Depressing recoveries. *RISK*, 13(11):108–111, 2000.
- [50] P. Georges, A-G. Maly, E. Nicolas, G. Quibel, and T. Roncalli. Multivariate survival modelling: a unified approach with copulas. <http://gro.creditlyonnais.fr>, 1996.
- [51] M. B. Gordy. A comparative anatomy of credit risk models. *Journal of Banking and Finance*, 24:119–149, 2000.
- [52] M. B. Gordy. A risk-factor model foundation for ratings-based bank capital rules. Draft, February 2001.
- [53] G. R. Grimmet and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, 2nd edition, 1992.
- [54] G. M. Gupton, C. C. Finger, and M. Bhatia. CreditMetrics – Technical Document. Morgan Guaranty Trust Co., http://www.defaultrisk.com/pp_model_20.htm, 1997.
- [55] J.M. Harrison and S. Pliska. Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and Their Applications*, 11:215–260, 1981.
- [56] D. Heath, R. Jarrow, and A. Morton. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60:77–105, 1992.
- [57] J.C. Hull. *Option, Futures, and Other Derivatives*. Prentice Hall, 4th edition, 2000.
- [58] J.C. Hull and A. White. Valuing credit default swaps II: Modeling default correlations. <http://www.defaultrisk.com>.
- [59] J.C. Hull and A. White. The impact of default risk on options and other derivatives. *Journal of Banking and Finance*, 19(2):299–322, 1995.

- [60] J.C. Hull and A. White. Valuing credit default swaps I: No counterparty default risk. *Journal of Derivatives*, 8(1):29–40, 2000.
- [61] International Swap and Derivatives Association Inc. *Credit derivatives documentation*, 1998.
- [62] R. Israel, J. Rosenthal, and J. Wei. Finding generators for Markov chains via empirical transition matrices with application to credit ratings. *Mathematical Finance*, 11(2):245–265, 2001.
- [63] J. James. Total return swaps. *The Electronic Journal of Financial Risk*, 1(2), 1998.
- [64] R.A. Jarrow, D. Lando, and S.M. Turnbull. A Markov model for the term structure of credit risk spreads. *Review of Financial Studies*, 10:481–523, 1997.
- [65] R.A. Jarrow and S.M. Turnbull. Pricing options on financial securities subject to default risk. *Journal of Finance*, 50:53–86, 1995.
- [66] S. Jaschke and U. Küchler. Coherent risk measures, valuation bounds, and (μ, σ) -portfolio optimization. *Finance and Stochastics*, 5(2):181–200, 2000.
- [67] H. Joe. *Multivariate Models and Dependence Concepts*. Chapman & Hall, 1997.
- [68] J.P. Morgan Securities. *The J.P. Morgan guide to credit derivatives*, 1999.
- [69] J.P. Morgan Securities. *CDO Handbook*, May 2001.
- [70] O. Kallenberg. *Foundations of Modern Probability*. Probability and its Applications. Springer-Verlag, 1997.
- [71] I. Karatzas and S. E. Shreve. *Brownian Motion and Stochastic Calculus*. Springer, 2nd edition, 1991.
- [72] I. Karatzas and S. E. Shreve. *Methods of Mathematical Finance*. Springer, 1998.
- [73] S. Kealhofer, S. Kwonk, and W. Weng. Uses and abuses of bond default rates. KMV Corporation, 1998.
- [74] H. U. Koyluoglu and A. Hickman. Reconcilable differences. *RISK*, 56, October 1998.

- [75] A. Kreinin and M. Sidelnikova. Regularization algorithms for transition matrices. *Algo Research Quarterly*, 4(1/2):25–40, 2001.
- [76] D. Lamberton and B. Lapeyre. *Introduction to Stochastic Calculus Applied to Finance*. Chapman & Hall, 1996.
- [77] D. Lando. *Modelling Bonds and Derivatives with Default Risk*. Mathematics of Derivatives Securities, M. Dempster and S. Pliska, eds. Cambridge University Press, 1997.
- [78] D. X. Li. The valuation of basket credit derivatives. *CreditMetricsTM Monitor*, pages 34–50, April 1999.
- [79] D. X. Li. On default correlation: A copula function approach. *Journal of Fixed Income*, 6:43–54, March 2000.
- [80] F. Longstaff and E. Schwartz. A simple approach to valuing risky fixed and floating rate debt. *Journal of Finance*, 50:789–819, 1995.
- [81] F. Longstaff and E. Schwartz. Valuing credit derivatives. *Journal of Fixed Income*, pages 6–12, June 1995.
- [82] R. Martin, K. Thompson, and C. Browne. Taking to the saddle. *Risk*, 14(7):91–94, June 2001.
- [83] R. Mashal and M. Naldi. Extreme vents and default baskets. *Risk*, 15(6):119–122, June 2002.
- [84] B. Masters. Credit derivatives and the management of credit risk. *The Electronic Journal of Financial Risk*, 1(2), 1998.
- [85] McKinsey & Company, German Office. *CreditPortfolioView 2.0*, June 2001. Technische Dokumentation.
- [86] R. Merton. On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29:449–470, 1974.
- [87] T. Mikosch. *Elementary Stochastic Calculus – with Finance in View*, volume 6 of *Advanced Series on Statistical Science & Applied Probability*. World Scientific Publishing, 1998.
- [88] Moody’s Investors Service. *Rating Cash Flow Transactions Backed by Corporate Debt 1995 Update*, April 1995.
- [89] Moody’s Investors Service. *The binomial expansion method applied to CBO/CLO*, December 1996.

- [90] Moody's Investors Service. *The double binomial method and its application to a special case of CBO structure*, March 1998.
- [91] Moody's Investors Service. *Stability of Ratings of CBO/CLO Tranches: What Does It Take to Downgrade a CBO Tranche?*, November 1998.
- [92] Moody's Investors Service. *Another Perspective on Risk Transference and Securitization*, July 1999.
- [93] Moody's Investors Service. *Moody's Approach to rating multisector CDOs*, September 2000.
- [94] Moody's Investors Service. *Moody's Approach to Rating Multi-sector CDOs*, September 2000.
- [95] Moody's Investors Service. *Default and Recovery Rates of Corporate Bond Issuers: 2000*, February 2001.
- [96] Moody's Investors Service. *Collateralized Debt Obligations Performance Overview Compilation*, March 2002.
- [97] D. Murphy. Keeping credit under control. *Risk*, 9, September 1996.
- [98] S.R. Neal. Credit derivatives: New financial instruments for controlling credit risk. *Economic Review*, 1996.
- [99] R. Nelsen. *An Introduction to Copulas*. Springer, New York, 1999.
- [100] C. Nelson and A. Siegel. Parsimonious modeling of yield curves. *Journal of Business*, 60:473–489, 1987.
- [101] P. Nickell, W. Perraudin, and S. Varotto. Ratings- versus equity-based credit risk modeling: an empirical analysis. <http://www.bankofengland.co.uk/workingpapers/>, 1999. Working paper.
- [102] J. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.
- [103] Basel Committee on Banking Supervision. *The Internal Ratings-Based Approach*. Supporting Document to the New Basel Capital Accord, 2001.
- [104] M. K. Ong. *Internal Credit Risk Models*. Risk Books, 1999.

- [105] L. Overbeck. Allocation of economic capital in loan portfolios. In U. Franke, W. Härdle, and G. Stahl, editors, *Measuring Risk in Complex Stochastic Systems*. Springer, New York, 2000.
- [106] W. R. Pestman. *Mathematical Statistics*. de Gruyter, 1998.
- [107] PriceWaterhouseCoopers. *The PriceWaterhouseCoopers Credit Derivatives Primer*, 1999.
- [108] D. Revuz and M. Yor. *Continuous Martingales and Brownian Motion*. Springer-Verlag, 1991. Chapter IV (3.13).
- [109] J. A. Rice. *Mathematical Statistics and Data Analysis*. Duxbury Press, 2nd edition, 1995.
- [110] W. Schmidt and I. Ward. Pricing default baskets. *Risk*, 15(1):111–114, January 2002.
- [111] Ph. J. Schoenbucher. Factor models for portfolio credit risk. *Preprint, University of Bonn, Germany*, 2001.
- [112] J. Skarabot. Asset securitization and optimal asset structure of the firm. *Working Paper*, July 2001.
- [113] A. Sklar. Fonction de repartition à n dimension et leur marges. *Publications de l'Institut Statistique de l'Université de Paris*, 8:229–231, 1959.
- [114] A. Sklar. Random variables, joint distribution functions and copulas. *Kybernetika*, 9:449–460, 1973.
- [115] J. R. Sobehart and S. C. Keenan. An introduction to market-based credit analysis. *Moody's Risk Management Services*, November 1999.
- [116] Standard & Poor's. *Global CBO/CLO Criteria*.
- [117] Standard & Poor's. *Global Synthetic Securities Criteria*.
- [118] Standard & Poor's. *Standard & Poor's Corporate Ratings Criteria 1998*.
- [119] W. Stromquist. Roots of transition matrices. Daniel H. Wagner Associates, 1996. Practical Paper.
- [120] D. Tasche. Risk contributions and performance measurement. <http://www.ma.tum.de/stat/>, 1999.

- [121] D. Tasche. Calculating value-at-risk contributions in CreditRisk+. <http://arxiv.org/abs/cond-mat/0112045>, 2000.
- [122] D. Tasche. Conditional expectation as quantile derivative. <http://www.ma.tum.de/stat/>, 2000.
- [123] O. A. Vasicek. Probability of loss on loan portfolio. KMV Corporation, 1987.
- [124] O.A. Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5:177–188, 1977.
- [125] S. S. Wang. Aggregation of correlated risk portfolios: Models & algorithms.
- [126] T. Wilde. IRB approach explained. *RISK*, 14(5), 2001.
- [127] T. Wilson. Portfolio credit risk, part i. *RISK*, pages 111–117, 1997.
- [128] T. Wilson. Portfolio credit risk, part ii. *RISK*, pages 56–61, 1997.
- [129] D. Zwillinger. *Handbook of Differential Equations*. Academic Press, Boston, 1995.