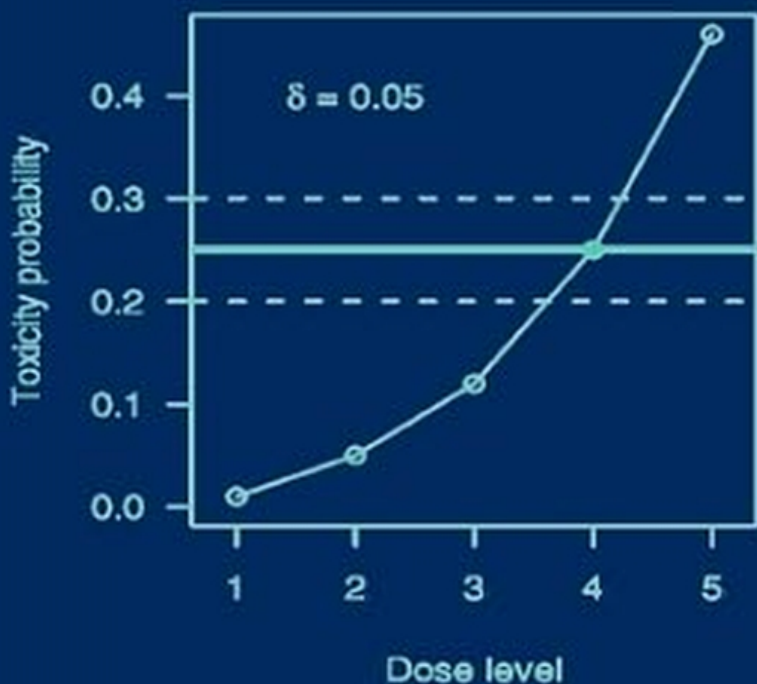


Chapman & Hall/CRC Biostatistics Series

Dose Finding by the Continual Reassessment Method



Ying Kuen Cheung



CRC Press

Taylor & Francis Group

A CHAPMAN & HALL BOOK

**Dose Finding
by the Continual
Reassessment Method**

Chapman & Hall/CRC Biostatistics Series

Editor-in-Chief

Shein-Chung Chow, Ph.D.

Professor

Department of Biostatistics and Bioinformatics

Duke University School of Medicine

Durham, North Carolina

Series Editors

Byron Jones

Senior Director

Statistical Research and Consulting Centre

(IPC 193)

Pfizer Global Research and Development

Sandwich, Kent, U.K.

Jen-pei Liu

Professor

Division of Biometry

Department of Agronomy

National Taiwan University

Taipei, Taiwan

Karl E. Peace

Georgia Cancer Coalition

Distinguished Cancer Scholar

Senior Research Scientist and

Professor of Biostatistics

Jiann-Ping Hsu College of Public Health

Georgia Southern University

Statesboro, Georgia

Bruce W. Turnbull

Professor

School of Operations Research

and Industrial Engineering

Cornell University

Ithaca, New York

Chapman & Hall/CRC Biostatistics Series

Adaptive Design Theory and Implementation Using SAS and R, Mark Chang

Advances in Clinical Trial Biostatistics, Nancy L. Geller

Applied Statistical Design for the Researcher, Daryl S. Paulson

Basic Statistics and Pharmaceutical Statistical Applications, Second Edition, James E. De Muth

Bayesian Adaptive Methods for Clinical Trials, Scott M. Berry, Bradley P. Carlin, J. Jack Lee, and Peter Muller

Bayesian Methods for Measures of Agreement, Lyle D. Broemeling

Bayesian Missing Data Problems: EM, Data Augmentation and Noniterative Computation, Ming T. Tan, Guo-Liang Tian, and Kai Wang Ng

Bayesian Modeling in Bioinformatics, Dipak K. Dey, Samiran Ghosh, and Bani K. Mallick

Causal Analysis in Biomedicine and Epidemiology: Based on Minimal Sufficient Causation, Mikel Aickin

Clinical Trial Data Analysis using R, Ding-Geng (Din) Chen and Karl E. Peace

Clinical Trial Methodology, Karl E. Peace and Ding-Geng (Din) Chen

Computational Methods in Biomedical Research, Ravindra Khattree and Dayanand N. Naik

Computational Pharmacokinetics, Anders Källén

Data and Safety Monitoring Committees in Clinical Trials, Jay Herson

Design and Analysis of Animal Studies in Pharmaceutical Development, Shein-Chung Chow and Jen-pei Liu

Design and Analysis of Bioavailability and Bioequivalence Studies, Third Edition, Shein-Chung Chow and Jen-pei Liu

Design and Analysis of Clinical Trials with Time-to-Event Endpoints, Karl E. Peace

Difference Equations with Public Health Applications, Lemuel A. Moyé and Asha Seth Kapadia

DNA Methylation Microarrays: Experimental Design and Statistical Analysis, Sun-Chong Wang and Arturas Petronis

DNA Microarrays and Related Genomics Techniques: Design, Analysis, and Interpretation of Experiments, David B. Allison, Grier P. Page, T. Mark Beasley, and Jode W. Edwards

Dose Finding by the Continual Reassessment Method, Ying Kuen Cheung

Elementary Bayesian Biostatistics, Lemuel A. Moyé

Frailty Models in Survival Analysis, Andreas Wienke

Handbook of Regression and Modeling: Applications for the Clinical and Pharmaceutical Industries, Daryl S. Paulson

Measures of Interobserver Agreement and Reliability, Second Edition, Mohamed M. Shoukri

Medical Biostatistics, Second Edition, A. Indrayan

Meta-Analysis in Medicine and Health Policy, Dalene Generalized Linear Models: A Bayesian Perspective, Dipak K. Dey, Sujit K. Ghosh, and Bani K. Mallick

Monte Carlo Simulation for the Pharmaceutical Industry: Concepts, Algorithms, and Case Studies, Mark Chang

Multiple Testing Problems in Pharmaceutical Statistics, Alex Dmitrienko, Ajit C. Tamhane, and Frank Bretz

Sample Size Calculations in Clinical Research, Second Edition, Shein-Chung Chow, Jun Shao, and Hansheng Wang

Statistical Design and Analysis of Stability Studies, Shein-Chung Chow

Statistical Methods for Clinical Trials, Mark X. Norleans

Statistics in Drug Research: Methodologies and Recent Developments, Shein-Chung Chow and Jun Shao

Statistics in the Pharmaceutical Industry, Third Edition, Ralph Buncher and Jia-Yeong Tsay

Translational Medicine: Strategies and Statistical Methods, Dennis Cosmatos and Shein-Chung Chow

Chapman & Hall/CRC Biostatistics Series

Dose Finding by the Continual Reassessment Method

Ying Kuen Cheung

Columbia University

New York, New York, USA



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

A CHAPMAN & HALL BOOK

Chapman & Hall/CRC
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2011 by Taylor and Francis Group, LLC
Chapman & Hall/CRC is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

Printed in the United States of America on acid-free paper
10 9 8 7 6 5 4 3 2 1

International Standard Book Number: 978-1-4200-9151-9 (Hardback)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

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 Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Preface

Despite its poor statistical properties, the 3+3 algorithm remains the most commonly used dose finding method in phase I clinical trials today. However, as clinicians begin to realize the important role of dose finding in the drug development process, there is an increasing openness to “novel” methods proposed in the past two decades. In particular, the continual reassessment method (CRM) and its variations have drawn much attention in the medical community. To ride on this momentum and overcome the status quo in the phase I practice, it is critical for us (statisticians) to be able to design a trial using the CRM in a timely and reproducible manner. This is the impetus to writing a detailed exposition on the calibration of the CRM for applied statisticians who need to deal with dose finding in phase I trials while having many other duties to attend to.

A natural approach to such a writing project is to write a *how-to* book. By the time I started this book project in the summer of 2008, I had helped design half a dozen CRM trials (three of which are included as examples in this book). In retrospect, I found some general patterns of how I calibrated the CRM parameters in these trials. These patterns, characterized collectively as a trial-and-error approach in Chapter 7, worked well in the sense that they gave reasonable operating characteristics to a design. However, it was time-consuming (weeks of simulation) and would require an intimate understanding of the CRM (I wrote a PhD dissertation on the CRM). I realized that some automation and step-by-step guidelines in this calibration process would be crucial and appreciated if the CRM was to be used on a regular basis by a wide group of statisticians. Chapters 7–10 try to address this need by breaking a CRM design into a list of design parameters, each of which is to be calibrated in a prescribed manner.

Despite my pragmatic approach, I hope this book is not only a cookbook. I intend to provide a full coverage of the CRM. This book includes a comprehensive review of the CRM (Chapter 3) and elaborate properties of the CRM (Chapters 5 and 6). While this book is based on my previous publications on the CRM, I have introduced new material so as to present the CRM under a unified framework (Chapter 4). These chapters serve as the theoretical foundation of the calibration techniques presented in the later chapters. I also reflect on what *not* to do with the CRM (Chapter 12) and when *not* to use the CRM (Chapter 13). From a practical viewpoint, these *not-to* chapters are as important as, if not more important than, the *how-to* chapters, because they avoid abuses and pitfalls in applying the CRM. I believe that using the CRM in a wrong way or in the wrong trial is no better, or arguably worse, than falling back to the 3+3 algorithm. The time-to-event aspect of the toxicity endpoint has been a

recurring concern in my previous CRM trials, and so is included as an extension of the CRM (Chapter 11). All in all, while this is not intended to be a cookbook, the inclusion of materials is based on their practical relevance.

This book does not cover dose finding in all possible clinical settings. In fact, it has a singular focus on the simplest and the most common phase I trial setting, where the study endpoint is defined as a binary outcome and the subjects are assumed to come from a homogeneous population. I make no mention of the concerns with multiple toxicity and the gradation of severe toxicities. The topic of individualized dosing is omitted. While some basic ideas of dose finding using both efficacy and toxicity are outlined in Chapter 13, the discussion is brief and does not do full justice to this fast-growing area. All these are important topics in which I am intellectually interested. Their omission, however, is mainly due to my limited *practical* experience in dealing with these “nonstandard” situations in real dose finding studies; dealing with these issues simply from a *methodological* and *theoretical* viewpoint does not fit my intent of writing a practical book (although I think such a book is interesting in its own right and hope someone more qualified than I will deliver it). I do have a word or two to add from a methodological and theoretical viewpoint here, if not already alluded to in the book’s final chapter (Section 14.4, to be precise). First, a complete theoretical framework is crucial for these nonstandard methods to be successfully translated into actual practice. In this book, I try to explicate possible pathological behaviors (e.g., incoherence and rigidity) of some CRM modifications for the simplest setting; it is reasonable to infer that these pathologies will multiply for methods more complex than the CRM for the more complicated clinical settings. Solid theoretical investigation will help us navigate the potential pitfalls. I also hope the theoretical framework developed in this book for the simplest case will prove useful when extended to the complicated settings. Second, and more specifically, I think stochastic approximation offers partial solutions (albeit mostly theoretical) to many of these nonstandard dose finding settings. This is why I close this book with a chapter where I try to connect and compare the CRM with the rich stochastic approximation literature.

The last points I just made give a hint about my methodological and theoretical interests. I hope that this book will in some way simulate research in the CRM and general dose finding methods, despite its practical nature. As I try to present the CRM and the dose finding criteria at a rigorous level, and to cover the CRM literature as comprehensively as possible, I also hope this book can serve as an introduction for those interested in doing research in this area. I taught a course on sequential experimentation at Columbia University from an early unpublished version of this book. This final manuscript is, in turn, adapted from the course notes, and is suitable for use in a course on sequential experimentation or clinical trials.

There are several statistics books on dose finding. The two most popular ones are the edited volumes by Chevret [26] and by Ting [105]. Both give surveys of dose finding methods and are good introductions to the dose finding literature. By comparison, this book is a single-authored work on a specific dose finding method, which I think is necessary if we are to get down to the nuts and bolts of the method.

By writing a book on the CRM, I do not imply that it is the best method out there.

In fact, for the dose finding objective considered here, it is unlikely that there is one method that is best or optimal in a *uniform* sense. While some methods may work best under certain scenarios according to some criterion, the others are optimal under a different criterion. There have been numerous proposals in the last two decades. These proposals can be good alternatives against the 3+3 algorithm as long as they are calibrated properly. And, the CRM is one of these methods. Furthermore, the CRM has been worked out and discussed in the statistical and medical literature so extensively that I believe we are getting close to translating this method into practice. This book hopefully will be a catalyst in this translational process.

I owe a debt of gratitude to Tom Cook, Bin Cheng, and an anonymous reviewer who have been generous with their time and given detailed comments on earlier versions of the book. I am grateful for Jimmy Duong for his help to maintain the R package ‘dfcrm’ (a companion software with this book). I would also like to thank Rick Chappell who introduced me to the CRM and clinical trials when I was a student at University of Wisconsin–Madison. This book would not be possible without his mentoring. Finally, my most heartfelt thanks go to my wife, Amy, for her support and enthusiasm during this writing process.

*New York
October 2010*

Contents

I	Fundamentals	1
1	Introduction	3
2	Dose Finding in Clinical Trials	7
2.1	The Maximum Tolerated Dose	7
2.2	An Overview of Methodology	10
2.3	Bibliographic Notes	15
2.4	Exercises and Further Results	16
3	The Continual Reassessment Method	17
3.1	Introduction	17
3.2	One-Stage Bayesian CRM	17
3.2.1	General Setting and Notation	17
3.2.2	Dose–Toxicity Model	17
3.2.3	Dose Labels	18
3.2.4	Model-Based MTD	20
3.2.5	Normal Prior on β	21
3.2.6	Implementation in R	21
3.3	Two-Stage CRM	22
3.3.1	Initial Design	22
3.3.2	Maximum Likelihood CRM	23
3.4	Simulating CRM Trials	25
3.4.1	Numerical Illustrations	25
3.4.2	Methods of Simulation	25
3.5	Practical Modifications	27
3.5.1	Dose Escalation Restrictions	27
3.5.2	Group Accrual	28
3.5.3	Stopping and Extension Criteria	30
3.6	Bibliographic Notes	31
3.7	Exercises and Further Results	31

4	One-Parameter Dose–Toxicity Models	33
4.1	Introduction	33
4.2	ψ -Equivalent Models	33
4.3	Model Assumptions [†]	36
4.4	Proof of Theorem 4.1 [†]	40
4.5	Exercises and Further Results	40
5	Theoretical Properties	41
5.1	Introduction	41
5.2	Coherence	41
5.2.1	Motivation and Definitions	41
5.2.2	Coherence Conditions of the CRM	42
5.2.3	Compatibility	43
5.2.4	Extensions	45
5.3	Large-Sample Properties	46
5.3.1	Consistency and Indifference Interval	46
5.3.2	Consistency Conditions of the CRM	48
5.3.2.1	Home Sets	48
5.3.2.2	Least False Parameters	48
5.3.2.3	Main Result	49
5.3.2.4	A Relaxed Condition	49
5.3.3	Model Sensitivity of the CRM	51
5.3.4	Computing Model Sensitivity in R	53
5.4	Proofs [†]	54
5.4.1	Coherence of One-Stage CRM	54
5.4.2	Consistency of the CRM	55
5.5	Exercises and Further Results	56
6	Empirical Properties	57
6.1	Introduction	57
6.2	Operating Characteristics	57
6.2.1	Accuracy Index	57
6.2.2	Overdose Number	59
6.2.3	Average Toxicity Number	59
6.3	A Nonparametric Optimal Benchmark	60
6.4	Exercises and Further Results	62
II	Design Calibration	63
7	Specifications of a CRM Design	65
7.1	Introduction	65
7.2	Specifying the Clinical Parameters	66
7.2.1	Target Rate θ	66
7.2.2	Number of Test Doses K	66
7.2.3	Sample Size N	66

7.2.4	Prior MTD v_0 and Starting Dose x_1	67
7.3	A Roadmap for Choosing the Statistical Component	68
7.4	The Trial-and-Error Approach: Two Case Studies	69
7.4.1	The Bortezomib Trial	69
7.4.2	NeuSTART	71
7.4.3	The Case for an Automated Process	73
8	Initial Guesses of Toxicity Probabilities	75
8.1	Introduction	75
8.2	Half-width (δ) of Indifferent Interval	75
8.3	Calibration of δ	77
8.3.1	Effects of δ on the Accuracy Index	77
8.3.2	The Calibration Approach	78
8.3.3	Optimal δ for the Logistic Model	79
8.4	Case Study: The Bortezomib Trial	81
8.5	Exercises and Further Results	87
9	Least Informative Normal Prior	89
9.1	Introduction	89
9.2	Least Informative Prior	89
9.2.1	Definitions	89
9.2.2	Rules of Thumb	91
9.3	Calibration of σ_β	93
9.3.1	Calibration Criteria	93
9.3.2	An Application to the Choice of v_0	93
9.3.3	Optimality Near σ_β^{LI}	95
9.4	Optimal Least Informative Model	97
9.5	Revisiting the Bortezomib Trial	99
10	Initial Design	103
10.1	Introduction	103
10.2	Ordering of Dose Sequences	103
10.3	Building Reference Initial Designs	106
10.3.1	Coherence-Based Criterion	106
10.3.2	Calibrating Compatible Dose Sequences	107
10.3.3	Reference Initial Designs for the Logistic Model	109
10.4	Practical Issues	109
10.4.1	Sample Size Constraint	109
10.4.2	Dose Insertion [†]	112
10.5	Case Study: NeuSTART	113
10.6	Exercises and Further Results	115

III CRM and Beyond	117
11 The Time-to-Event CRM	119
11.1 Introduction	119
11.2 The Basic Approach	119
11.2.1 A Weighted Likelihood	119
11.2.2 Weight Functions	120
11.2.3 Individual Toxicity Risks	122
11.3 Numerical Illustration	123
11.3.1 The Bortezomib Trial	123
11.3.2 Implementation in R	124
11.4 Enrollment Scheduling	125
11.4.1 Patient Accrual	125
11.4.2 Interim Suspensions	127
11.5 Theoretical Properties [†]	129
11.5.1 Real-Time Formulation	129
11.5.2 Real-Time Coherence	129
11.5.3 Consistency	130
11.6 Two-Stage Design	131
11.6.1 Waiting Window	131
11.6.2 Case Study: The Poly E Trial	132
11.7 Bibliographic Notes	135
11.8 Exercises and Further Results	136
12 CRM with Multiparameter Models	139
12.1 Introduction	139
12.2 Curve-Free Methods	139
12.2.1 The Basic Approach	139
12.2.2 Product-of-Beta Prior Distribution	140
12.2.3 Dirichlet Prior Distribution	143
12.2.4 Isotonic Design	144
12.3 Rigidity	146
12.3.1 Illustrations of the Problem	146
12.3.2 Remedy 1: Increase m	147
12.3.3 Remedy 2: Increase Prior Correlations	147
12.4 Two-Parameter CRM [†]	149
12.4.1 The Basic Approach	149
12.4.2 A Rigid Two-Parameter CRM: Illustration	150
12.4.3 Three-Stage Design	151
12.4.4 Continuous Dosage	153
12.5 Bibliographic Notes	154
12.6 Exercise and Further Results	154

13 When the CRM Fails	155
13.1 Introduction	155
13.2 Trade-Off Perspective of MTD	155
13.2.1 Motivation	155
13.2.2 Maximum Safe Dose and Multiple Testing	156
13.2.3 A Sequential Stepwise Procedure	157
13.2.4 Case Study: The ASCENT Trial	159
13.2.5 Practical Notes	161
13.3 Bivariate Dose Finding	162
14 Stochastic Approximation	167
14.1 Introduction	167
14.2 The Past Literature	167
14.2.1 The Robbins-Monro Procedure	167
14.2.2 Maximum Likelihood Recursion	168
14.2.3 Implications on the CRM	169
14.3 The Present Relevance	170
14.3.1 Practical Considerations	170
14.3.2 Dichotomized Data	171
14.3.3 Virtual Observations	174
14.3.4 Quasi-Likelihood Recursion	175
14.4 The Future Challenge	176
14.5 Assumptions on $M(x)$ and $Y(x)^\dagger$	177
14.6 Exercises and Further Results	178
References	179
Index	187

Part I

Fundamentals

Introduction

The clinical development of a new drug or a new treatment proceeds through three phases of testing in human subjects. Phase I trials are small studies that evaluate safety and identify a safe dose range of the treatment. Once a dose range is chosen, its therapeutic efficacy will be examined in a phase II trial. Regimens that are shown promising in phase II trials will be moved to multi-institutional phase III clinical trials for randomized comparison to standard treatments. The ultimate goal of this entire process is to translate promising discoveries in the laboratory into new medical procedures that can be used in the general clinical settings. This division of clinical trials, however, may give an oversimplified picture of the actual drug development process. Often, several phase I-II trial sequels of a drug, possibly with minor variations in the treatment schedule and patient populations, are needed before a phase III trial is warranted. This process is necessarily iterative rather than linear, as the phase I-II-III paradigm appears to suggest. In addition, the taxonomy of trials is not universal across disciplines, and may include finer divisions such as phase IA, IB, IIA, and IIB. The recent trend to combine phases of trials, the so-called combined phase I/II trials and seamless phase II/III trials, renders further refinement of the drug development process.

This having been said, the phase I-II-III paradigm provides a conceptual framework for in-depth study of statistical methods. The subject matter of this book is dose finding using the continual reassessment method (CRM). The CRM [78] is among the first model-based designs for phase I cancer trials in which toxicity is the primary study endpoint. The role of toxicity in early-phase cancer trials had long been a subject for discussion in the medical literature [93, 85]. In particular, for cytotoxic drugs, toxicity serves as evidence that the drug has reached a level that does harm not only to the cancer cells but also to a patient's normal organs. In other words, a therapeutic dose is expected to cause a significant amount of severe but reversible toxicities in the cancer patient population. Therefore, a primary goal of phase I cancer trials is to identify the so-called maximum tolerated dose (MTD). For other disorders such as acute stroke and HIV, identifying the MTD is also a primary objective of early-phase safety studies (usually called phase IB trials). In addition, dose finding is important in phase II proof-of-concept trials where the goal is to identify a dose range with demonstrated biological activity. This objective is usually achieved through the estimation of the minimum effective dose (MED) [106, 27]. From a statistical viewpoint,

the MTD in safety studies and the MED in efficacy studies can be formulated in an analogous way. Therefore, this book is relevant to the design of phase I and II dose finding trials. Under the modernized paradigm, the dose finding principles discussed here also address the design issues in the combined phase I/II trials, in which both the safety and the efficacy endpoints are considered as co-primary (cf. Section 13.3).

Another advantage of dividing the drug development process into phases is that by doing so, we can set a clear and manageable benchmark to achieve in a particular study. This entails a clearly defined set of study endpoints and an interpretable study objective. Since clinical trials are conducted in human subjects, each benchmark is to be reached within certain ethical constraints. In particular, in phase I dose finding studies, randomization is not entirely appropriate because it may expose subjects to excessively high and toxic doses without sufficiently testing the lower doses. (Some would also argue randomization exposes subjects to low and inefficacious doses, although this aspect is apparently not as alarming.)

We illustrate these points using a bortezomib dose finding trial [62]. Bortezomib is a proteasome inhibitor with proven activity in lymphoma. In the trial, bortezomib was given in combination with the standard chemotherapy as a first-line treatment for patients with diffuse large B cell or mantle cell non-Hodgkin's lymphoma. Each patient would receive up to six 21-day cycles of the treatment combination. Table 1.1 describes the five dose schedules of bortezomib tested in the trial. The primary safety concerns related to bortezomib were neuropathy, low platelet count, and symptomatic non-neurologic or non-hematologic toxicity. Toxicity was graded according to the National Cancer Institute Common Terminology Criteria for Adverse Events [71], with grade 3 or higher defined as *dose limiting*. Generally, a grade 3 toxicity is severe but can be resolved by symptomatic treatment, whereas a grade 4 toxicity is irreversible; toxic death due to the treatment is invariably defined as grade 5. The primary endpoint of each patient was the indicator of whether any dose-limiting toxicity (DLT) was experienced at any time during the six cycles. The objective of the trial was to determine the MTD, defined as a dose associated with a 25% DLT rate. Table 1.1 gives the number of patients and the number of DLTs per dose in the bortezomib trial. The data show strong evidence that the highest dose is adequately safe: we pool the observations in dose levels 4 and 5 by assuming an increasing dose-toxicity relationship; based on 1 DLT out of 16 patients, we obtain a 95% confidence upper bound of 0.26 for the DLT probability.

Table 1.1 The bortezomib trial [62]: dose schedules of bortezomib, sample size (n), and the number of DLT (z) at each dose

Level	Dose and schedule within cycle	n	z
1	0.7 mg/m ² on day 1 of each cycle	0	0
2	0.7 mg/m ² on days 1 and 8 of each cycle	0	0
3	0.7 mg/m ² on days 1 and 4 of each cycle	4	0
4	1.0 mg/m ² on days 1 and 4 of each cycle	9	1
5	1.3 mg/m ² on days 1 and 4 of each cycle	7	0

While simple analyses are usually adequate to address the primary scientific questions in a phase I study, the summary statistics in Table 1.1 ignore how the data were collected. Figure 1.1 shows the dose assignments of the trial in chronological order. The trial started at level 3, a dose schedule that the investigators believed to be safe to treat patients. Escalation to the next higher dose occurred after four patients had been followed for several weeks without signs of toxicity, and another escalation took place after three following patients. Shortly after the eighth patient entered the trial at the highest dose, patient 7 at dose level 4 experienced a DLT, thus leading to a deescalation for the ninth patient. Subsequent patients were enrolled in a staggered fashion, allowing months to pass before reescalating to the highest level. A central feature of this dose assignment scheme is its outcome adaptiveness. Specifically, in the bortezomib trial, the dose assignments were made in accordance with the time-to-event continual reassessment method (TITE-CRM), an extension of the CRM to be discussed in Chapter 11. For ethical reasons, most dose finding trials are conducted in an outcome-adaptive manner, so that the dose assignment of the current patient depends on those of the previous patients. As such, the focus of this book is the *design* (as opposed to analysis) of a dose finding study using the CRM and its variants.




Figure 1.1 Dose assignments in the bortezomib trial. Each number indicates a patient: An unmarked number represents the patient's entry time; a number marked with "O" indicates the time when a DLT occurs, and "D" indicates the time of dropout. Vertical positions of some numbers are jittered for clarification.

This book is organized into three parts. Part I (Chapters 2–6) contains the background and introductory material of the CRM. Specifically, Chapter 2 provides the clinical background, outlines the problem of dose finding in the context of several real trial examples, and reviews the dose finding literature. Chapter 3 introduces the basic approach of the CRM and presents its major modifications. The method will be developed along with a description of an R package ‘dfcrm’. Chapter 4 presents a unified framework for dose–toxicity models used in the CRM. Chapters 5 and 6, respectively, discuss the theoretical and empirical properties of the CRM. The objective of Part I is for the readers to develop a basic understanding of the CRM and be able to implement the method using a simple R code. Readers familiar with the basic CRM methodology are also encouraged to review the materials, as they are reorganized and presented in a unified framework in this book.

Part II (Chapters 7–10) details the calibration process of the CRM based on the notation and the theory introduced in Part I. Chapter 7 introduces a system of design parameters involved in the CRM, and classifies them into two categories: *clinical parameters* and *model parameters*. The subsequent chapters then present fine-tuning techniques of the model parameters: the initial guesses of the toxicity probabilities (Chapter 8), the prior distribution of the model parameter (Chapter 9), and the initial design of a two-stage CRM (Chapter 10). The objective of Part II is for the readers to develop the ability to design a “good” CRM trial within a reasonable timeline.

Part III (Chapters 11–14) contains a variety of advanced topics related to the CRM. Chapter 11 presents the TITE-CRM to deal with situations in which the toxicity outcome is defined with respect to a nontrivial duration. Chapter 12 gives a critical review of CRM using multiparameter models. Chapter 13 considers situations where the CRM is an inappropriate design, and puts forward some alternatives. Chapter 14 connects the CRM and modern dose finding trials to the large literature of stochastic approximation. The objective of Part III is to stimulate further research in the CRM and general dose finding methodology.

The materials in this book are presented at a level that requires college algebra and some basic calculus concepts. Sections marked with “†” in the table of contents contain technical details that may be skipped without affecting the reading of the other chapters. Exposition in the book will be supplemented by illustrations of the usage of R functions in the ‘dfcrm’ package. While some basic knowledge of R will enhance the reading experience, proficiency in R is not required. Interested readers can find out more information about R from the Comprehensive R Archive Network (CRAN) [83] at

<http://www.r-project.org>.

Dose Finding in Clinical Trials

2.1 The Maximum Tolerated Dose

The primary objective of phase I trials of a new anticancer drug is to assess the toxic side effects of the drug and to recommend a dose for the subsequent phase II trials. This recommended dose is typically the maximum test dose that does not exceed an acceptable level of toxicity, the so-called *maximum tolerated dose* (MTD). Traditional chemotherapy takes the cytotoxic therapeutic mechanism under which toxicity may be viewed as a surrogate for anti-tumor activity. Toxicity, therefore, is in a sense a desirable endpoint, so the trial objective is to find a dose that is associated with a given level of toxicity probability. Also in this sense, this MTD is presumed optimal in the absence of information about efficacy and clinical response.

Definition 2.1 (MTD—surrogacy perspective). *In a trial with K test doses, let p_k denote the toxicity probability associated with dose level k for $k = 1, \dots, K$. The MTD from a surrogate-for-efficacy perspective, denoted by v , is defined as the dose level with toxicity probability closest to a prespecified target probability θ , that is, $v \equiv \arg \min_k |p_k - \theta|$.*

Since the late 1980s, most dose finding designs have been proposed to specifically address dose finding of cytotoxic drugs in patients with solid tumors and other forms of malignancies. The bortezomib trial in lymphoma patients introduced in Chapter 1 is one such example. As a result, the phase I method literature has focused on the surrogacy definition of MTD. See Exercise 2.2 for an alternative definition of MTD from a surrogacy perspective.

For noncytotoxic target anticancer agents and treatments for other diseases such as acute ischemic stroke, toxicity does not play a therapeutic role, but safety remains the primary concern in the early drug development phase. For these agents, it is still useful to define an upper safety limit of the dose range for further clinical research. Furthermore, under the assumption that efficacy of a drug increases with dose, there is merit in pushing the dose as high as safety will permit. In view of this trade-off between safety and efficacy, one may seek to maximize the dose administered to patients subject to toxicity constraints. This may lead to a slightly different dose recommendation for the next study phase:

Definition 2.2 (MTD—trade-off perspective). *In a trial with K test doses, let p_k*

denote the toxicity probability associated with dose level k for $k = 1, \dots, K$. The MTD from a trade-off-for-efficacy perspective, denoted by γ , is defined as the largest dose level with toxicity probability no greater than a prespecified threshold θ , that is, $\gamma \equiv \max\{k : p_k \leq \theta\}$.

Both the surrogacy and the trade-off perspectives define the MTD with respect to a target toxicity rate θ , and as such formulate dose finding as a percentile estimation problem. However, from a statistical viewpoint, it is generally easier to estimate v than γ (if v is well defined) by making use of monotonicity of the dose-toxicity curve. Therefore, it is in some cases pragmatic to take v as the operative objective of a trial, even though the toxicity endpoint is generally not a surrogate for efficacy for noncytotoxic drugs. In practice, it is important to discern the appropriate objective for a given clinical setting, and the degree of tolerance in terms of the target θ . Apparently, these decisions need to be made on a trial-by-trial basis.

Example 2.1 (acute ischemic stroke). A number of statins, when administered early after stroke in animal models, have demonstrated neuroprotective effects in a dose-dependent manner, with the greatest effects at the highest doses. The NeuSTART (Neuroprotection with Statin Therapy for Acute Recovery Trial) drug development program aimed to translate preclinical research and test the role of high-dose statins in stroke patients. In a phase IB dose finding study under the NeuSTART program [34], high-dose lovastatin was given to patients for 3 days after stroke followed by a standard dose for 27 days. The primary safety concerns for giving high-dose lovastatin included elevated liver enzyme; a toxicity was said to occur if the peak enzyme levels at any posttreatment time points exceeded a prespecified threshold. There were five test doses in the NeuSTART and the trial objective was to identify a dose with toxicity rate closest to 10%, that is, $\theta = 0.10$.

Example 2.2 (cancer prevention). Polyphenon E (Poly E) is a tea catechin extract that is thought to block tumor promotion by inhibiting cell proliferation and inducing cell cycle arrest and apoptosis. The multiple mechanisms of Poly E make it a good candidate agent for chemoprevention. On the other hand, the agent has been shown toxic and causing mortality in female beagle dogs after an overnight fast. Toxicity generally involved the gastrointestinal (GI) system, producing vomiting and damage to the lining of the GI tract, with hemorrhage and necrosis apparent at autopsy. A Poly E trial was conducted in women with a history of hormone receptor-negative breast cancer. Three different doses of Poly E were administered to subjects over 6 months. The study objective was to find the MTD defined as a dose that causes 25% DLT during the six-month period, where a DLT here meant any grade 2 or higher toxicity that would persist for at least one week or requires stopping the study drug. As a secondary objective, this study included biologic correlates such as tissue-based biomarkers and mammography. Therefore, some subjects were randomized to receive a placebo, to which the identified MTD would be compared.

Example 2.3 (early neuro-rehabilitation). Early rehabilitation was conjectured to

enhance recovery in stroke patients. On the other hand, premature physical therapy might cause neurologic worsening, cardiac complication, or even death in the short term. Based on historical data, the adverse event rate is estimated to be 25% in the untreated patients during the first 4 days after stroke. The ASCENT (Acute Stroke Collaboration with Early Neurorehabilitation Therapy) trial was a clinical trial of early physical therapy (PT) in stroke patients, its objective being to identify the largest PT dose that could be instituted without causing adverse events in excess of the 25% background rate. Table 2.1 displays the six PT regimens in ASCENT. In this example, a dose is composed of the timing and duration of therapy.

Table 2.1 Physical therapy (PT) regimen in ASCENT

Regimen	Minutes of PT on			Total PT dose (minutes)
	Day 2	Day 3	Day 4	
1	0	0	30	30
2	0	30	30	60
3	30	30	30	90
4	30	30	60	120
5	30	60	60	150
6	60	60	60	180

In the NeuSTART in Example 2.1, liver enzyme elevation is expected in about 3% of the stroke population, which consists mainly of the elderly. However, since the toxicity endpoint is quite mild and reversible upon drug withdrawal, a toxicity rate higher than the background would be tolerated for the potential benefit in efficacy, and it would also seem reasonable to accept a dose above the target as long as the toxicity probability at this dose was close to θ . Therefore, Definition 2.1, that is, v was chosen as the operative MTD objective for the trial, even though in this case, the surrogacy perspective is far from reality: Liver enzyme elevation was by no means a surrogate for any clinical benefits in stroke patients. Similar argument may be made for using v as the objective in the Poly E trial; cf. Exercise 2.1.

On the other hand, in the ASCENT trial in Example 2.3, functional recovery due to early physical therapy may not warrant an elevated adverse event rate, and the target rate θ should be set at the background 25%. In this situation, Definition 2.2 appears more appropriate than Definition 2.1. For one thing, since we expect that the dose–toxicity probability is about 25% for doses below and up to the MTD, the objective v is not uniquely defined. Figure 2.1c displays the plausible shape of the dose–toxicity curve in the ASCENT trial, under which γ is still well defined. Generally, the use of v requires a strictly increasing dose–toxicity relationship around the MTD, whereas γ always exists regardless of the shape of the dose–toxicity curve.

Figure 2.1 also shows plausible dose–toxicity curves for the bortezomib trial and the Poly E trial, where the MTD is defined with respect to a 25% target toxicity rate. In the bortezomib trial, since untreated lymphoma patients will be at no risk of getting a grade 3 neuropathy, the probability (y -intercept) in Figure 2.1a approaches zero as




Figure 2.1 Plausible shapes of dose–toxicity curve on a conceptual continuous dose range in three studies and the corresponding MTD, defined as a dose associated with a 25% target toxicity rate. Definition 2.1 is used to define MTD in (a) and (b), and Definition 2.2 in (c).

dose decreases. In contrast, because low-grade GI toxicities are not uncommon in subjects with a history of cancer in the Poly E trial, there is a nonzero intercept in Figure 2.1b. However, it is generally believed that the dose–toxicity curve will be strictly increasing around the 25th percentile in both trials, and thus v is well defined.

It is easy to verify that $v = \gamma$ when the dose–toxicity curve is a continuous and strictly increasing function of dose. Thus, if a continuum of test doses is available for the trial, it will be a practical choice to use v as the operative objective, because estimating v is generally easier than estimating γ . On the other hand, most trials in practice allow only a discrete number of test doses; in which case, Definition 2.1 may yield a slightly more aggressive recommendation than Definition 2.2, because $v \geq \gamma$. Therefore, the choice between v and γ depends on whether it is acceptable to be (slightly) more aggressive than the target θ , given the clinical factors such as the nature of treatment, the severity of the disease, and the seriousness of the study endpoint.

As alluded to earlier, the CRM is motivated by applications in cancer trials with a surrogacy view. Thus, the book will naturally focus on the estimation of v using the CRM. However, Chapter 13 will explore situations in which the CRM is not applicable, and introduce an alternative approach that uses γ as trial objective. Also, as the MED in efficacy trials may be defined analogously to v , the CRM may be applicable to dose finding in phase II efficacy trials. However, this book will focus on the MTD finding by the CRM according to the method’s originally intended use; the design strategy for the MED can be derived by analogy.

2.2 An Overview of Methodology

This section gives a brief overview of the development of the dose finding literature since the late 1980s so as to put the CRM in a historical light. Limited by the scope of this book, the review will be cursory. Interested readers can find additional topics in the edited volume by Chevret [26] and the article by Le Tourneau et al. [59].

The 3+3 algorithm. Traditionally, a 3+3 algorithm is used to dictate dose escalation and to approach the eventual recommended dose. The method starts the trial at a low dose (e.g., one-tenth of LD10 in mice) and escalates after every three to six patients per dose; the recommended dose is defined as the largest dose with fewer than two patients experiencing a predefined DLT during the first course of treatment. Table 2.2 describes the dose escalation rules of this algorithm. In practice, there may be slight variations from institution to institution. For example, in order to obtain preliminary information about efficacy of the drug, it is common to treat additional patients (usually 6 to 12) at the identified MTD.

Table 2.2 Escalation rules at a given dose with cumulative sample size n and total number z of DLT in accordance with the 3+3 algorithm

n	z	Action
3	0	Escalate to the next higher dose
3	1	Treat three additional patients at the current dose
6	1	Escalate to the next higher dose
3 or 6	≥ 2	Stop escalation and terminal trial ^a

^a*The MTD is estimated by the dose immediately below the terminating dose.*

The 3+3 algorithm has historically been the most widely used phase I trial design. Its main advantage is simplicity. Since the dose escalation rules can be tabulated before a trial starts, the clinical investigators can make dose decisions during a trial without help from a statistician.

However, the algorithm has two major shortcomings. First, due to a low starting dose and the conservative escalation scheme, the 3+3 algorithm tends to treat many patients at low and inefficacious doses. Since phase I cancer trials typically enroll patients as subjects, as opposed to healthy volunteers, there is an intent to treat the subjects at a therapeutic dose that is likely higher than the lowest test dose. As such, the algorithm is discordant with the therapeutic intent of these trials [84, 55]. Second, the 3+3 algorithm has no statistical justification. There is no intrinsic property in the method to stop escalation at any given percentile, and thus the distribution of the recommended MTD depends arbitrarily on the underlying dose–toxicity curve and the number of test doses [97]. This deficiency is due to the fact that there is no correspondence between the method and any quantitative definition of the MTD, as the 3+3 algorithm does not involve an explicit choice of the target θ . As poor dose selection in the early phase will likely be carried over to its subsequent developmental phases, the use of the 3+3 algorithm will have lingering financial implications and adverse scientific consequences. And thus, the simplicity of the method does not justify its widespread application.

Stochastic approximation. Although discussions on phase I designs for cancer trials can be traced back to the 1960s [93], a formal statistical formulation of the MTD appeared at a much later time. Among the earliest discussions was Anbar [2], who in 1984 proposed the use of stochastic approximation [86] in phase I trials. The

procedure assigns dose sequentially by the recursion

$$x_{i+1} = x_i - \frac{1}{ib}(Y_i - \theta) \quad (2.1)$$

for some prespecified constant $b > 0$, where x_i is the dose assigned to patient i , and Y_i is the toxicity indicator. The stochastic approximation is a nonparametric method in that it does not assume any parametric structure on the dose–toxicity relationship. Let $\pi(x) = \Pr(Y_i = 1 \mid x)$ denote the probability of toxicity at dose x . Under very mild assumptions of $\pi(x)$, the dose sequence $\{x_i\}$ generated by (2.1) will converge with probability 1 to a dose x^* such that $\pi(x^*) = \theta$.

However, the use of recursion (2.1) implicitly assumes a continuum of doses is available for testing in the trial. This is not always feasible in practice. In many situations such as the bortezomib trial, there may be no natural scale of dosage; rather, “dose” is composed of drug dosage and treatment schedule. There are also other difficulties from a statistical viewpoint. First, the stochastic approximation has been shown to be inferior to model-based methods such as the maximum likelihood recursion for binary data [112, 50]. Second, the choice of the constant b in (2.1) has a large impact on the performance of the procedure. As a consequence, the stochastic approximation has seldom been used in dose finding trials. We will return to this in Chapter 14.

Up-and-down designs. Subsequent to Anbar’s work [2], Storer [98] considered the up-and-down schemes originally described by Wetherill [110]. An example of the up-and-down design, which Storer called design D, is to enroll patients in groups of three; then, escalate dose for the next group if there is no toxicity in the most recent group, deescalate if there is more than one toxicity in the group, and stay at the same level if one of three patients has toxicity. By Markov chain representation, design D can be shown to sample around a dose that causes toxicity with a probability $\theta = 0.33$. The group size and the decision rules in the up-and-down schemes can apparently be modified to accommodate other target θ ; for example, for a target $\theta = 0.20$, one will intuitively enroll patients in groups of five. Also, combinations of schemes can be applied in stages. Storer, in particular, suggested using a group size of one initially in the trial and switching to design D upon the first observed toxicity. The idea here is to move the trial quickly through the low doses so as to avoid treating many patients at low and inefficacious doses.

Durham et al. [32] proposed a randomized version of the up-and-down rule for any target at or below the 50th percentile, that is, $\theta \leq 0.50$. The design deescalates the dose for the next patient if the current patient has toxicity, and escalates according to a biased coin with probability $\theta/(1 - \theta)$ if there is none. The method is thus called a *biased coin design*.

In a trial that uses design D or the biased coin design for dose escalation, when the enrollment is complete, we may naturally estimate the dose–toxicity curve using logistic regression [98] or isotonic regression [99], and estimate the MTD by the 100θ th percentile of the fitted curve. The consistency conditions for the estimation of the MTD hold, because the design points $\{x_i\}$ “spread out” under this type of random walk sampling plans. That is, technically, it can be shown that the number of patients

treated at each dose will grow indefinitely as sample size increases. It can also be shown by the properties of random walk that the asymptotic distribution of dose allocation has a mode near the target MTD. However, from a *design* viewpoint, the “memoryless” property of random walk may cause ethical difficulties: Since these up-and-down rules make dose decisions based only on the most recent patient or group of patients at the current dose, previously accrued data are ignored and a trial will likely reescalate to a dose that appears toxic.

Model-based designs. In brief, a model-based design makes dose decisions based on a dose–toxicity model, which is being updated repeatedly throughout a trial as data are accrued. The continual reassessment method (CRM), proposed by O’Quigley et al. [78] in 1990, is the first model-based design in the modern dose finding literature. Several model-based designs proposed since 1990 share a similar notion with the CRM. One example is the escalation with overdose control (EWOC) design [4], which takes the continual reassessment notion but estimates the MTD with respect to an asymmetric loss function that places heavier penalties on overdosing than underdosing; see (2.2) below. A list of model-based methods is given in Section 2.3.

Most model-based methods take the myopic approach by which dose assignment is optimized with respect to the next immediate patient without regard to the future patients. For example, the EWOC at each step minimizes the Bayes risk with respect to the loss function:

$$x_{i+1} = \arg \min_x E_i \left\{ \alpha(v - x)^+ + (1 - \alpha)(x - v)^+ \right\} \quad (2.2)$$

where $x^+ = \max(x, 0)$ and $E_i(\cdot)$ denotes expectation computed with respect to the posterior distribution of the MTD v , given the first i observations. As the EWOC is intended to control overdose, the loss function (2.2) should be specified with a value of α , that is, $\alpha < 0.50$. More recently, Bartroff and Lai [6] take a stochastic optimization approach that minimizes the global risk and propose to choose the doses $\{x_1, x_2, \dots, x_N\}$ sequentially so as to minimize

$$E \left\{ \sum_{i=1}^N \alpha(v - x_i)^+ + (1 - \alpha)(x_i - v)^+ \right\},$$

where the expectation is taken with respect to the joint distribution of x_i s and v . Such sequential optimization is implemented by backward induction and requires dynamic programming which can be computationally intensive. However, this approach presents a new direction for model-based design and warrants further research.

The model-based approach facilitates borrowing strength from information across doses through the parametric assumptions on the dose–toxicity curve. This is especially important in early-phase dose finding trials where sample sizes are small and informational content is low [41]. In addition, since these methods allow starting a trial at a dose higher than the lowest level, the in-trial allocation tends to concentrate around the target dose.

Many model-based designs take a Bayesian approach. They update sequentially

the uncertainty about the dose–toxicity curve with respect to the posterior distribution given the interim data. This approach is by nature automated, so far as the posterior computations can be efficiently programmed and reproduced. An advantage of such automation is that a carefully *calibrated* dose–toxicity model can handle *unplanned* contingencies in a manner that is coherent with the trial objective. In contrast, we may imagine the predicament arising with the 3+3 rule in a trial where there are two toxic outcomes among 7 patients at a dose; this contingency can be caused by overaccrual due to administrative delays, and is not uncommon in practice.

Several practical difficulties may hinder the use of a model-based design. First, there is skepticism among the clinicians because of the “blackbox” approach of these designs. Second, there is the perception that the success of the method is sensitive to the choice of the dose–toxicity model. Third, as a model-based design requires specialized computations, the clinical team will need to interact regularly with the study statistician for interim dose decisions. Such interaction may be perceived as adding unnecessary burdens on both parties, in light of the fact that the standard 3+3 algorithm requires minimal statistical inputs.

This book attempts to address the second difficulty with a focus on the CRM. By theoretical and empirical arguments, we will see that the method’s performance does not depend on correctness of the model specification. This book also partially addresses the third difficulty by illustrating how a CRM design can be calibrated and implemented in practice. The ultimate goal is to alleviate the statistician’s burden during the planning stage and the conduct of the trial. This endeavor is facilitated by availability of software. This book focuses on the R package ‘dfcrm’ for the CRM and its major variants. Software for some model-based designs is also available to public access; see Table 2.3.

Table 2.3 Some software links for model-based approaches

Description	Section	Reference
Escalation with overdose control (EWOC) ^a	2.2	[4]
Late-onset toxicity monitor using predicted risks ^b	11.4.2	[8]
Modified CRM on a continuous dosage range ^c	12.4.4	[81]
Phase I/II dose finding based on efficacy and toxicity ^b	13.3	[102]

^a<http://www.sph.emory.edu/BRI-WCI/ewoc.html>

^b<http://biostatistics.mdanderson.org/softwaredownload>

^c<http://www.cancerbiostat.unc.jhmi.edu/software.cfm>

Algorithm-based designs. Because of the above-mentioned difficulties with the model-based designs, there seems to be a renewed interest in the algorithm-based designs since the late 1990s. Generally, an algorithm-based design prescribes a set of escalation rules for any given dose without regard to the outcomes at other doses. As a result of the independence among observations across doses, the rules can be tabulated and made accessible to the clinical investigators before a trial starts.

The 3+3 algorithm is the most prominent example of algorithm-based designs. Efforts have recently been made to extend this traditional method so as to obtain

well-defined statistical properties. In particular, Cheung [18] formulates dose finding as a multiple testing problem and introduces a class of stepwise test procedures that operate in a manner similar to the 3+3 algorithm. This approach has practical appeal because clinicians are familiar with the 3+3 algorithm. Chapter 13 gives further details of the stepwise procedures.

Ji et al. [49] propose a class of up-and-down designs that make dose decisions based on the posterior toxicity probability intervals. Specifically, the parameter space of the toxicity probability p_k at dose k is partitioned into three sets: for some prespecified constants $K_1, K_2 > 0$,

$$\begin{aligned}\Theta_{k,E} &= \{p_k - \theta < -K_1\sigma_k\} \\ \Theta_{k,S} &= \{-K_1\sigma_k \leq p_k - \theta \leq K_2\sigma_k\} \\ \Theta_{k,D} &= \{p_k - \theta > K_2\sigma_k\}\end{aligned}$$

so that $\Theta_{k,E} \cup \Theta_{k,S} \cup \Theta_{k,D} = [0, 1]$, where σ_k is the posterior standard deviation of p_k . Escalation from a current dose, say dose k , will take place if the posterior probability of $\Theta_{k,E}$ is largest among the three sets. Similarly, deescalation occurs if $\Theta_{k,D}$ is the most probable event according to the posterior distribution of p_k . Otherwise, the next dose will remain at level k . An important difference between this design and the random walk up-and-down is that the posterior interval uses all observations accrued to a dose to make a dose decision and avoids the memoryless problem of a random walk design.

2.3 Bibliographic Notes

The practice and design of phase I trials are discussed in the medical community by Schneiderman [93], Carbone et al. [13], Ratain et al. [85, 84], and Kurzrock and Benjamin [55]. Korn [53] discusses the relevance of MTD in noncytotoxic targeted cancer treatments. Discussion of dose finding strategies in the other disease areas is relatively sporadic. Fisher et al. [38] present some phase I and II trial designs in the context of acute stroke MRI trials, and make a case against the use of the 3+3 algorithm in the phase I safety studies. Cheung et al. [23] make analogous arguments for early-phase trials in patients with amyotrophic lateral sclerosis. The dose finding design of NeuSTART is reported in Elkind et al. [34] as a case study.

Robbins and Monro [86] introduce the first stochastic approximation method, which has been studied extensively and has motivated a large number of subsequent modifications. (See for example Sacks [92], Venter [107], Lai and Robbins [57, 58], and Wu [112, 113].) In the more recent literature, Lai [56] gives a thorough review of the advances of stochastic approximation. Cheung [19] draws a specific tie between this area and modern dose finding methods.

Lin and Shih [64] study the operating characteristics of a class of $A + B$ designs that include the 3+3 algorithm as a special case. The theoretical properties of the biased coin design are established by Durham and colleagues [29, 30, 31].

Several model-based designs have been proposed since the 1990s. These include the Bayesian decision-theoretic design [111], the logistic dose-ranging strategy [70],

and the Bayesian c -optimal design [44]. The CRM has generated a large literature and will be reviewed in the next chapter. Chapter 12 will comment on two CRM-like dose finding designs: the curve-free method [40] and the isotonic design [63]. For the EWOC, Zacks et al. [117] prove that the method is Bayesian-feasible, Bayesian-optimal, and consistent under the assumption that the specified dose–toxicity model is correct. These Bayesian criteria are introduced in the previous work by Eichhorn and Zacks [33].

2.4 Exercises and Further Results

Exercise 2.1. Discuss the MTD objective (surrogacy versus trade-off perspectives) for cancer prevention in the context of the Poly E trial.

Exercise 2.2. Definition 2.1 formulates dose finding as estimating a percentile on a dose–toxicity curve. Another possible alternative to define the MTD according to the surrogacy perspective is

$$v' = \arg \min_k |\pi^{-1}(p_k) - \pi^{-1}(\theta)|.$$

Show that $|v - v'| \leq 1$, that is, the two definitions can differ by at most one dose level. Discuss why v' may not be applicable for the bortezomib trial.

Exercise 2.3. By computer simulations, generate the outcomes of a trial using the 3+3 algorithm with $K = 5$ doses and true toxicity probabilities $p_1 = 0.02$, $p_2 = 0.04$, $p_3 = 0.10$, $p_4 = 0.25$, and $p_5 = 0.50$. Observe the recommended MTD. Repeat the simulations 1000 times, and record the distribution of the recommended MTD in the 1000 simulated trials.

The Continual Reassessment Method

3.1 Introduction

In this book, we consider two types of dose finding strategies using the continual reassessment method (CRM): a one-stage design that necessitates the use of Bayesian CRM (Section 3.2) and a two-stage design (Section 3.3). This chapter outlines the basic CRM approach and introduces the necessary notation for further development in the later chapters. Section 3.4 presents some simulation outputs of the CRM to illustrate how the method may operate in practice. Section 3.5 reviews some common modifications of the CRM. Section 3.6 gives key references in the CRM literature.

3.2 One-Stage Bayesian CRM

3.2.1 General Setting and Notation

Consider a trial with K test doses with numerical labels d_1, \dots, d_K . In a dose finding trial, patients are enrolled in small groups of size $m \geq 1$. Let $x_i \in \{d_1, \dots, d_K\}$ denote the dose assigned to the i th group of patients, so that patients in the same group receive the same dose. In what follows, we first consider a *fully sequential* enrollment plan (i.e., $m = 1$), where we observe a binary toxicity outcome Y_i from the i th patient, and postulate Y_i as a Bernoulli variable with toxicity probability $\pi(x_i)$, where $\pi(x)$ is a monotone increasing function in x . We will consider group accrual enrollment, that is, $m > 1$, in Section 3.5.2.

In accordance with Definition 2.1, the trial objective is to identify the dose level $v \in \{1, \dots, K\}$ that is associated with a toxicity probability θ .

3.2.2 Dose–Toxicity Model

The CRM assumes a dose–toxicity model $F(x, \beta)$; that is, the true dose–toxicity curve $\pi(x)$ is postulated to be $F(x, \beta_0)$ for some true parameter value β_0 . Generally, the CRM does not require F to be a correct model for π , and β_0 may then be viewed as a “least false” value. Briefly, we require $F(x, \beta)$ to be strictly increasing in the dose x , in addition to some regularity conditions. Details of the assumptions will be given in Chapters 4 and 5. The two most commonly used models in the CRM literature are the empiric function

$$F(x, \beta) = x^\beta \text{ for } 0 < x < 1 \tag{3.1}$$

and a one-parameter logistic function

$$F(x, \beta) = \frac{\exp(a_0 + \beta x)}{1 + \exp(a_0 + \beta x)} \text{ for } -\infty < x < \infty$$

where the intercept a_0 is a fixed constant. Another common dose–toxicity model is the hyperbolic tangent function [78, 67]

$$F(x, \beta) = \left(\frac{\tanh x + 1}{2} \right)^\beta \text{ for } -\infty < x < \infty.$$

To ensure an increasing dose–toxicity relationship, the parameter β in these models is restricted to taking on positive values. The positivity constraint could present some difficulty in estimation, especially when the sample size is small. Hence, it is useful to consider the following parameterization:

$$\text{empiric: } F(x, \beta) = x^{\exp(\beta)} \text{ for } 0 < x < 1 \quad (3.2)$$

$$\text{logistic: } F(x, \beta) = \frac{\exp\{a_0 + \exp(\beta)x\}}{1 + \exp\{a_0 + \exp(\beta)x\}} \text{ for } -\infty < x < \infty \quad (3.3)$$

$$\text{hyperbolic tangent: } F(x, \beta) = \left(\frac{\tanh x + 1}{2} \right)^{\exp(\beta)} \text{ for } -\infty < x < \infty \quad (3.4)$$

under which the parameter β is free to take on any real values while $F(x, \beta)$ is strictly increasing.

The original formulation of the CRM uses a Bayesian approach by which the model parameter β is assumed random and follows a prior distribution $G(\beta)$. We will focus on the normal prior distribution, that is,

$$\beta \sim N(\hat{\beta}_0, \sigma_\beta^2)$$

where $\hat{\beta}_0$ and σ_β^2 are, respectively, the prior mean and variance.

3.2.3 Dose Labels

An important point about the CRM is that the numerical dose labels d_1, \dots, d_K are *not* the actual doses administered, but rather are defined on a conceptual scale that represents an ordering of the risks of toxicity. Consider the dose schedules used in the bortezomib trial (Table 1.1). The first three levels prescribe bortezomib at a fixed dose 0.7 mg/m^2 with increasing frequency, whereas the next two levels apply the same frequency with increasing bortezomib dose. While it is reasonable to assume that the toxicity risk increases with each level, there is no natural unit for dose (e.g., mg/m^2) in this application. Similarly, in the ASCENT trial (Example 2.3), “dose” is a composite of timing and duration of physical therapy given. In these examples, it is artificial to assume the dose–toxicity curve $\pi(x)$ is well defined on a continuous

dose range. Instead, one will have access only to a discrete set of increasing doses. As the CRM operates on a discrete set of dose levels, a physical interpretation for the dose labels d_1, \dots, d_K is not required, as long as they constitute a *strictly* increasing sequence.

In practice, to ensure monotonicity, the label d_k can be obtained by substituting the initial guess of toxicity probability p_{0k} for dose level k into the dose–toxicity model, that is, solving

$$p_{0k} = F(d_k, \hat{\beta}_0). \quad (3.5)$$

The set of initial guesses $\{p_{0k}\}$ is sometimes called the ‘skeleton’ of the CRM, and is a *strictly* increasing sequence, that is, $p_{01} < p_{02} < \dots < p_{0K}$.

Suppose there is a prior belief that dose level v_0 is the MTD. We may set the initial guess

$$p_{0, v_0} = \theta. \quad (3.6)$$

Consider, for instance, a trial with $K = 5$ dose levels and a target probability $\theta = 0.25$. Suppose we use the logistic function (3.3) with $a_0 = 3$ and prior mean $\hat{\beta}_0 = 0$, and we believe that $v_0 = 3$ is the prior MTD such that $p_{03} = 0.25$. Then we can solve

$$0.25 = \frac{\exp\{3 + \exp(0)d_3\}}{1 + \exp\{3 + \exp(0)d_3\}}$$

and obtain $d_3 = -4.10$. Suppose further that $p_{01} = 0.05, p_{02} = 0.12, p_{04} = 0.40$, and $p_{05} = 0.55$. Then in the same manner, we obtain $d_1 = -5.94, d_2 = -4.99, d_4 = -3.41$, and $d_5 = -2.80$. Table 3.1 shows the dose labels obtained by *backward substitution* (3.5) under various dose–toxicity functions using the same skeleton. Note that the range of d_k varies with the model, and may take on negative values. However, it is always true that $d_1 < \dots < d_K$.

Table 3.1 Dose labels via backward substitution under four CRM models for $K = 5$ with $p_{01} = 0.05, p_{02} = 0.12, p_{03} = 0.25, p_{04} = 0.40, p_{05} = 0.55$, and prior mean $\hat{\beta}_0 = 0$

Model	d_1	d_2	d_3	d_4	d_5
Empiric (3.2)	0.05	0.12	0.25	0.40	0.55
Logistic (3.3) with $a_0 = 0$	-2.94	-1.99	-1.10	-0.41	0.20
Logistic (3.3) with $a_0 = 3$	-5.94	-4.99	-4.10	-3.41	-2.80
Hyperbolic tangent (3.4)	-1.47	-1.00	-0.55	-0.20	0.10

The backward substitution (3.5) ensures the dose–toxicity model F provides an exact fit over the initial guesses of toxicity probabilities, which ideally should reflect the clinicians’ prior beliefs. This is a crucial step particularly because of the use of underparameterized (one-parameter) model. In practice, it is often unrealistic for the clinicians to provide reliable guesses for all test doses prior to a study. Rather, we take the approach by which the skeleton $\{p_{0k}\}$ is numerically calibrated to yield good design’s operating characteristics. We shall return to this in Chapter 8.

3.2.4 Model-Based MTD

The CRM starts a trial by treating the first patient at the prior MTD v_0 , that is, $x_1 = d_{v_0}$. By (3.6), this starting dose is the dose initially believed to have toxicity probability (closest to) θ . Each subsequent x_i is determined sequentially based on the previous observation history $H_i = \{(x_j, Y_j) : j < i\}$ for $i \geq 2$. A CRM design \mathcal{D}_1 can be viewed as a function defined on the increasing H_i . The basic idea is to treat the next patient at the model-based MTD estimate, given H_i . Precisely,

$$x_i = \mathcal{D}_1(H_i) = \arg \min_{d_k} |F(d_k, \hat{\beta}_{i-1}) - \theta| \quad (3.7)$$

where

$$\hat{\beta}_{i-1} = \frac{\int_{-\infty}^{\infty} \beta L_{i-1}(\beta) dG(\beta)}{\int_{-\infty}^{\infty} L_{i-1}(\beta) dG(\beta)}$$

is the posterior mean of β given H_i and

$$L_{i-1}(\beta) = \prod_{j=1}^{i-1} \{F(x_j, \beta)\}^{Y_j} \{1 - F(x_j, \beta)\}^{1-Y_j} \quad (3.8)$$

is the binomial likelihood. The assignment (3.7) continues in a sequential fashion until a prespecified sample size N is reached. For the CRM, the final MTD estimate is given by $x_{N+1} = \mathcal{D}_1(H_{N+1})$, that is, a dose would have been given to the $(N+1)$ st patient enrolled to the trial.

In other words, the CRM attempts to treat the next patient at the current best guess of the MTD, a dose with toxicity probability estimated to be closest to the target θ . The motivation of this algorithm is to correct for the deficiency of the 3+3 algorithm, which treats the majority of the subjects at low and inefficacious doses. While this is ethically sound on a conceptual level, there may be various ways to calculate the “best” dose on the implementation level. The model-based MTD (3.7) is obtained by estimating $p_k = \pi(d_k)$ with a plug-in estimate $F(d_k, \hat{\beta}_{i-1})$. An alternative MTD estimate is

$$\mathcal{D}_1^*(H_i) = \arg \min_{d_k} |E_{i-1}\{F(d_k, \beta)\} - \theta| \quad (3.9)$$

where $E_{i-1}(\cdot)$ denotes expectation computed with respect to the posterior given H_i , that is,

$$E_{i-1}\{F(d_k, \beta)\} = \frac{\int_{-\infty}^{\infty} F(d_k, \beta) L_{i-1}(\beta) dG(\beta)}{\int_{-\infty}^{\infty} L_{i-1}(\beta) dG(\beta)}.$$

The MTD estimate (3.9) involves the computation of K integrals at each interim, and is a more formal Bayesian estimate of p_k than (3.7). In the early CRM literature, the plug-in estimate (3.7) emerged to be the convention because of its computational ease (although computational consideration is of much less importance today). Also, the estimate (3.9) as a formal Bayesian estimate is conceptually advantageous only when F is a correct model for π . As we will see in Chapter 5, an attractive feature of the model-based CRM is that its performance does not rely on correct specification of F . Hence, we will focus on the plug-in CRM (3.7), which is studied more thoroughly and systematically than the other estimators in the literature.

3.2.5 Normal Prior on β

Now, suppose that β has a normal prior distribution with mean $\hat{\beta}_0$. For the logistic function (3.3), applying backward substitution (3.5) gives

$$d_k = \frac{\text{logit}(p_{0k}) - a_0}{\exp(\hat{\beta}_0)}$$

where $\text{logit}(p) = \log\{p/(1-p)\}$. As a result, we have

$$F(d_k, \beta) = \frac{\exp\left[a_0 + \exp(\beta - \hat{\beta}_0) \{\text{logit}(p_{0k}) - a_0\}\right]}{1 + \exp\left[a_0 + \exp(\beta - \hat{\beta}_0) \{\text{logit}(p_{0k}) - a_0\}\right]}.$$

Since $F(d_k, \beta)$ depends on the parameter β only via $\beta - \hat{\beta}_0$, which is mean zero normal, we may arbitrarily set $\hat{\beta}_0 = 0$ without affecting the computation. The logistic model (3.3) is therefore invariant to the mean of a normal prior distribution. This invariance property holds for the general class of dose–toxicity models described in Chapter 4.

3.2.6 Implementation in R

The R package ‘dfcrm’ consists of functions for the implementation and the design of the Bayesian CRM using the empiric (3.2) and logistic (3.3) models. In particular, the function `crm` takes cumulative patient data and returns a dose for the next patient according to the model-based estimate (3.7).

```
> ### Return the recommended dose level for patient 6
> ### based on data from five patients
> library(dfcrm)
> p0 <- c(0.05,0.12,0.25,0.40,0.55) # initial guesses
> theta <- 0.25 # target toxicity rate
> y <- c(0, 0, 1, 0, 0) # toxicity indicators
> lev <- c(3, 5, 5, 3, 4) # dose levels
> fooB <- crm(p0,theta,y,lev,model="logistic",intcpt=3)
> fooB$estimate # posterior mean of beta
[1] 0.2794614
> fooB$mtd
[1] 4
> fooB$doses # Dose labels
[1] -5.944439 -4.992430 -4.098612 -3.405465 -2.799329
>
```

The above R code illustrates the usage of `crm` for a trial with $K = 5$ test doses and a target $\theta = 0.25$; the function is applied to observations from the first 5 subjects who receive dose levels 3, 5, 5, 3, and 4, where the third patient has a toxic outcome. The one-parameter logistic function (3.3) is used when the argument `model` is specified

as “logistic”. The default intercept value is 3, and can be modified by the argument `intcpt`; thus, the specification “`intcpt=3`” in the above illustration is redundant. When no value for the argument `model` is provided, `crm` will use the empiric model (3.2) as the default. The function computes the posterior mean of β by using a normal prior with mean 0. The prior standard deviation is specified by the argument `scale`; if `scale` is not specified (as in the last illustration), the default value is $\sqrt{1.34}$.

One-stage Bayesian CRM

The one-stage Bayesian CRM requires the specification of an array of *design parameters* which can be classified into clinical parameters (or clinician-input parameters) and model parameters. Planning and implementation of the method in a dose finding trial takes three steps:

1. *Setting the clinical parameters:*

- Target toxicity probability θ
- Number of test doses K
- Prior MTD v_0
- Sample size N

2. *Calibrating the model parameters:*

- Functional form of the dose–toxicity model $F(\cdot, \beta)$
- Skeleton $\{p_{0k}\}$ and hence the dose labels d_k s via backward substitution
- Prior distribution $G(\beta)$ of β

3. *Execution:* Treat the first patient at v_0 , and the subsequent patients at the most recent model-based MTD (3.7) as data accrue throughout the trial.

3.3 Two-Stage CRM

3.3.1 Initial Design

There are two practical difficulties associated with the use of the one-stage Bayesian CRM. First, treating the first patient at the prior MTD rather than the lowest dose level may raise safety concerns. Second, the use of a prior distribution $G(\beta)$ may be viewed as subjective and arbitrary.

To address the first difficulty, several authors suggest starting a CRM trial at the lowest dose, and applying dose escalation restrictions when the model-based MTD appears aggressive. This approach can be represented by a two-stage CRM, defined as follows. First, specify an initial design as a predetermined nondecreasing dose sequence $\{x_{i,0}\}$ such that $x_{i-1,0} \leq x_{i,0}$. Then a two-stage CRM $\mathcal{D}_2(H_i)$ is defined as

$$\mathcal{D}_2(H_i) = \begin{cases} x_{i,0}, & \text{if } Y_j = 0 \text{ for all } j < i, \\ \mathcal{D}_1(H_i) & \text{if } Y_j = 1 \text{ for some } j < i. \end{cases} \quad (3.10)$$

In other words, the initial design is in effect until the first observed toxicity; once a

toxic outcome is observed, the trial turns to the model-based CRM for dose assignments. Because the 3+3 algorithm is familiar, there is an inclination to consider the “group-of-three” initial design by which escalation takes place after every group of three nontoxic outcomes. That is,

$$x_{1,0} = x_{2,0} = x_{3,0} = d_1, x_{4,0} = x_{5,0} = x_{6,0} = d_2, x_{7,0} = x_{8,0} = x_{9,0} = d_3 \dots$$

and so on. There is, however, no clear justification for using the group-of-three rule apart from convention. In fact, this initial design is sometimes not in line with the motivation of the CRM. We will study the calibration of the initial dose sequences in Chapter 10.

3.3.2 Maximum Likelihood CRM

In response to the difficulty associated with the subjectivity of the Bayesian approach, we may use maximum likelihood estimation in conjunction with the CRM [80]. The idea is simple and analogous to the Bayesian CRM: with data observed in the first $i - 1$ patients, the dose for the next subject is computed as

$$x_i = \tilde{\mathcal{D}}_1(H_i) = \arg \min_{d_k} |F(d_k, \tilde{\beta}_{i-1}) - \theta| \quad (3.11)$$

where $\tilde{\beta}_{i-1} = \arg \max_{\beta} L_{i-1}(\beta)$ is the maximum likelihood estimate (mle) of β for given H_i .

The R function `crm` also implements the maximum likelihood CRM through the specification of the argument `method`:

```
> ### Compute the recommended dose level for patient 6
> ### using maximum likelihood CRM on the same data
> ### i.e., same values of p0, theta, y, lev
> fooL <- crm(p0,theta,y,lev,model="logistic",method="mle")
> fooL$estimate # mle of beta
[1] 0.3142946
> fooL$mtd
[1] 5
>
```

The function `crm` evaluates the maximum likelihood estimate of MTD (3.11) through the specification of the `method` argument as “mle”. When `method` is not specified, the Bayesian CRM is assumed.

Using (3.11) presupposes the existence of mle of β . For a one-parameter model F , the mle $\tilde{\beta}_{i-1}$ exists if and only if there is heterogeneity in the toxicity outcomes among patients, that is, $Y_j = 0$ and $Y_{j'} = 1$ for some $j, j' < i$. In the last R illustration, there is one toxic outcome out of 5 patients, and thus $\tilde{\beta}_5$ exists. In general, when one plans to use the maximum likelihood CRM, it is necessary to consider a two-stage design, which can be formed by replacing $\mathcal{D}_1(H_i)$ with $\tilde{\mathcal{D}}_1(H_i)$ in (3.10).

The posterior mean $\hat{\beta}_{i-1}$ and the mle $\tilde{\beta}_{i-1}$ are generally different, thus leading

to potentially different dose recommendations. For example, in the above R codes, the Bayesian CRM chooses dose level 4 for the sixth patient, whereas the maximum likelihood CRM chooses level 5. However, the posterior mean $\hat{\beta}_5 = 0.279$ and mle $\tilde{\beta}_5 = 0.314$ are only slightly different. Looking at the model-based estimates of the toxicity probabilities reveals that both dose levels 4 and 5 are roughly equally apart from the target $\theta = 0.25$:

```
> round(fooB$ptox,digits=2) # Posterior toxicity rates
[1] 0.01 0.03 0.08 0.18 0.33
> round(fooL$ptox,digits=2) # Maximum likelihood estimates
[1] 0.01 0.02 0.07 0.16 0.30
>
```

The difference in the estimation of β will diminish as sample size increases; and eventually, the choice of the estimation method per se will have minimal impact on the dose assignments. However, the estimation method may have different implied performance due to the fact that the maximum likelihood CRM always requires a two-stage strategy whereas Bayesian CRM is typically used as a one-stage design.

Two-stage CRM

The two-stage CRM requires the specification of an initial design as part of the design parameters in addition to those required in the one-stage CRM. Planning and implementation of a two-stage CRM in a dose finding trial takes four steps:

1. *Setting the clinical parameters:*

- Target toxicity probability θ
- Number of test doses K
- Prior MTD v_0
- Sample size N

2. *Calibrating the model parameters:*

- Functional form of the dose–toxicity model $F(\cdot, \beta)$
- Skeleton $\{p_{0k}\}$ and hence the dose labels d_{kS} via backward substitution[†]
- Prior distribution $G(\beta)$ of β , if Bayesian CRM is used

3. *Specifying an initial dose sequence:* $x_{1,0} \leq x_{2,0} \leq \dots \leq x_{N,0}$.

4. *Execution:* Treat patients initially according to $\{x_{i,0}\}$. Upon the first observed toxic outcome, treat the subsequent patient at the most recent model-based MTD (3.7) or (3.11) as data accrue throughout the trial.

[†]Backward substitution for the maximum likelihood CRM can be carried out as in (3.5) using an arbitrary initial value $\tilde{\beta}_0$ without affecting the dose calculations. More details will be given in Chapter 4.

3.4 Simulating CRM Trials

3.4.1 Numerical Illustrations

Computer simulation is a primary tool for evaluating the aggregate performance of the method. To get a sense of how the CRM works, it is also useful to examine individual simulated trials. Figure 3.1 shows the outcomes of simulated trials using a one-stage and two-stage CRM for a trial with target $\theta = 0.25$ and $K = 5$ under the true dose–toxicity scenario

$$p_1 = 0.02, p_2 = 0.04, p_3 = 0.10, p_4 = 0.25, p_5 = 0.50. \quad (3.12)$$

Hence, dose level 4 is the true MTD. The dose assignments by the two methods in the figure are quite different, although both select the correct MTD (dose level 4) based on the 20 simulated subjects. The one-stage CRM, taking advantage of a high starting dose, treats the majority of the subjects at the MTD but also seven patients at an overdose. In contrast, the two-stage CRM takes a conventional escalation approach initially and potentially treats many patients at low doses. As a result, the one-stage CRM causes two toxic outcomes (patients 3 and 9) more than the two-stage CRM. It is debatable whether the one-stage CRM is unsafe. On the one hand, the CRM has been criticized on account of ethical concerns, as it is shown to cause more toxic outcomes than the standard 3+3 method under some dose–toxicity scenarios [54]. On the other hand, overdosing is not necessarily a worse mistake than underdosing when treating patients with severe diseases such as cancer. Rosa et al. [89] describe a case study in which starting at a low dose (per the 3+3 algorithm) causes an ethical dilemma when consenting a cancer patient. Overall, the risk–benefit trade-off should be considered on a case-by-case basis. While the bortezomib trial in Chapter 1 started at the third level, the NeuSTART (Example 2.1) exercised caution via a low starting dose. In addition, the one-stage CRM in Figure 3.1 causes a 25% observed toxicity rate (5 out of 20), which is on target. From a numerical viewpoint, the two-stage design appears to be overconservative in this particular simulated trial.

3.4.2 Methods of Simulation

When simulating toxicity outcomes in a trial, each patient may be viewed to be carrying a latent toxicity tolerance u_i that is uniformly distributed on the interval $[0, 1]$. If the uniform variate is smaller than the true toxicity probability of the dose assigned to the patient, the patient has a toxic outcome; otherwise, the patient does not have a toxic outcome. That is,

$$Y_i = \begin{cases} 1 & \text{if } u_i \leq \pi(x_i) \\ 0 & \text{otherwise} \end{cases}$$

Table 3.2 displays the latent tolerance of the 20 simulated patients used in the simulated trials in Figure 3.1. Consider, for example, patient 1 who receives dose level 3 according to the one-stage CRM (left panel of the figure). He does not have a toxic outcome because the tolerance $u_1 = .571 > p_3 = .10$. Consequently, using the

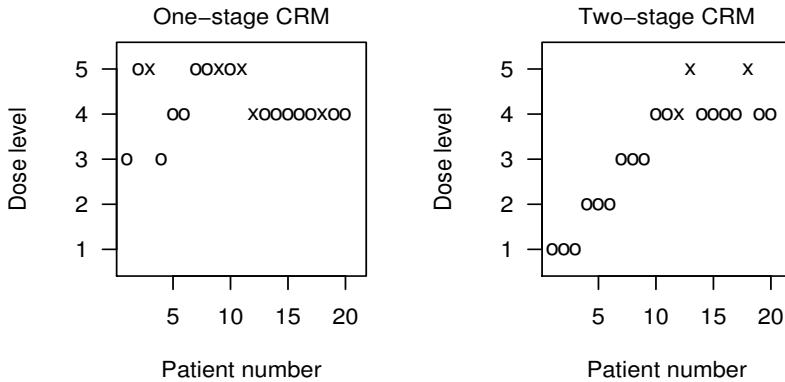


Figure 3.1 Simulated trials using the one-stage and two-stage Bayesian CRM in 20 subjects. The logistic model (3.3) with $a_0 = 3$ in Table 3.1 is used with $\beta \sim N(0, 1.34)$ a priori. For the two-stage CRM, the initial design escalates according to the group-of-three rule. Each point represents a patient, with “o” indicating no toxicity and “x” indicating toxicity.

same CRM model as in Figure 3.1, we obtain the posterior mean $\hat{\beta}_1 = 0.60$, which implies $x_2 = 5$. Table 3.2 also gives the numerical outputs of the one-stage CRM based on these 20 simulated patients.

Table 3.2 A simulated CRM trial through latent toxicity tolerance of 20 simulated patients

i	x_i	$\pi(x_i)$	u_i	y_i	$\hat{\beta}_i$	i	x_i	$\pi(x_i)$	u_i	y_i	$\hat{\beta}_i$
1	3	0.10	.571	0	0.60	11	5	0.50	.321	1	0.25
2	5	0.50	.642	0	0.93	12	4	0.25	.099	1	0.15
3	5	0.50	.466	1	0.04	13	4	0.25	.383	0	0.18
4	3	0.10	.870	0	0.18	14	4	0.25	.995	0	0.21
5	4	0.25	.634	0	0.28	15	4	0.25	.628	0	0.24
6	4	0.25	.390	0	0.34	16	4	0.25	.346	0	0.26
7	5	0.50	.524	0	0.41	17	4	0.25	.919	0	0.28
8	5	0.50	.773	0	0.47	18	4	0.25	.022	1	0.21
9	5	0.50	.175	1	0.31	19	4	0.25	.647	0	0.22
10	5	0.50	.627	0	0.35	20	4	0.25	.469	0	0.24

In a real trial, the latent tolerance u_i is not observable. In computer simulation, on the other hand, toxicity tolerance can be easily generated and is a useful tool to make different designs comparable in experiments where the dose assignments are made adaptively. For example, the same latent tolerance sequence in Table 3.2 can be used to generate the two-stage CRM in Figure 3.1 (right panel) so that both the

one-stage and two-stage designs are treating the same patients, while the patients are not necessarily treated at the same doses. The concept of toxicity tolerance is also instrumental to the construction of a nonparametric optimal design. We will return to this in Chapter 5.

The function `crmsim` in the ‘`dferm`’ package can be used to simulate multiple CRM trials under a given dose–toxicity curve. The following R code runs 10 trials using the one-stage Bayesian CRM specified in Figure 3.1 under the dose–toxicity scenario (3.12):

```
> ### Generate 10 CRM trials
> theta <- 0.25
> PI <- c(0.02,0.04,0.10,0.25,0.50) # True MTD = 4
> N <- 20 # sample size
> x0 <- 3 # starting dose level
> foo10 <- crmsim(PI,p0,theta,N,x0,nsim=10,model="logistic",restrict=F)
simulation number: 1
simulation number: 2
simulation number: 3
simulation number: 4
simulation number: 5
simulation number: 6
simulation number: 7
simulation number: 8
simulation number: 9
simulation number: 10
> foo10$MTD # Display the distribution of recommended MTD
[1] 0.0 0.0 0.1 0.8 0.1
>
```

In this illustration, dose level 4 is selected as the MTD in 8 of the 10 simulated trials.

3.5 Practical Modifications

3.5.1 Dose Escalation Restrictions

The one-stage CRM in Figure 3.1 assigns dose level 5 to the second patient after the nontoxic outcome in the first patient who receives dose level 3. This escalation may raise safety concerns because a high dose is tested without testing an intermediate dose level. Several authors have noted the potential problem with dose skipping by the CRM, and proposed the restricted CRM by imposing an escalation restriction: The dose level for the next patient cannot be more than one level higher than that of the current patient. Likewise, it is possible for the unrestricted CRM to skip doses in deescalation; cf. patient 4 under the one-stage CRM in Figure 3.1. While skipping doses in deescalation has not been perceived to be as problematic or unsafe, we may apply a similar restriction that the dose level for the next patient cannot be more than one level lower than that of the current patient. At any rate, these restrictions against dose skipping will typically be applied, if ever, only to the first few patients because

the change in the model-based estimates diminishes as the number of observations increases (try plotting the sequence of $\hat{\beta}_i$ in Table 3.2).

Another pathology in escalation is illustrated in the two-stage CRM in Figure 3.1: the dose for patient 13 is escalated from that of patient 12, who has a toxic outcome. Such an escalation is called incoherent as it puts patient 13 at undue risk of toxicity in light of the outcome in the previous patient [16]. To avoid the potential incoherent moves by the CRM, we may apply the restrictions:

Coherence in escalation: If a toxicity is observed in the current patient, then the dose level for the next patient cannot be higher than that of the current patient.

Coherence in deescalation: If the current patient does not experience toxicity, then the dose level for the next patient cannot be lower than that of the current patient.

It is noteworthy that there is no incoherent move by the one-stage CRM in Figure 3.1. That is, escalation occurs only after a nontoxic outcome, and deescalation after a toxic outcome. In fact, the one-stage CRM will never induce an incoherent move—even if no restriction is applied. This property indicates that the model-based CRM is doing the right thing ethically. In contrast, an unrestricted two-stage CRM may yield incoherent escalation as seen in Figure 3.1. It turns out that the problem is due to poor calibration of the initial design. In this particular example, the group-of-three escalation sequence is not an appropriate choice. This topic will be further discussed in Chapters 5 and 10.

Figure 3.2 shows the outcomes of the *restricted* CRM trials using the same model specifications as in Figure 3.1 and the same 20 patients as in Table 3.2. The dose assignment patterns are similar to that of the unrestricted counterparts. In general, the dose escalation restrictions have little impact on the CRM’s aggregate behaviors in terms of the probability of correctly selecting the MTD and the average toxicity number; try Exercise 3.4. Instead, these restrictions modify the *pointwise properties* of the design: a pointwise property of a dose finding design concerns the behaviors of individual outcome sequences (“points”). The pointwise properties of a design should agree with principles deemed sensible to clinicians. Violations of any such principles (e.g., making incoherent moves) will be perceived as worse than showing poor statistical properties. After all, only a single outcome sequence is observed in an actual trial. The aggregate operating characteristics are comparatively abstract to clinicians.

In the rest of this book, unless specified otherwise, we assume that the restriction of no dose skipping in escalation and the coherence restrictions are in effect. This is the assumption used in the function `crmsim` if no value is given to the argument `restrict`.

3.5.2 Group Accrual

The fully sequential CRM follows the current patient over an evaluation period, called an *observation window*, before enrolling the next patient. A long observation

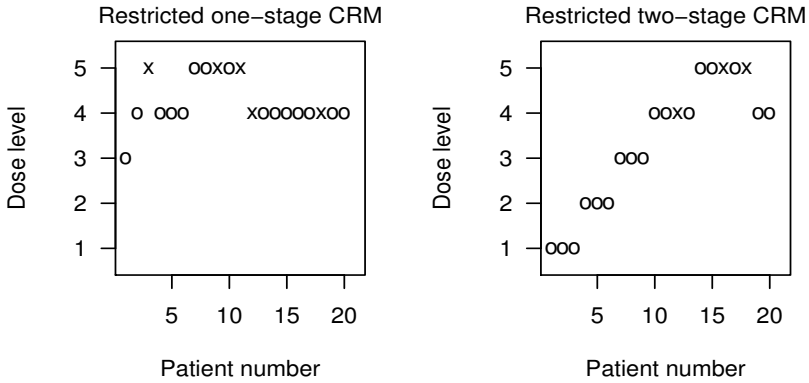


Figure 3.2 Simulated trials using the one-stage and the two-stage Bayesian CRM with dose escalation restrictions. Each point represents a patient, with “o” indicating no toxicity and “x” indicating toxicity.

window (e.g., 6 months) will likely result in repeated interim accrual suspensions, impose excessive administrative burdens, and cause long trial duration.

Goodman et al. [43] address this problem by assigning $m > 1$ patients at a time to each dose, and updating the model-base estimate (3.7) between every group of m patients. Figure 3.3 shows the outcomes of the simulated trials by group accrual CRM, using the models in Figure 3.1 and the same 20 patients in Table 3.2. In the left panel of Figure 3.3, the model-based MTD estimate (3.7) is updated after every $m = 2$ patients. The dose assignment pattern is similar to that of the fully sequential CRM in Figure 3.1: half of the study subjects receive the MTD. The use of a larger group size with $m = 4$, as shown in the right panel of the figure, gives the CRM fewer occasions to adapt to the previous outcomes. This may partly explain the findings that the accuracy of MTD estimation decreases as a large group size is used, when the assumed model F is not a correct specification of the true dose–toxicity curve [43]. Therefore, although a large group size reduces the number of interim calculations and shortens trial duration, it also undermines the advantage of the adaptiveness of the CRM.

When the observation window is long in comparison to the recruitment period, the reduction in trial duration by group accrual may not be adequate. Suppose, for instance, each patient in the current group is to be followed for 6 months before the next group is enrolled in Figure 3.3. Assuming instantaneous patient availability, a group CRM with size 4 will take roughly 24 months to enroll 20 patients, whereas a group size of 2 will take 54 months. We will return to consider clinical settings with a long observation window due to late toxicities in Chapter 11, where we introduce the time-to-event continual reassessment method (TITE-CRM) as an alternative to the group accrual CRM.

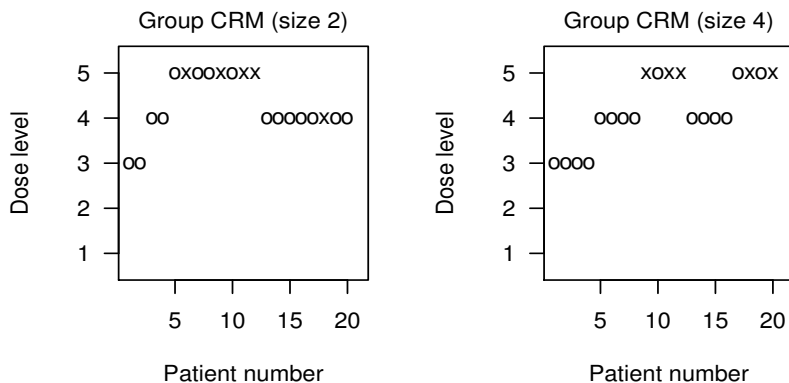


Figure 3.3 Simulated trials using the one-stage group accrual CRM. The MTD estimate is updated after every two patients in the left panel, and every four patients in the right panel. Dose escalation restrictions are applied. Each point represents a patient, with “o” indicating no toxicity and “x” indicating toxicity.

3.5.3 Stopping and Extension Criteria

It is common to have provisions for early termination in clinical trials for economic and ethical reasons. The standard 3+3 design for dose finding trials stops a trial when two toxic outcomes are observed at a dose. In contrast, the original CRM works with a fixed sample size N . From a practical viewpoint, neither the economic nor ethical reason is sufficiently compelling for stopping a CRM trial early. Economically, the sample size of a phase I trial is already small, and there is not much room to reduce the sample size unless the early stopping rules are liberal. Ethically, the CRM is expected to converge to the MTD as the trial accrues patients. Therefore, patients enrolled towards the end of the trial will likely receive a good dose, and thus there will be a reduced ethical imperative to stop the trial.

On the other hand, there may be good reasons for extending enrollment with the CRM in some situations. Goodman et al. [43] consider continuing a CRM trial beyond N until the recommended MTD has at least a certain number of patients assigned to it. This is to avoid the case where only few patients are treated at the recommended MTD. Consider the two-stage CRM in Figure 3.2 (right panel). Based on the 20 observations shown in the figure, the recommended MTD is dose level 5, which is one level higher than the dose given to the final patients, and has only had 5 patients with an observed toxicity rate of 40%. These are signs that the recommended MTD is not adequately assessed. Figure 3.4 shows the outcomes of an extended CRM trial with a minimum sample size criterion: the trial will go beyond 20 patients until at least 9 patients have been treated at the recommended MTD. As a result of this modification, the trial continues for another 7 patients. With a total of 27

subjects, there is evidence that dose level 5 exceeds the acceptable toxicity level (44% observed toxicity rate) and the final MTD is dose level 4.

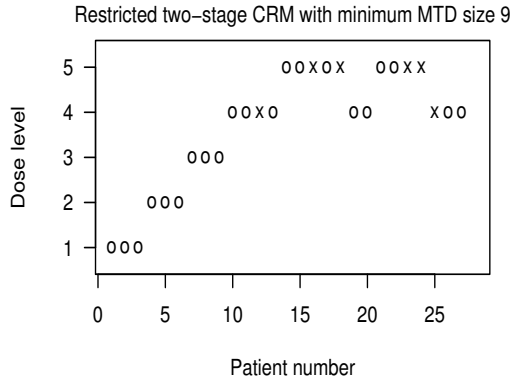


Figure 3.4 Simulated trials using the restricted two-stage CRM with a minimum sample size criterion. Dose escalation restrictions are applied. Each point represents a patient, with “o” indicating no toxicity and “x” indicating toxicity.

3.6 Bibliographic Notes

O’Quigley and Chevret [75] and Chevret [25] examine the operating characteristics of the one-stage Bayesian CRM by simulation study. Moller [67] and O’Quigley and Shen [80] study the two-stage CRM; the latter also propose the use of maximum likelihood estimation. Several authors, including Faries [36], Korn et al. [54], and Goodman et al. [43], note the potential problem with dose skipping by the CRM and propose to modify the CRM to limit escalation by no more than one level at a time. Faries [36] also suggests enforcing coherence in escalation by restriction; the notion of coherence is subsequently formalized in Cheung [16]. Ahn [1] compares various variants of the CRM by simulations. Early stopping of the CRM is considered in Heyd and Carlin [45] and O’Quigley and Reiner [79]. A recent overview of the CRM is given in Garrett-Mayer [39] and Iasonos et al. [48].

3.7 Exercises and Further Results

Exercise 3.1. Apply backward substitution (3.5) to the empiric function (3.2) with $\hat{\beta}_0 = 0$, and show that

$$F(d_k, \beta) = p_{0k}^{\exp(\beta)}$$

for a given skeleton $\{p_{01}, \dots, p_{0K}\}$.

Exercise 3.2. Use the toxicity tolerance in Table 3.2 to verify the outcomes of the two-stage CRM in Figure 3.1.

Exercise 3.3. The function `crmsim` can implement the two-stage CRM by specifying the argument `x0` by the initial dose sequence, whose length should equal the specified sample size `N`. Modify the R code in Section 3.4.2 to run 20 simulated trials using the two-stage CRM in Figure 3.1 with a group-of-three initial rule.

Exercise 3.4. By default, the function `crmsim` runs the restricted CRM with group size 1. These options can be modified through the arguments `restrict` and `mcohort`. Usage of these options is documented in the reference manual of ‘`dfcrm`’ at the R Web site:

<http://www.r-project.org>.

Run simulations, and examine the effects of escalation restrictions on the distribution of recommended MTD under dose–toxicity curve (3.12).

One-Parameter Dose–Toxicity Models

4.1 Introduction

One-parameter functions are predominantly used in the CRM literature to model the dose–toxicity relationship. An introduction on the use of one-parameter functions is in order—for two reasons. First, these functions are underparameterized and cannot be expected to produce a realistic fit on the entire dose range. Rather, they should be flexible enough to provide a reasonable approximation *locally* around the target dose. Second, the use of one-parameter models is not common in the other statistical areas, and modeling in the context of CRM is quite different from the conventional approach. This chapter presents a unified framework for CRM models. Section 4.2 introduces a class of dose–toxicity functions that includes the most commonly used CRM models. The required assumptions on a one-parameter CRM model are given in Section 4.3. Proof of the main result in this chapter is given in Section 4.4.

4.2 ψ -Equivalent Models

Consider the following class of dose–toxicity functions:

$$F(x, \beta) = \psi\{c(\beta)h(x)\} \quad (4.1)$$

where the parameter β is a scalar, and the functions ψ , c , and h are strictly monotone and known. This class of functions includes the most commonly used dose–toxicity curves in the CRM literature. For example, the empiric function (3.2) corresponds to

$$\psi(z) = \exp(z), c(\beta) = \exp(\beta), h(x) = \log(x) \quad (4.2)$$

so that

$$F(x, \beta) = \exp\{\exp(\beta) \log(x)\} = x^{\exp(\beta)}.$$

Also, the logistic function (3.3) with a fixed intercept a_0 is represented by

$$\psi(z) = \frac{\exp(a_0 + z)}{1 + \exp(a_0 + z)}, c(\beta) = \exp(\beta), h(x) = x.$$

Recall that the dose labels in the CRM are obtained by backward substitution so that $F(d_k, \hat{\beta}_0) = p_{0k}$, where p_{0k} is the initial guess of toxicity probability for dose

level k , and that $\hat{\beta}_0$ is the prior mean of β in a Bayesian CRM or an initial value of β in a maximum likelihood CRM. Therefore, a CRM model is defined by the function form $F(\cdot, \beta)$ of the dose-toxicity curve, the skeleton $\{p_{01}, \dots, p_{0K}\}$, and the initial value $\hat{\beta}_0$. (However, see Theorem 4.1 below.) Denote

$$F_k(\beta) = F(d_k, \beta) \text{ for } k = 1, \dots, K.$$

Under (4.1), the ψ -representation of a CRM model can be written as

$$F_k(\beta) = \psi\{c(\beta)h(d_k)\} = \psi\left\{\frac{c(\beta)}{c(\hat{\beta}_0)}\psi^{-1}(p_{0k})\right\} \quad (4.3)$$

which depends on the dose index k via the initial guess p_{0k} . In order for the model to distinguish between doses, therefore, the skeleton needs to be specified as a *strictly* increasing sequence. Furthermore, the representation (4.3) does not depend on the function $h(x)$. Thus, two CRM models are identical if their dose-toxicity functions are represented by the same $\psi(z)$ and $c(\beta)$. For example, the hyperbolic tangent function (3.4) corresponds to

$$\psi(z) = \exp(z) \text{ and } c(\beta) = \exp(\beta)$$

with $h(x) = \log\{(\tanh x + 1)/2\}$. Since the function is represented by the same $\psi(z)$ and $c(\beta)$ as in (4.2), the CRM model defined by the hyperbolic tangent function is identical to that by the empiric function. The ψ -representation (4.3) of both functions is

$$F_k(\beta) = \exp\left\{\frac{\exp(\beta)}{\exp(\hat{\beta}_0)}\log(p_{0k})\right\} = p_{0k}^{\exp(\beta - \hat{\beta}_0)}.$$

Example 4.1. The logistic function with a fixed slope $a_1 > 0$:

$$F(x, \beta) = \frac{\exp(\beta + a_1x)}{1 + \exp(\beta + a_1x)} \quad (4.4)$$

corresponds to $\psi(z) = z/(1+z)$ and $c(\beta) = \exp(\beta)$, which are free from the choice of a_1 . After backward substitution, the model can be represented by

$$F_k(\beta) = \frac{\exp(\beta - \hat{\beta}_0)p_{0k}}{1 - p_{0k} + \exp(\beta - \hat{\beta}_0)p_{0k}}. \quad (4.5)$$

The function (4.4) leads to the same CRM model regardless of the value of a_1 . One may therefore set $a_1 = 1$ arbitrarily.

Example 4.1 illustrates how the ψ -representation can be used to identify identical CRM models, and reduce unnecessary comparisons of different functions (e.g., a_1).

Definition 4.1 (ψ -equivalent models). Two CRM models under class (4.1) are said to be ψ -equivalent if their dose-toxicity functions can be represented by the same $\psi(z)$.

While ψ -equivalent models are not necessarily identical, the difference arises only as a result of different parameterizations. For example, another form (3.1) of the empiric function,

$$F(x, \beta) = x^\beta,$$

can be represented by $\psi(z) = \exp(z)$ and $c(\beta) = \beta$, and is ψ -equivalent to (3.2). While (3.1) and (3.2) are not equivalent, two functions differ only in terms of how the parameter appears in the functions. Hence, both functions will lead to identical estimation of the toxicity probabilities when maximum likelihood estimation is used.

Theorem 4.1. *Suppose that $F_k^{(1)}(\beta)$ and $F_k^{(2)}(\phi)$ are derived from two ψ -equivalent models and the same skeleton $\{p_{0k}\}$. Let $\hat{\beta}_{i-1}$ and $\tilde{\phi}_{i-1}$ be the respective maximum likelihood estimates of β and ϕ for the two models given the observation history H_i . Then*

$$F_k^{(1)}(\hat{\beta}_{i-1}) = F_k^{(2)}(\tilde{\phi}_{i-1})$$

for all k .

In words, Theorem 4.1 states that the maximum likelihood CRM is invariant within a ψ -equivalent class of CRM models. The proof is an extension of the invariance property of maximum likelihood estimation, and is given in Section 4.4. Interestingly, Theorem 4.1 holds regardless of the initial values of β and ϕ , thus implying that we can arbitrarily choose the initial values, $\hat{\beta}_0$ and $\hat{\phi}_0$, used in the backward substitution step (3.5). Without loss of generality, from now on, we will set the initial values so that $c(\hat{\beta}_0) = c(\hat{\phi}_0) = 1$.

Example 4.2. The dose–toxicity function

$$F(x, \phi) = \frac{\phi x^2}{a_2 + \phi x^2} \text{ for some } a_2 > 0 \tag{4.6}$$

can be represented by $\psi(z) = z/(1+z)$ with $c(\phi) = \phi$ and $h(x) = x^2/a_2$. Therefore, the CRM model generated by (4.6) is ψ -equivalent to that by the logistic function with a fixed slope; cf., equation (4.4) in Example 4.1. That is, both models will yield identical dose assignments if maximum likelihood CRM is used.

In contrast, the Bayesian CRM (3.7) is not invariant among ψ -equivalent models. The model (4.6) in Example 4.2 can be represented by

$$F_k(\phi) = \frac{(\phi/\hat{\phi}_0) p_{0k}}{1 - p_{0k} + (\phi/\hat{\phi}_0) p_{0k}} \tag{4.7}$$

where $\hat{\phi}_0$ is the prior mean of ϕ . Bayesian CRM using models (4.5) and (4.7) would lead to identical posterior computations if $\exp(\beta - \hat{\beta}_0)$ and $\phi/\hat{\phi}_0$ had the same prior distribution. However, this is impossible because $E_0(\phi/\hat{\phi}_0) = 1$ by definition of $\hat{\phi}_0$, whereas

$$E_0 \left\{ \exp(\beta - \hat{\beta}_0) \right\} > \exp \left\{ E_0(\beta - \hat{\beta}_0) \right\} = 1$$

by Jensen's inequality. This example illustrates that parameterization via $c(\beta)$ can be an important consideration when the Bayesian CRM (3.7) is used.

When $c(\beta) = \exp(\beta)$, the ψ -representation (4.3) of a CRM model depends on the parameter β only via $\beta - \hat{\beta}_0$. If the prior distribution of β is normal, the centered variable $\beta - \hat{\beta}_0$ is always mean zero normal. Therefore, the posterior computations will be identical regardless of the specified prior mean $\hat{\beta}_0$, and the Bayesian CRM is invariant to the mean of a normal prior. In general, this invariance property holds for Bayesian CRM when model (4.1) is used and the prior distribution of β constitutes a location-scale family (Lehmann, 1983, page 20). In this book, we will focus on the model class (4.1) with $c(\beta) = \exp(\beta)$ and a normal prior distribution for β , so that the model parameters of the Bayesian CRM are the function $\psi(z)$, the skeleton $\{p_{0k}\}$, and the variance σ_β^2 of the normal prior. Table 4.1 gives some simple examples of $\psi(z)$ and the corresponding CRM models. The calibration of $\{p_{0k}\}$ will be detailed in Chapter 8, and specification of σ_β^2 in Chapter 9.

4.3 Model Assumptions

In this section, we state the regularity conditions on the one-parameter CRM model $F_k(\beta)$ assumed in this book. The role of these conditions in various theoretical results will be discussed in the following chapters. In this chapter, we focus on the intuition behind the conditions and their implications. A practical point: all these assumptions are verifiable for any given $F_k(\beta)$, and thus can serve as preliminary conditions to remove certain models from consideration for use.

Condition 4.1. $F(x, \beta)$ is strictly increasing in x for all β .

Condition 4.2. $F_k(\beta)$ is monotone in β in the same direction for all k .

Conditions 4.1 and 4.2 are satisfied by many dose-toxicity functions. In particular, the model class (4.1) satisfies these conditions when either $h(d_k) > 0$ for all k or < 0 for all k . It is easy to verify that the logistic model with $a_0 = 0$ in Table 3.1 does not satisfy Condition 4.2 because $h(d_k) < 0$ for $k \leq 4$ but $h(d_5) > 0$. Thus, this model should not be considered for use. Graphically, Condition 4.2 implies that the family of curves induced by $F(x, \beta)$ do not cross each other; see Figure 4.1.

For the next condition, we first define $g_{ij}(\beta) = \{1 - F_i(\beta)\} / \{1 - F_j(\beta)\}$.

Condition 4.3. The derivatives $F'_k(\beta)$ and $g'_{ij}(\beta)$ exist and $g'_{ij}(\beta)F'_k(\beta) \leq 0$ for all k and $i > j$.

Condition 4.3 can be equivalently stated as

$$|F'_i(\beta)| \geq |F'_j(\beta)| \{1 - F_i(\beta)\} / \{1 - F_j(\beta)\}$$

for all $i > j$, which puts a lower bound on $|F'_i(\beta)|$ relative to $|F'_j(\beta)|$. Using the fact that $\log\{1 - F_k(\beta)\} = \int_{-\infty}^{\beta} [-F'_k(\phi) / \{1 - F_k(\phi)\}] d\phi$ when $F'_k(\phi) > 0$ for all k , one can show that Condition 4.3 implies Condition 4.1 and hence is a stronger condition.

Table 4.1 Dose–toxicity functions and their ψ –representations with $c(\beta) = \exp(\beta)$, $h(x) = x$, and $\hat{\beta}_0 = 0$

Model	$\psi(z)$	$F_k(\beta)$	Range of β
Empiric (or hyperbolic tangent)	$\exp(z)$	$p_{0k}^{\exp(\beta)}$	$(-\infty, \infty)$
Complementary log-log	$1 - \exp(-z)$	$1 - (1 - p_{0k})^{\exp(\beta)}$	$(-\infty, \infty)$
Logistic; fixed intercept a_0	$\frac{\exp(a_0 + z)}{1 + \exp(a_0 + z)}$	$\frac{\exp[a_0 + \exp(\beta)\{\text{logit}(p_{0k}) - a_0\}]}{1 + \exp[a_0 + \exp(\beta)\{\text{logit}(p_{0k}) - a_0\}]}$	$(-\infty, \infty)$
Logistic; fixed slope	$\frac{z}{1 + z}$	$\frac{\exp(\beta)p_{0k}}{1 - p_{0k} + \exp(\beta)p_{0k}}$	$(-\infty, \infty)$
Product-of-beta (12.10)	$1 - \frac{1}{z}$	$1 - \exp(-\beta)(1 - p_{0k})$	$> \log(1 - p_{01})$
Probit; fixed intercept a_0	$\Phi(a_0 + z)$	$\Phi[a_0 + \exp(\beta)\{\Phi^{-1}(p_{0k}) - a_0\}]$	$(-\infty, \infty)$
Probit; fixed slope	$\Phi\{\log(z)\}$	$\Phi\{\beta + \Phi^{-1}(p_{0k})\}$	$(-\infty, \infty)$

Note: $\text{logit}(p) = \log\{p/(1 - p)\}$; $\Phi(z)$ indicates cdf of standard normal at z .

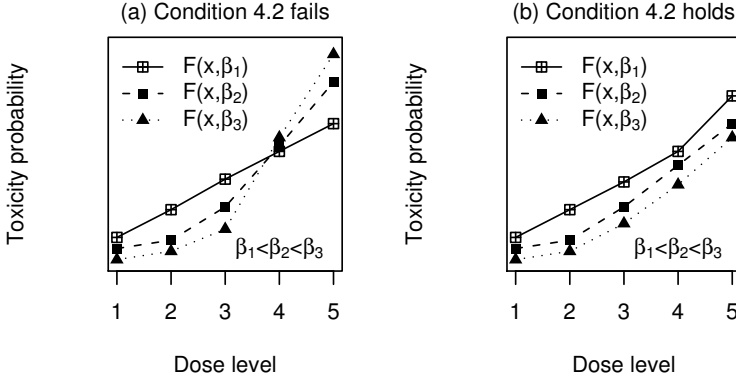


Figure 4.1 Graphical illustration of Condition 4.2. (a) A model $F_k(\beta)$ that is decreasing in β for $k = 1, 2, 3$, and increasing in β for $k = 4, 5$ (i.e., Condition 4.2 is violated). (b) A model $F_k(\beta)$ that is uniformly decreasing in β for all dose level k .

Example 4.3. The logistic model with fixed slope in Example 4.1 satisfies both Conditions 4.1 and 4.2, because (4.5) is increasing in k and increasing in β for all k . To further show that the model satisfies Condition 4.3, we first obtain that $F'_k(\beta) = F_k(\beta)\{1 - F_k(\beta)\} > 0$. Next, for $i > j$,

$$\begin{aligned}
 g'_{ij}(\beta) &= \frac{F'_j(\beta)\{1 - F_i(\beta)\} - F'_i(\beta)\{1 - F_j(\beta)\}}{\{1 - F_j(\beta)\}^2} \\
 &= \frac{F_j(\beta)\{1 - F_j(\beta)\}\{1 - F_i(\beta)\} - F_i(\beta)\{1 - F_i(\beta)\}\{1 - F_j(\beta)\}}{\{1 - F_j(\beta)\}^2} \\
 &= \frac{\{F_j(\beta) - F_i(\beta)\}\{1 - F_i(\beta)\}\{1 - F_j(\beta)\}}{\{1 - F_j(\beta)\}^2} < 0.
 \end{aligned}$$

Hence, $g'_{ij}(\beta)F'_k(\beta) < 0$ and Condition 4.3 is met.

Condition 4.4. For any given $\theta \in (0, 1)$ and a given v , there exists β in the parameter space such that $F_v(\beta) = \theta$.

Condition 4.5. Suppose $F_k(\beta)$ is decreasing in β . There exist the constants \underline{b} and \bar{b} such that $F_1(\underline{b}) > \theta$ and $F_K(\bar{b}) < \theta$ and $[\underline{b}, \bar{b}]$ is a compact subset of the parameter space of β .

Consider the ψ -representation of the empiric model in Table 4.1. Condition 4.4 is satisfied because we can set $\beta = \log\{\log(\theta)/\log(p_{0v})\} \in (-\infty, \infty)$ for any given v and $\theta \in (0, 1)$. Also, Condition 4.5 is met by setting

$$\underline{b} < \log\{\log(\theta)/\log(p_{01})\} \text{ and } \bar{b} > \log\{\log(\theta)/\log(p_{0K})\}.$$

Intuitively, Conditions 4.4 and 4.5 ensure that the model is sufficiently flexible to approximate dose–toxicity scenarios with all possible MTD v and any given θ , so that there always exists x^* that solves the equation $F(x^*, \beta) = \theta$; the uniqueness of x^* is guaranteed by Condition 4.1.

Condition 4.5 assumes $F_k(\beta)$ is decreasing in β . An analogous condition for $F_k(\beta)$ that is increasing in β can be also stated.

Condition 4.6. $F_k(\beta)$ is bounded away from 0 and 1 on $[\underline{b}, \bar{b}]$ for all k ; and $F'_k(\beta)$ is uniformly bounded in β .

Condition 4.7. For any given $0 < p < 1$ and each k , the function

$$p \frac{F'_k(\beta)}{F_k(\beta)} + (1-p) \frac{-F'_k(\beta)}{1-F_k(\beta)}$$

is continuous and strictly monotone in β .

Condition 4.8. For $\beta \in [\underline{b}, \bar{b}]$,

$$F_k(\beta)F''_k(\beta) - \{F'_k(\beta)\}^2 \leq 0 \tag{4.8}$$

and

$$-F''_k(\beta) + \left[F_k(\beta)F''_k(\beta) - \{F'_k(\beta)\}^2 \right] \leq 0 \tag{4.9}$$

with at least one inequality being strict for all k .

Condition 4.6 is introduced for technical brevity (in the proof of consistency), and can be relaxed by modifying the tails of the function F . Condition 4.7 is mild and is met by any dose–toxicity functions generated by a continuous $\psi(z)$ under model class (4.1). Condition 4.8 is an additional assumption required for the consistency of the Bayesian CRM using posterior mean, and is not required for the maximum likelihood CRM or Bayesian CRM using *posterior mode* (cf. Exercise 4.4).

If a dose–toxicity function satisfies the above conditions, its ψ -equivalent models will also satisfy the conditions. Therefore, verification of these conditions may be done under a parameterization that allows simple algebraic manipulations.

Example 4.4. To verify that the empiric model (3.2) satisfies Condition 4.8, we may instead work with model (3.1) for which $F'_k(\beta) = F_k(\beta) \log d_k$ and $F''_k(\beta) = F_k(\beta)(\log d_k)^2$. Hence, $F_k(\beta)F''_k(\beta) - \{F'_k(\beta)\}^2 = 0$, which satisfies (4.8); and $-F''_k(\beta) < 0$, which satisfies (4.9) and the inequality is strict under Condition 4.6.

One may also prove the empiric model satisfies Condition 4.8 using the functional form (3.2) directly; see Exercise 4.3.

4.4 Proof of Theorem 4.1

To prove Theorem 4.1, we first express the two models as

$$F_k^{(1)}(\beta) = \psi \left\{ \frac{c_1(\beta)}{c_1(\tilde{\beta}_0)} \psi^{-1}(p_{0k}) \right\} \text{ and } F_k^{(2)}(\phi) = \psi \left\{ \frac{c_2(\phi)}{c_2(\tilde{\phi}_0)} \psi^{-1}(p_{0k}) \right\},$$

where $\tilde{\beta}_0$ and $\tilde{\phi}_0$ are the respective initial values used in backward substitution for the two models. Now, since the models depend on the parameters through $c_1(\beta)/c_1(\tilde{\beta}_0)$ and $c_2(\phi)/c_2(\tilde{\phi}_0)$ in the same way, the likelihood is maximized at the same value

$$\frac{c_1(\tilde{\beta})}{c_1(\tilde{\beta}_0)} = \frac{c_2(\tilde{\phi})}{c_2(\tilde{\phi}_0)}, \text{ which implies } \frac{c_1(\tilde{\beta}_{i-1})}{c_1(\tilde{\beta}_0)} = \frac{c_2(\tilde{\phi}_{i-1})}{c_2(\tilde{\phi}_0)}$$

for a given H_i by invariance of maximum likelihood estimation. It thus follows that $F_k^{(1)}(\tilde{\beta}_{i-1}) = F_k^{(2)}(\tilde{\phi}_{i-1})$.

4.5 Exercises and Further Results

Exercise 4.1. Verify Conditions 4.1–4.3 for the empiric model (3.2).

Exercise 4.2. Verify Conditions 4.4 and 4.7 for the logistic model (3.3) with a fixed intercept a_0 . For given θ and p_{0k} s, determine the range of a_0 such that Condition 4.2 is met.

Exercise 4.3. Verify Condition 4.8 for the empiric model using the form (3.2), and compare your result with Example 4.4.

Exercise 4.4 (invariance of CRM using posterior mode). Consider the CRM that assigns the doses based on the posterior mode of β , that is,

$$x_i = \arg \min_{d_k} |F_k(\check{\beta}_{i-1}) - \theta|,$$

where $\check{\beta}_{i-1}$ is the maximum of the posterior density given H_i . Suppose that $F_k^{(1)}(\beta)$ and $F_k^{(2)}(\phi)$ are ψ -equivalent models with respect to the mode, and can be written as

$$F_k^{(1)}(\beta) = \psi \left\{ \frac{c_1(\beta)}{c_1(\check{\beta}_0)} \psi^{-1}(p_{0k}) \right\} \text{ and } F_k^{(2)}(\phi) = \psi \left\{ \frac{c_2(\phi)}{c_2(\check{\phi}_0)} \psi^{-1}(p_{0k}) \right\}.$$

Let $G_1(\beta)$ and $G_2(\phi)$ denote the prior distributions of β and ϕ . Prove that

$$F_k^{(1)}(\check{\beta}_{i-1}) = F_k^{(2)}(\check{\phi}_{i-1}) \text{ for all } k \quad (4.10)$$

if

$$G_1 \left[c_1^{-1} \left\{ c_1(\check{\beta}_0)z \right\} \right] = G_2 \left[c_2^{-1} \left\{ c_2(\check{\phi}_0)z \right\} \right] \text{ for all } z.$$

A direct consequence of (4.10) is that the CRM using the two models $F^{(1)}$ and $F^{(2)}$ will be identical if the prior distributions are matched. Similar results can also be established for CRM using the median.

Theoretical Properties

5.1 Introduction

This chapter presents some theoretical criteria for the evaluation of dose finding methods, and discusses the properties of the CRM in light of these criteria. First, the concept of coherence is introduced in Section 5.2. The asymptotic property of the CRM is then discussed in Section 5.3, where we also demonstrate its relevance in finite sample settings. Finally, Section 5.4 sketches the proofs of the main results in this chapter.

5.2 Coherence

5.2.1 Motivation and Definitions

Although the 3+3 dose escalation algorithm has fundamental deficiencies in terms of MTD estimation, it reflects an important ethical principle whereby dose escalation for the next patient is appropriate only when the most recent patient (or group of patients) does not show any sign of toxicity. This design feature, called coherence, limits the risk of unduly exposing patients to high and toxic doses. Conversely, if the most recent patient does not have a toxic outcome, it is counterintuitive to reduce the dose for the next patient, while the trial objective is to maximize the dose given to patients within safety limit. The use of outcome-adaptive design in dose finding studies indeed underlines these coherence considerations. Formally, coherence may be formulated as follows. Recall that Y_i denotes the toxicity indicator of patient i in a fully sequential trial. An escalation for patient i is said to be coherent only when $Y_{i-1} = 0$; and a dose finding design \mathcal{D} is said to be coherent in escalation if, with probability one,

$$P_{\mathcal{D}}(U_i > 0 | Y_{i-1} = 1) = 0 \text{ for all } i \quad (5.1)$$

where $U_i = x_i - x_{i-1}$ is the dose increment from patient $i-1$ to patient i , and $P_{\mathcal{D}}(\cdot)$ denotes the probability computed under the design \mathcal{D} . Likewise, a deescalation for patient i is said to be coherent only when $Y_{i-1} = 1$; and a design \mathcal{D} is said to be coherent in deescalation if with probability one

$$P_{\mathcal{D}}(U_i < 0 | Y_{i-1} = 0) = 0 \text{ for all } i. \quad (5.2)$$

In other words, a coherent design does not induce any incoherent move. To illustrate, consider the unrestricted two-stage CRM in Figure 3.1. The design escalates the dose

for patient 13 to dose level 5 after a toxic outcome is observed in patient 12. This is an incoherent escalation, and by definition, this two-stage CRM is not a coherent design. On the one hand, one could argue that the incoherent escalation for patient 13 cannot be justified ethically because it exposes the patient to a high dose without clearly demonstrating safety at the lower doses. On the other hand, one may anticipate that based on the outcomes in the first 11 patients, patient 12 would receive dose level 5 regardless of the outcome of patient 12; thus, with the objective to treat as many patients at the right dose as possible, it seems that the escalation should have occurred for patient 12. Either way, the fact that a design is incoherent indicates something not entirely in line with the ethics and objective of the trial.

5.2.2 Coherence Conditions of the CRM

The one-stage Bayesian CRM \mathcal{D}_1 is coherent.

This fact indicates the model-based CRM is intrinsically compatible with the ethical principles in dose finding studies, although there is no apparent rule to enforce coherence. The proof of the method's coherence is provided in Section 5.4.

In contrast, the two-stage CRM \mathcal{D}_2 is not coherent in escalation in general; cf., the two-stage design in Figure 3.1. Consider another two-stage design that uses the same CRM model as in Figure 3.1, but with an initial design that escalates after groups of two nontoxic outcomes. Figure 5.1 shows a simulated trial of size 20 generated by this design. There is no indication of incoherent moves. However, to show that this design is indeed coherent, we need to verify that the design does not induce any incoherent move in all 2^{20} possible outcome sequences.

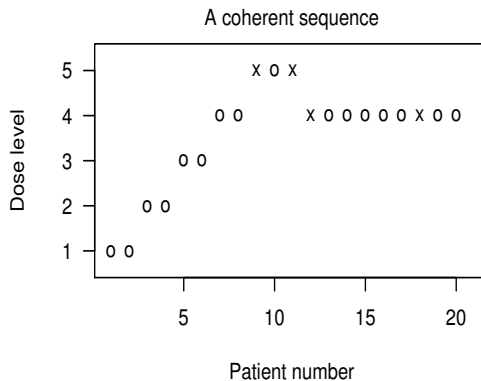


Figure 5.1 Simulated trials generated by a coherent two-stage CRM, without applying dose escalation restrictions. Each point represents a patient, with “o” indicating no toxicity and “x” indicating toxicity.

Theorem 5.1 (coherence condition). *Let $M_0 = \min\{i : Y_i = 1\}$ be the index for the first patient who experiences a toxicity. If $\mathcal{D}_2(H_{M_0+1}) \leq \mathcal{D}_2(H_{M_0})$ with probability one, then the two-stage CRM \mathcal{D}_2 is coherent.*

In words, Theorem 5.1 states that a two-stage CRM is coherent if it does not induce an incoherent escalation at the transition from the initial design (stage 1) to the model-based CRM (stage 2). Since there are only $N - 1$ possible ways that a transition may occur, this theorem reduces the computational needs for coherence checking significantly. For example, we can consider the outcome sequence

$$\mathbf{Y}^{(i)} = (Y_1^{(i)}, \dots, Y_N^{(i)})$$

where $Y_j^{(i)} = 0$ for $j \neq i$, and $Y_i^{(i)} = 1$, and generate the corresponding dose sequence

$$\mathbf{X}^{(i)} = (X_1^{(i)}, \dots, X_N^{(i)})$$

by a two-stage CRM \mathcal{D}_2 . Then, the design \mathcal{D}_2 is coherent if

$$X_{i+1}^{(i)} \leq X_i^{(i)} \text{ for } i = 1, \dots, N-1,$$

that is, no incoherent escalation is induced for patient $i + 1$.

The R function `cohere` in the package ‘`dfcrm`’ executes this algorithm and checks the coherence status of a two-stage CRM. The design used in Figure 5.1 is confirmed to be coherent using this function:

```
> ### Verify coherence of two-stage CRM in Figure 5.1
>
> p0 <- c(0.05,0.12,0.25,0.40,0.55)      # initial guesses
> theta <- 0.25                          # target DLT rate
> x0 <- c(1,1,2,2,3,3,4,4,rep(5,12))    # initial design
> foo <- cohere(p0,theta,x0,model="logistic")
> foo$message
[1] "Coherent"
>
```

Faries [36] notes the potential problem of incoherent escalation in the modified CRM, which starts at a dose below the prior MTD, and proposes enforcing coherence by *restrictions*. A natural question: what advantage checking the coherence condition in Theorem 5.1 offers over the simple use of dose escalation restrictions. Before answering this question, we will explore a related concept called compatibility next.

5.2.3 Compatibility

The initial design sequence $\{x_{i,0}\}$ in a two-stage CRM dictates the pace of escalation when there is no observed toxicity throughout the trial. Intuitively, should any toxic outcome occur, dose escalation may proceed at a rate slower than that prescribed by this initial sequence. With this in mind, the sequence $\{x_{i,0}\}$ should be set so as to reflect the “fastest” escalation pace for a target toxicity rate θ . (The notion of speed of escalation will be made specific in Chapter 10.)

Definition 5.1 (compatibility). An initial design $\{x_{i,0}\}$ is said to be compatible with the CRM in a two-stage CRM \mathcal{D}_2 in (3.10) with respect to θ if $\mathcal{D}_2(H_i) \leq x_{i,0}$ for all i with probability one.

Figure 5.2 shows an outcome sequence of 20 patients by a two-stage CRM with a group-of-four initial design. In the left panel, an escalation occurs for patient 15 after patient 14 has a toxic outcome. Furthermore, the toxic outcome of patient 14 seems to have speeded up the pace of escalation: patient 15 would have received dose level 4 if patient 11 did not experience toxicity. To mend the incoherence problem, one may simply adopt dose escalation restrictions as discussed in Section 3.5. The right panel of Figure 5.2 indicates no violation of coherence. However, the problem of incompatibility remains: patient 16 would have received dose 4 if no toxicity occurred prior to his.

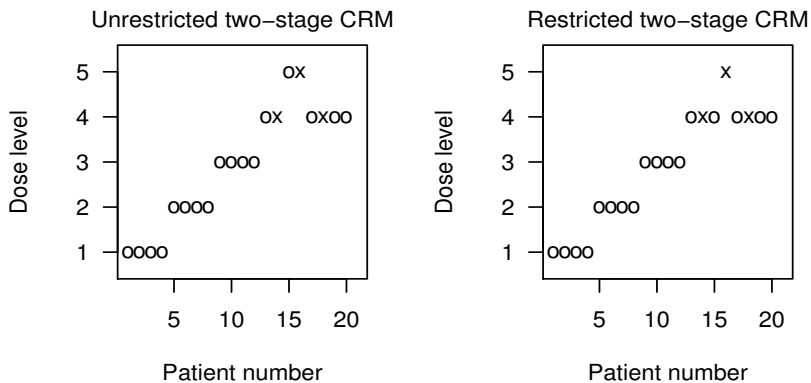


Figure 5.2 Simulated trials by an incompatible two-stage CRM. Dose escalation restrictions are applied in the right panel to enforce coherence, but not in the left. Each point represents a patient, with “o” indicating no toxicity and “x” indicating toxicity.

The property of compatibility is more subtle than coherence in that violation against compatibility does not cause any immediate ethical concern. However, it does cast doubt on the soundness of the overall design when the model-based CRM and the initial sequence give conflicting dose assignments. The fact that a one-stage CRM is always coherent whereas a two-stage CRM is not indicates that the problem may arise as a result of the improper choice of the initial design.

Theorem 5.2 (compatibility condition). Suppose that Conditions 4.1–4.3 hold. If the initial design $\{x_{i,0}\}$ is chosen such that an unrestricted two-stage CRM \mathcal{D}_2 is coherent, then $\{x_{i,0}\}$ is compatible with the CRM with respect to θ .

Theorem 5.2 establishes compatibility as a consequence of coherence, which in turn may be assessed by Theorem 5.1. In contrast, enforcing coherence by restrictions

does not guarantee compatibility; cf., Figure 5.2. It motivates us toward achieving our goal to calibrate the initial design so that the two-stage CRM is coherent, without applying any dose escalation restrictions.

A pragmatic approach may enforce coherence and compatibility by restriction. This is justified, in the sense that any statistical and model-based decision should be overruled if it contradicts sound clinical judgment. However, by imposing ad hoc restrictions on the CRM out of convenience, we may undermine the advantage a model-based design may offer. While model-based decisions should not replace sound clinical judgments, a model-based design should be calibrated so that it can best mimic clinical judgments in a systematic manner, without the need to introduce ad hoc rules. In this regard, the criteria of coherence and compatibility can serve as tools to evaluate whether an initial design is appropriate for a trial's objective when used with a particular CRM model. As a case in point, in practice, it is common to use a group-of-three initial design with a two-stage CRM. There is no justification for this choice apart from its resemblance to the traditional 3+3 algorithm. In fact, on the basis of coherence, we observe in Figure 3.1 that a group-of-three initial design is not appropriate for $\theta = 0.25$. In Chapter 10, we will further exploit the interdependence of coherence, the initial design, and the target θ as a tool to calibrate a two-stage CRM design.

5.2.4 Extensions

For the general enrollment situations with $m \geq 1$, we may view each group as an experimental unit and modify the coherence definitions in (5.1) and (5.2), respectively as

$$P_{\mathcal{D}}(U_j > 0 | \bar{Y}_{j-1} \geq \theta) = 0 \text{ and } P_{\mathcal{D}}(U_j < 0 | \bar{Y}_{j-1} \leq \theta) = 0$$

for all j , where U_j in this context denotes the dose increment from group $j - 1$ to j , and \bar{Y}_j is the observed proportion of toxicity in group j . In other words, a design is group-coherent in escalation if it escalates only when the observed toxicity rate in the most current cohort is strictly less than the target θ . Likewise, a design is group-coherent in deescalation if it deescalates only when the observed toxicity rate in the most current group is strictly greater than θ . In the special case when group size is equal to 1, the group-coherence conditions reduce to (5.1) and (5.2). Cheung [16] proves the group coherence of a one-stage CRM that assigns doses based on the posterior mode of β with a unimodal prior.

The concept of group coherence makes practical sense only if no information is available between when the first patient in a group arrives and when the last patient of the group is treated. This assumption does not reflect the clinical situations, where the toxicity outcome may occur at a random time during the observation window. In this case, a real-time extension of the coherence concept is needed. As the definition of coherence becomes complicated, however, it also becomes difficult to interpret. See Chapter 11 for a discussion on the real-time coherence of a specific extension of the CRM.

In a combined phase I/II trial, dose escalation decisions depend on the efficacy outcomes such as tumor shrinkage as well as toxicity outcomes of the patients. There

are various ways to extend the CRM to use both the toxicity and efficacy outcomes for dose finding, and a consensus is yet to be reached. However, the extension of coherence is quite straightforward. Namely, a deescalation is incoherent if the most current patient does not have a positive efficacy outcome. Section 13.3 will further explore the notion of bivariate coherence and its application.

5.3 Large-Sample Properties

5.3.1 Consistency and Indifference Interval

A primary estimation concern in a dose finding trial is whether the true MTD is correctly identified. If a design \mathcal{D} utilizes information properly, we expect that it will select the correct dose when the sample size becomes sufficiently large. Technically, we are concerned about strong consistency of the design; that is, whether $\mathcal{D}(H_N) = d_v$ eventually as $N \rightarrow \infty$ with probability one, where

$$v = \arg \min_k |p_k - \theta|.$$

Since $\mathcal{D}(H_N)$ denotes the dose assignment of patient N , consistency has an ethical connotation that as a trial keeps enrolling, all patients will eventually be treated at the MTD.

However, it can be difficult for a design to achieve consistency under all possible dose–toxicity scenarios, when there are other practical constraints such as coherence and discrete dosing. Alternatively, therefore, we may try to measure how close the eventual selected dose is to the target θ on the probability scale. A way to measure such “closeness” is by the *indifference interval* of a dose finding design.

Definition 5.2 (indifference interval). *The interval (θ_L, θ_U) , for $\theta_L < \theta < \theta_U$, is said to be an indifference interval of a dose finding design \mathcal{D} if for some $N > 0$,*

$$P_{\mathcal{D}} \{x_n \in \mathcal{I}_{\pi}(\theta_L, \theta_U) \text{ for all } n \geq N\} = 1, \quad (5.3)$$

where $\mathcal{I}_{\pi}(p_L, p_U) = \{x : p_L \leq \pi(x) \leq p_U\}$.

In other words, an indifference interval is an interval in which the toxicity probability of the selected dose will eventually fall. In other words, a design with an indifference interval (θ_L, θ_U) may fail to differentiate the true MTD from a dose nearby, as long as the toxicity probability of the dose falls within (θ_L, θ_U) .

Not every dose finding design has an indifference interval; Chapter 12 will give examples that do not have indifference intervals. A design is said to be δ -sensitive, if its indifference interval can be expressed as $\theta \pm \delta$ for some $\delta \in (0, \theta)$, that is, $\theta_L = \theta - \delta$ and $\theta_U = \theta + \delta$. If a design has an indifference interval (θ_L, θ_U) that is not symmetric about θ , it is δ -sensitive by defining $\delta = \max(\theta_U - \theta, \theta - \theta_L)$.

The CRM is δ -sensitive. Furthermore, it can be calibrated to achieve any δ level of sensitivity. The smaller the half-width δ of a design, the higher resolution the design will have asymptotically: In Figure 5.3, the design with $\delta = 0.02$ will yield consistent estimates for the true MTD under both curves (bottom panel), whereas the

design with $\delta = 0.05$ may fail to converge to the MTD under Curve 2 (top panel). This having been said, a half-width 0.05 is arguably sufficiently small for $\theta = 0.25$ in many practical situations. Again, under Curve 2 in Figure 5.3, the design with $\delta = 0.05$ may not select the MTD, but will eventually select a dose in the neighborhood. In addition, since indifference interval is an asymptotic concept, a smaller δ is not necessarily a better choice than a larger value in finite sample settings. We will return to the choice of δ in Chapter 8.

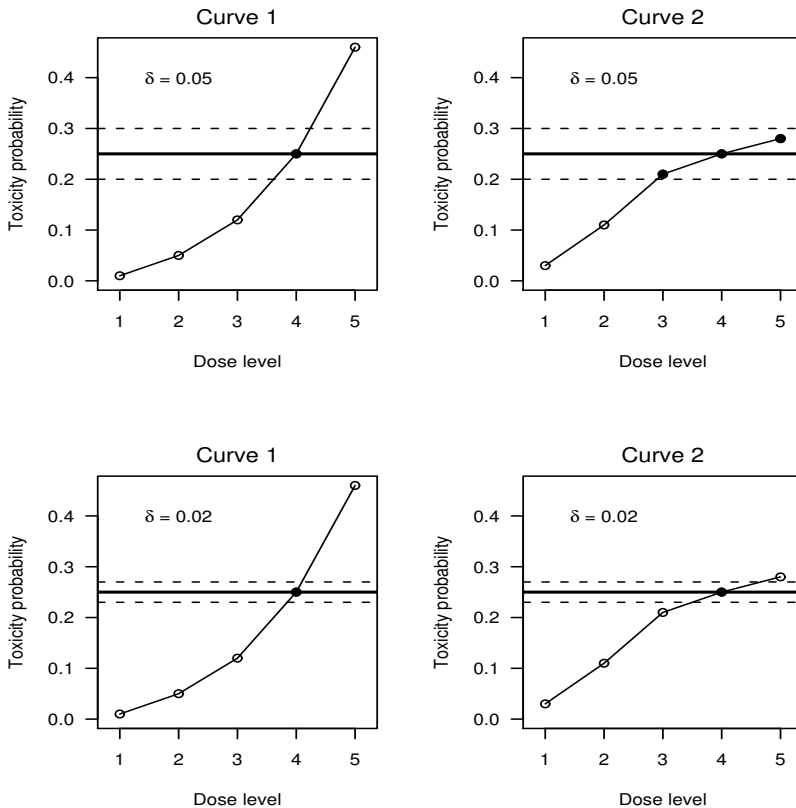


Figure 5.3 Two dose–toxicity curves under which dose level 4 is the true MTD with $\theta = 0.25$. Top: A design with an indifference interval $\theta \pm 0.05$ will eventually select the true MTD under Curve 1 (left), but may select a dose among levels 3, 4, and 5 under Curve 2 (right). Bottom: A design with an indifference interval $\theta \pm 0.02$ will yield consistent estimates for the MTD under both curves. The dark horizontal line in each plot indicates the target rate $\theta = 0.25$; the dotted lines indicate the indifference interval $\theta \pm \delta$.

5.3.2 Consistency Conditions of the CRM

Shen and O’Quigley [94] study the conditions under which the maximum likelihood CRM is strongly consistent. It is useful to bring up two points here. First, since we are dealing with an asymptotic property, the conditions do not depend on the initial design in a two-stage CRM. Second, in this chapter, we will establish the consistency of the Bayesian CRM under Conditions 4.1–4.8 stated in Section 4.3. The proof will be sketched in Section 5.4. As the initial design is irrelevant in the consideration of consistency, so is the prior distribution $G(\beta)$ on the model parameter β —as long as $G(\beta)$ is not a degenerated distribution.

5.3.2.1 Home Sets

Now recall that $F_k(\beta) = F(d_k, \beta)$ for $k = 1, \dots, K$. Assume without loss of generality that $F_k(\beta)$ is monotone decreasing in β for all k (Condition 4.2). Then, define for dose level j

$$B_j = \{\beta \in [\underline{b}, \bar{b}] : |F_j(\beta) - \theta| < |F_k(\beta) - \theta| \text{ for } k \neq j\}, \quad (5.4)$$

where \underline{b} and \bar{b} are constants so that $F_1(\underline{b}) > \theta$ and $F_K(\bar{b}) < \theta$, and that $[\underline{b}, \bar{b}]$ is a compact subset of the parameter space of β (Condition 4.5). We call B_j the *home set* for dose level j , because the CRM will select dose level j if $\hat{\beta}_n \in B_j$. Under these two conditions on the model $F_k(\beta)$, the home sets are mutually exclusive intervals defined by $B_1 = [\underline{b}, b_2)$, $B_j = (b_j, b_{j+1})$ for $j = 2, \dots, K-1$, and $B_K = (b_K, \bar{b}]$, where b_j solves

$$F_{j-1}(b_j) + F_j(b_j) = 2\theta \text{ for } j = 2, \dots, K. \quad (5.5)$$

Example 5.1. Consider the ψ -representation of the logistic model

$$F_k(\beta) = \frac{\exp[3 + \exp(\beta)\{\logit(p_{0k}) - 3\}]}{1 + \exp[3 + \exp(\beta)\{\logit(p_{0k}) - 3\}]} \quad (5.6)$$

with skeleton $p_{01} = 0.05$, $p_{02} = 0.12$, $p_{03} = 0.25$, $p_{04} = 0.40$, and $p_{05} = 0.55$. For a trial with target $\theta = 0.25$, if we set $\underline{b} = -5.000$ and $\bar{b} = 5.000$, solving (5.5) gives us

$$\begin{aligned} B_1 &= [-5.000, -0.280), B_2 = (-0.280, -0.093), B_3 = (-0.093, 0.097), \\ B_4 &= (0.097, 0.288), \text{ and } B_5 = (0.288, 5.00]. \end{aligned} \quad (5.7)$$

5.3.2.2 Least False Parameters

Next, define

$$\beta_j = F_j^{-1}(p_j) \text{ for } j = 1, \dots, K.$$

The parameter β_j may be called a *least false parameter* for dose level j , because it is the parameter value with which the model F_j at dose level j equals the true toxicity probability p_j . In the special case where all least false parameters are identical, that is, $\beta_j \equiv \beta_0$ for some β_0 , the model $F(x, \beta)$ is a correct specification of the dose–toxicity relationship and β_0 is the true parameter. In general, a true parameter value β_0 may not exist when the model is misspecified. However, the least false parameters always exist, because F_j^{-1} exists by monotonicity of F_j (Condition 4.4).

5.3.2.3 Main Result

In Chapter 4, Section 4.3 lists the regularity conditions (Conditions 4.1–4.8) for a CRM model. These assumptions are mild and verifiable. For the CRM to be consistent, we require another assumption:

Condition 5.1 (consistency condition). $\beta_j \in B_l$ for all j and some l .

Theorem 5.3 (consistency). *Suppose that the dose–toxicity model $F(x, \beta)$ satisfies the regularity conditions given in Section 4.3. If Condition 5.1 holds, the CRM will select dose d_l eventually with probability one, and $l = v$ is the true MTD.*

When $F(x, \beta)$ is a correct specification of the dose–toxicity curve, Condition 5.1 is satisfied. That is, the CRM is consistent when we use a correct model. However, in general, consistency of the CRM requires that the least false parameters β_j s are sufficiently close to each other so that they belong to the same home set.

Table 5.1 gives the least false parameter values β_j s for the logistic model (5.6) under four dose–toxicity curves, where $v = 4$ for a target $\theta = 0.25$. This model is a correct specification of Curve 0, indicated by the fact that $\beta_j \equiv 0.18$ for all j . As a consequence, Condition 5.1 is met with $l = 4$, and $0.18 \in B_4$; the home sets B_j s of model (5.6) are given in (5.7) in Example 5.1. Theorem 5.3 guarantees the CRM using model (5.6) will consistently select the true MTD under Curve 0. In general, for consistency to hold, Condition 5.1 requires only that $\beta_j \in B_4$: Model (5.6) satisfies this condition under Curve 1 in Table 5.1, but not under Curves 2 and 3.

Table 5.1 Least false parameters β_j s of logistic model (5.6) under four dose–toxicity curves

j	Curve 0		Curve 1		Curve 2		Curve 3	
	p_j	β_j	p_j	β_j	p_j	β_j	p_j	β_j
1	0.02	0.18	0.01	0.21	0.03	0.09	0.01	0.30
2	0.05	0.18	0.05	0.18	0.11	0.01	0.02	0.35
3	0.13	0.18	0.12	0.19	0.21	0.05	0.09	0.25
4	0.25	0.18	0.25	0.18	0.25	0.18	0.25	0.18
5	0.41	0.18	0.46	0.12	0.28	0.34	0.54	0.02

As illustrated in the above computations, the least false parameters of a model depend on the underlying unknown dose–toxicity curve. Therefore, unlike the other regularity conditions, Condition 5.1 cannot be directly used in practice to choose between models. We will shortly return to this issue in Section 5.3.3.

5.3.2.4 A Relaxed Condition

When the underlying dose–toxicity curve is steep, the MTD is very distinct from its neighboring doses on the probability scale; cf., Curve 3 in Table 5.1. Intuitively, then, a reasonable dose finding method will likely identify the MTD from the other doses under a steep dose–toxicity curve. In other words, if a CRM design yields consistent

MTD estimate under Curve 1 in Table 5.1, we expect that the same design will also be consistent under Curve 3. Cheung and Chappell [21] thus postulate the steepness condition as a sufficient condition for consistency:

Condition 5.2 (steepness condition).

$$\beta_k \in \bigcup_{j=k+1}^K B_j \text{ for } k < l, \beta_l \in B_l, \text{ and } \beta_k \in \bigcup_{j=1}^{k-1} B_j \text{ for } k > l.$$

Condition 5.1 is stronger than Condition 5.2, and implies the latter. Applying the steepness condition to model (5.6) in Example 5.1 for $l = 4$, we obtain

$$\begin{aligned} \beta_1 \in (-0.280, 5.000], \beta_2 \in (-0.093, 5.000], \beta_3 = (0.097, 5.000], \\ \beta_4 \in (0.097, 0.288), \text{ and } \beta_5 \in [-5.000, 0.288) \end{aligned} \quad (5.8)$$

From Table 5.1, we see that model (5.6) satisfies the steepness condition (5.8) under Curve 3, and hence anticipate that the CRM will converge to the true MTD under this dose–toxicity curve. Figure 5.4 plots the probability of correctly selecting dose 4 by a two-stage CRM using model (5.6) under Curves 1–3 in Table 5.1. As anticipated, since Condition 5.2 is met under Curve 3 (as well as Curve 1), the CRM selects the correct MTD increasingly often as sample size increases.




Figure 5.4 Probability of selecting the MTD (dose level 4) versus sample size by a two-stage CRM under dose–toxicity Curve 1 (solid), Curve 2 (dashed), and Curve 3 (dotted) in Table 5.1. The maximum likelihood CRM is used with model (5.6) and a group-of-two initial design.

Curve 2 is not steep enough for model (5.6) as the steepness condition (5.8) fails under Curve 2: $\beta_5 = 0.34 \notin [-5.000, 0.288)$. Figure 5.4 shows that the probability of selecting the MTD plateaus at about 40% under Curve 2 even as sample size

grows. Thus, Condition 5.2 successfully predicts the failure of consistency under this dose–toxicity scenario. More important, we observe that the failure of consistency is predictive of the CRM’s poor performance at small samples. It thus demonstrates the relevance of asymptotics to dose finding trials with small-to-moderate sample sizes.

5.3.3 Model Sensitivity of the CRM

Ideally, Conditions 5.1 and 5.2 would allow us to specify the model $F_k(\beta)$ so that the CRM is consistent. Unfortunately, both conditions require knowledge of the unknown dose–toxicity curve, and cannot be used in practice. Therefore, we take a reverse approach that enumerates the set of dose–toxicity curves under which a given model $F_k(\beta)$ will satisfy Condition 5.2 [21]. To illustrate, consider model (5.6) in Example 5.1. Condition 5.2 for this model with $v = 4$ is given in (5.8), and is equivalent to

$$\begin{aligned} p_1 \in [0.00, 0.18), p_2 \in (0.00, 0.18], p_3 = (0.00, 0.18], \\ p_4 \in (0.18, 0.32), \text{ and } p_5 \in (0.32, 0.95] \end{aligned} \quad (5.9)$$

by converting the intervals onto the probability scale via $p_k = F_k(\beta_k)$. In other words, the use of model (5.6) will lead to consistent MTD estimation under any dose–toxicity curve with $v = 4$, $p_1, p_2, p_3 < 0.18$ and $p_5 > 0.32$. This condition is illustrated in Figure 5.5d. Condition 5.2 is satisfied if the dose–toxicity curve lies completely in the shaded region. In particular, Curves 1 and 3 in Table 5.1 are embedded in the shaded region, whereas Curve 2 lies outside the region at dose level 3, suggesting this CRM design is not consistent under Curve 2.

As the true dose–toxicity curve is unknown, so is the true MTD v . We therefore need to repeat the calculation analogous to that of (5.9) for each possible v . Figure 5.5 plots the regions of dose–toxicity curves under which the CRM is consistent under each possible v . We observe that the regions may cover curves that are not monotone increasing, but require that the neighboring doses of the true v are quite distinct from v in terms of the toxicity probability. The quantity of such distinction may be used as a numerical measure for the sensitivity of model (5.6), as follows.

Suppose again that $v = 4$ with $p_4 \in (0.18, 0.32)$. If p_3 is close to $\theta = 0.25$, though not as close as p_4 , that is, $p_3 \in (0.18, p_4)$ such that $|p_3 - \theta| > |p_4 - \theta|$, then condition (5.9) is violated and the CRM may converge to dose level 3. By the same token, if $p_5 \in (p_4, 0.32]$, the CRM may eventually select dose level 5. In other words, the CRM using model (5.6) may converge to a dose with toxicity probability falling on the interval $(0.18, 0.32]$ if the true $v = 4$. The interval $(0.18, 0.32]$ is thus called the indifference interval of this CRM design given $v = 4$.

Generally, for a trial with K test doses, for a given model $F_k(\beta)$, the indifference interval (θ_L, θ_U) of the CRM can be computed as follows: let

$$(\theta_{L,v}, \theta_{U,v}) = \begin{cases} (0, F_2(b_2)) & \text{if } v = 1 \\ (F_{v-1}(b_v), F_{v+1}(b_{v+1})) & \text{if } v = 2, \dots, K-1 \\ (F_{K-1}(b_K), 1) & \text{if } v = K \end{cases}$$

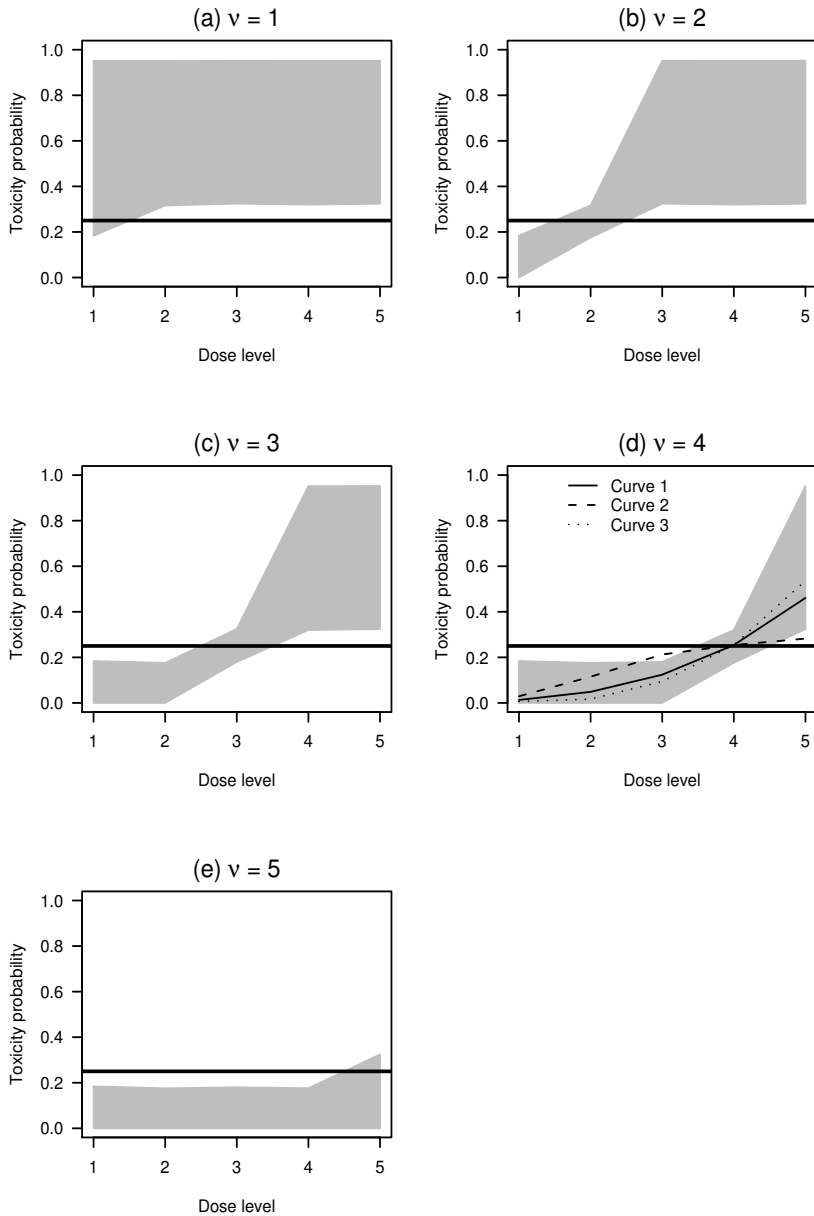


Figure 5.5 Condition 5.2 for logistic model (5.6) in Example 5.1 with $\theta = 0.25$. The model satisfies Condition 5.2 if the dose–toxicity curve lies completely in the shaded region.

where b_2, \dots, b_K solve (5.5). Then the overall indifference interval is given by $\theta_L = \min\{\theta_{L,v} : v = 2, \dots, K\}$ and $\theta_U = \max\{\theta_{U,v} : v = 1, \dots, K-1\}$.

Table 5.2a displays the indifference interval of model (5.6) for each v . From the table, we can deduce that this CRM model will eventually recommend a dose with toxicity probability somewhere between $(0.175, 0.325) = 0.250 \pm 0.075$. We note that the CRM may select a dose outside this range if all the test doses are either too low or too toxic. This is apparently an issue about how doses are chosen for experimentation, rather than an issue of model sensitivity.

Table 5.2 Indifference intervals of three CRM models for $\theta = 0.25$

v	(a) Model (5.6)		(b) Model (5.10)		(c) Model (5.17)	
	$\theta_{L,v}$	$\theta_{U,v}$	$\theta_{L,v}$	$\theta_{U,v}$	$\theta_{L,v}$	$\theta_{U,v}$
1	0.000	0.316	0.000	0.301	0.000	0.313
2	0.184	0.325	0.199	0.339	0.187	0.305
3	0.175	0.320	0.161	0.342	0.195	0.297
4	0.180	0.324	0.158	0.427	0.203	0.297
5	0.176	1.000	0.073	1.000	0.203	1.000
Overall	0.175	0.325	0.073	0.427	0.203	0.313

Consider another logistic CRM model:

$$F_k(\beta) = \frac{\exp[3 + \exp(\beta) \{\text{logit}(p'_{0k}) - 3\}]}{1 + \exp[3 + \exp(\beta) \{\text{logit}(p'_{0k}) - 3\}]} \quad (5.10)$$

with skeleton $p'_{01} = 0.05$, $p'_{02} = 0.10$, $p'_{03} = 0.25$, $p'_{04} = 0.45$, and $p'_{05} = 0.75$ for a trial with $\theta = 0.25$. Table 5.2b shows the indifference intervals of this model. In particular, if the true $v = 5$, the CRM using model (5.10) may fail to distinguish v from a dose with toxicity probability as low as 0.073. Also, this model generally has wider indifference intervals than model (5.6). These intervals may thus serve as the numerical measures for model sensitivity and the basis for comparison of various CRM models.

Another important point is that the model sensitivity in the CRM depends on the skeleton, generally to a greater extent than the functional form $F(\cdot, \beta)$. Models (5.6) and (5.10) have the same logistic form but their indifference intervals differ as a result of the choices of $\{p_{0k}\}$ and $\{p'_{0k}\}$. The concept of the indifference intervals will serve as the theoretical basis of how the skeleton is calibrated in Chapter 8.

5.3.4 Computing Model Sensitivity in R

Model sensitivity of the CRM is evaluated by the function `crmsens` in the ‘`dfcrm`’ package. The function requires input on the skeleton $\{p_{0k}\}$, the target θ , the dose–toxicity model (empiric or logistic), and an intercept for the logistic function when it applies. The following R code evaluates the sensitivity of model (5.6).

```
> ### Evaluate sensitivity of a CRM model
```

```

> p0 <- c(0.05, 0.12, 0.25, 0.40, 0.55)
> theta <- 0.25
> foo <- crmsens(p0, theta, model="logistic", intcpt=3)
> foo$H          # home sets
              [,1]      [,2]
[1,] -5.00000000 -0.28046846
[2,] -0.28046846 -0.09340689
[3,] -0.09340689  0.09723351
[4,]  0.09723351  0.28842988
[5,]  0.28842988  5.00000000
> foo$iint       # indifference intervals
              [,1]      [,2]
[1,]          NA  0.3161698
[2,]  0.1838308  0.3245127
[3,]  0.1754882  0.3201194
[4,]  0.1798811  0.3240478
[5,]  0.1759521          NA
>

```

5.4 Proofs

5.4.1 Coherence of One-Stage CRM

Coherence of the one-stage Bayesian CRM requires Conditions 4.1 and 4.2. In the following, assume without loss of generality that $F_k(\beta)$ is monotone decreasing in β for all k .

To establish coherence in escalation, we need to prove $x_{i+1} \leq x_i$ if $Y_i = 1$ for all i . Under Condition 4.1, it suffices to show that $Y_i = 1$ implies $\hat{\beta}_i \leq \hat{\beta}_{i-1}$ so that $F_k(\hat{\beta}_i) \geq F_k(\hat{\beta}_{i-1})$ for all k under Condition 4.2. We take the approach in Cheung [16], who proves coherence for an extended class of CRM. First,

$$\begin{aligned} \hat{\beta}_i - \hat{\beta}_{i-1} &= \frac{\int \beta L_i(\beta) dG(\beta)}{\int L_i(\beta) dG(\beta)} - \frac{\int \beta L_{i-1}(\beta) dG(\beta)}{\int L_{i-1}(\beta) dG(\beta)} \\ &= \frac{\int \beta L_i(\beta) dG(\beta) \int L_{i-1}(\beta) dG(\beta) - \int \beta L_{i-1}(\beta) dG(\beta) \int L_i(\beta) dG(\beta)}{\int L_i(\beta) dG(\beta) \int L_{i-1}(\beta) dG(\beta)}. \end{aligned}$$

Because $L_i(\beta) = F(x_i, \beta) L_{i-1}(\beta)$ for $Y_i = 1$, the numerator in the above expression can be rewritten as

$$A = \int \beta L_{i-1}(\beta) L_{i-1}(\phi) \{F(x_i, \beta) - F(x_i, \phi)\} dG(\beta) dG(\phi) \quad (5.11)$$

$$= \int \phi L_{i-1}(\phi) L_{i-1}(\beta) \{F(x_i, \phi) - F(x_i, \beta)\} dG(\phi) dG(\beta) \quad (5.12)$$

The expression (5.12) is obtained by symmetry. Adding (5.11) and (5.12) gives

$$2A = \int L_{i-1}(\beta) L_{i-1}(\phi) (\beta - \phi) \{F(x_i, \beta) - F(x_i, \phi)\} dG(\phi) dG(\beta) < 0$$

because $(\beta - \phi) \{F(x_i, \beta) - F(x_i, \phi)\} < 0$ under Condition 4.2. Hence, $\hat{\beta}_i - \hat{\beta}_{i-1} < 0$ when $Y_i = 1$.

Coherence of the one-stage Bayesian CRM is completed by analogous derivation for the case $Y_i = 0$.

5.4.2 Consistency of the CRM

Shen and O'Quigley [94] prove the consistency of the maximum likelihood CRM under Condition 5.1, in addition to Conditions 4.1, 4.2, and 4.4–4.7 in Chapter 4, Section 4.3. Readers are referred to their article for the details of the proof.

In this subsection, we sketch the proof of the consistency of the *Bayesian* CRM using posterior mean under an additional assumption that Condition 4.8 also holds. A key result in the proof of the maximum likelihood CRM is that the maximum likelihood estimate $\hat{\beta}_n$, given the observation history H_{n+1} , belongs to the open set B_V almost surely when the sample size n is sufficiently large. Here, we further note that $\hat{\beta}_n \in B_V$ for any possible history H_{n+1} without regard for the underlying dose assignment mechanism.

Let $l_n(\beta) = \log L_n(\beta)$ denote the log-likelihood of β with first and second derivatives, respectively, denoted by $l'_n(\beta)$ and $l''_n(\beta)$. Precisely,

$$l'_n(\beta) = \sum_{k=1}^K z_k \frac{F'_k(\beta)}{F_k(\beta)} + (n_k - z_k) \frac{-F'_k(\beta)}{1 - F_k(\beta)}$$

where z_k and n_k indicate the number of toxic outcomes and the number of patients treated at dose k . Since $l'_n(\tilde{\beta}_n) = 0$, we can apply Taylor's expansion and obtain

$$l_n(\beta) = l_n(\tilde{\beta}) + \frac{1}{2}(\beta - \tilde{\beta})^2 l''_n(\tilde{\beta}^*)$$

for some $\tilde{\beta}^*$ between β and $\tilde{\beta}$, where $l''_n(\beta) =$

$$\sum_{k=1}^K z_k \left[\frac{F_k(\beta)F''_k(\beta) - \{F'_k(\beta)\}^2}{\{F_k(\beta)\}^2} \right] + (n_k - z_k) \left[\frac{-F''_k(\beta) + F_k(\beta)F''_k(\beta) - \{F'_k(\beta)\}^2}{\{1 - F_k(\beta)\}^2} \right].$$

Under Condition 4.8, we have

$$\bar{l}''_n \equiv \sup_{\beta \in [\underline{b}, \bar{b}]} l''_n(\beta) \rightarrow -\infty \text{ as } n \rightarrow \infty \quad (5.13)$$

almost surely. Now, we express the posterior mean of β given history H_{n+1} as

$$\hat{\beta}_n = \frac{\int_{-\infty}^{\infty} \beta L_n(\beta) dG(\beta)}{\int_{-\infty}^{\infty} L_n(\beta) dG(\beta)} = \tilde{\beta}_n + \frac{\int_{-\infty}^{\infty} (\beta - \tilde{\beta}_n) L_n(\beta) dG(\beta)}{\int_{-\infty}^{\infty} L_n(\beta) dG(\beta)}. \quad (5.14)$$

Furthermore,

$$\left| \int_{-\infty}^{\infty} (\beta - \tilde{\beta}_n) L_n(\beta) dG(\beta) \right|$$

$$\begin{aligned}
&= \left| \int_{-\infty}^{\infty} (\beta - \tilde{\beta}_n) \exp \left\{ l_n(\tilde{\beta}_n) + \frac{1}{2}(\beta - \tilde{\beta}_n)^2 l_n''(\tilde{\beta}_n^*) \right\} dG(\beta) \right| \\
&\leq L_n(\tilde{\beta}_n) \int_{-\infty}^{\infty} |\beta - \tilde{\beta}_n| \exp \left\{ \frac{1}{2}(\beta - \tilde{\beta}_n)^2 T_n'' \right\} dG(\beta)
\end{aligned}$$

The above integral can be further broken into

$$\int_{\tilde{\beta}_n - \varepsilon}^{\tilde{\beta}_n + \varepsilon} |\beta - \tilde{\beta}_n| \exp \left\{ \frac{1}{2}(\beta - \tilde{\beta}_n)^2 T_n'' \right\} dG(\beta) \leq \varepsilon \quad (5.15)$$

for a given $\varepsilon > 0$, and

$$\left(\int_{-\infty}^{\tilde{\beta}_n - \varepsilon} + \int_{\tilde{\beta}_n + \varepsilon}^{\infty} \right) |\beta - \tilde{\beta}_n| \exp \left\{ \frac{1}{2}(\beta - \tilde{\beta}_n)^2 T_n'' \right\} dG(\beta) \rightarrow 0 \quad (5.16)$$

as $n \rightarrow \infty$ by (5.13). Since we can choose ε in (5.15) to be arbitrarily small independently of n , the second term in (5.14) is arbitrarily close to 0 as sample size increases. We have thus shown

$$|\hat{\beta}_n - \tilde{\beta}_n| \rightarrow 0$$

almost surely. Since $\tilde{\beta}_n \in B_V$ eventually with probability one, we can also draw the same conclusion for $\hat{\beta}_n$ because B_V is an open set.

5.5 Exercises and Further Results

Exercise 5.1. Using Definition (5.4) of the home sets, convince yourself that $\hat{\beta}_n \in B_k$ implies the Bayesian CRM will recommend dose k for the next patient.

Exercise 5.2. Verify that (5.7) in Example 5.1 solves (5.5).

Exercise 5.3. Verify the computation of the indifference intervals in Table 5.2c for the following logistic CRM model for a trial with $\theta = 0.25$:

$$F_k(\beta) = \frac{\exp [3 + \exp(\beta) \{ \text{logit}(p_{0k}'') - 3 \}]}{1 + \exp [3 + \exp(\beta) \{ \text{logit}(p_{0k}'') - 3 \}]} \quad (5.17)$$

with skeleton $p_{01}'' = 0.07$, $p_{02}'' = 0.15$, $p_{03}'' = 0.25$, $p_{04}'' = 0.35$, and $p_{05}'' = 0.45$.

Exercise 5.4. Using similar arguments developed in Section 5.4, prove that the Bayesian CRM with the posterior mode is coherent (cf. Exercise 4.4).

Empirical Properties

6.1 Introduction

The two main operating characteristics of a dose finding design are given by the distribution of the selected dose and the in-trial dose allocation. The former reflects accuracy of the design, and the latter safety. This chapter introduces some summary indices for accuracy and safety in Section 6.2; these indices will serve as the basis of design comparison. In Section 6.3, we present a nonparametric optimal benchmark for a method's performance.

6.2 Operating Characteristics

6.2.1 Accuracy Index

Table 6.1 compares the distribution of the selected dose of a maximum likelihood two-stage CRM using models (5.6), (5.10), and (5.17) for a trial with $\theta = 0.25$. The sample size is $N = 20$, and four dose-toxicity scenarios are considered. The probability of correct selection (PCS) of the true MTD is indicated in bold in the table. The PCS is the most immediate index for accuracy, which can be used to compare different designs. For example, the design using model (5.10) has the best PCS under Curves 1–3 in Table 6.1, but worst under Curve 4. In addition, the PCS behaves in line with asymptotic reasoning: the accuracy improves as the dose-toxicity curve becomes steep; the poor performance of model (5.10) under Curve 4 may be explained by its violation of Condition 5.2. For a similar asymptotic reason, all three models have low PCS under Curve 2.

On the other hand, the entire distribution of selected dose does provide more detailed information than what the PCS alone suggests. For example, it is of ethical importance to look at the probability of selecting an overly toxic dose; for example, dose level 5 under Curves 1 and 3 in Table 6.1. This having been said, for design comparison purposes, it is useful to have a single numerical index to summarize the entire distribution and reflects a design's accuracy. In particular, we will use the following accuracy index:

$$\mathcal{A}_N = 1 - K \times \frac{\sum_{k=1}^K \rho_k \times \text{Probability of selecting dose } k}{\sum_{k=1}^K \rho_k}. \quad (6.1)$$

where ρ_k is a discrepancy measure between the true toxicity probability p_k at dose

Table 6.1 Distribution of the selected dose and the accuracy indices of three two-stage CRM for a trial with a target $\theta = 0.25$ and $N = 20$

Model	Probability of selecting dose					Accuracy index \mathcal{A}_{20}			
	1	2	3	4	5	sq	abs	0-1	od ^a
Curve 1:	.01	.05	.12	.25	.46				
(5.6)	.00	.02	.23	.55	.19	.58	.51	.44	.28
(5.10)	.00	.02	.24	.63	.11	.69	.61	.53	.51
(5.17)	.00	.02	.24	.51	.24	.52	.45	.38	.16
Curve 2:	.03	.11	.21	.25	.28				
(5.6)	.01	.13	.33	.26	.28	.75	.52	.07	.34
(5.10)	.01	.13	.36	.33	.17	.75	.54	.16	.44
(5.17)	.01	.12	.27	.25	.37	.77	.54	.06	.27
Curve 3:	.01	.02	.09	.25	.54				
(5.6)	.00	.00	.20	.63	.16	.58	.56	.54	.36
(5.10)	.00	.00	.17	.72	.11	.70	.67	.65	.55
(5.17)	.00	.01	.21	.59	.20	.51	.49	.48	.25
Curve 4:	.03	.05	.08	.12	.25				
(5.6)	.00	.01	.07	.25	.66	.74	.66	.57	.66
(5.10)	.00	.01	.09	.41	.48	.62	.50	.35	.50
(5.17)	.00	.01	.05	.20	.74	.81	.75	.68	.75

Note: Numbers associated with the MTD are in bold.

^aThe overdose error is used with $\alpha = 0.20$.

level k and the target θ . The accuracy index (6.1) summarizes the distribution of the selected dose through its weighted average with weights $\rho_j / \sum_k \rho_k$. Generally, a large index indicates high accuracy, and the maximum value of the index is 1. Suppose we make a blind guess in the absence of any observation, that is, $N = 0$, by randomly selecting one of the K test doses with equal probability $1/K$. It can then be shown that $\mathcal{A}_0 = 0$. Therefore, the index (6.1) measures the information gained due to the use of the observations. Finally, the index \mathcal{A}_N may take on a negative value; that is, improper use of data could do more harm than good.

Among many possible choices of ρ_k , some possible discrepancy measures are:

1. Squared discrepancy (sq): $\rho_k = (p_k - \theta)^2$.
2. Absolute discrepancy (abs): $\rho_k = |p_k - \theta|$.
3. Classification error (0-1): $\rho_k = I(k \neq v)$, where $I(\cdot)$ is an indicator function.
4. Overdose error (od): $\rho_k = \alpha(\theta - p_k)^+ + (1 - \alpha)(p_k - \theta)^+$ for some $\alpha \leq 0.50$.

The discrepancy measure ρ_k is a characterization of the true dose-toxicity curve. Precisely, the denominator $\sum_k \rho_k$ in (6.1) measures the total deviation of the test doses from the target θ . A small value of this sum corresponds to less serious error one may commit under the dose-toxicity curve. Consider Curve 3 in Table 6.1. If we use the absolute discrepancy error, then $\rho_1 = 0.24, \rho_2 = 0.23, \rho_3 = 0.16, \rho_4 = 0.00,$

$\rho_5 = 0.29$ and $\sum_k \rho_k = 0.92$. In comparison, Curve 1, which is flatter around the MTD than Curve 3, has $\sum_k \rho_k = 0.78$; Curve 2, the flattest of all three, has $\sum_k \rho_k = 0.43$. Intuitively, it is easy to see that selecting dose level 5 under Curve 2 (which has a 0.28 toxicity probability) is not as serious a mistake as selecting dose level 5 under Curve 3 (which has a 0.54 toxicity probability).

Table 6.1 gives the accuracy index \mathcal{A}_{20} of the three CRM designs using various discrepancy measures. Most of the time, using different ρ_k does not change the relative rankings of the designs. For example, the CRM with model (5.10) is best for Curves 1 and 3 by all discrepancy measures, and model (5.17) is best for Curve 4. On the other hand, under Curve 2, model (5.10) is compared favorably to the other models when the classification error or overdose error ($\alpha = 0.20$) is used. As mentioned above, selecting dose 5 under Curve 2 is not a serious mistake. However, since there is a one-to-one correspondence between the classification error and the PCS, this discrepancy measure ignores the insignificance of the error and penalizes models (5.6) and (5.17) relatively severely for choosing dose 5 more often than model (5.10). Similarly, using the overdose error with $\alpha = 0.20$ overpenalizes models (5.6) and (5.17) for the same mistake. In contrast, the absolute error appears quite sensible: all three models get similar accuracy scores under Curve 2, while the frequent selection of toxic dose by models (5.6) and (5.17) is appropriately penalized under Curves 1 and 3. The squared error depicts a similar picture.

6.2.2 Overdose Number

The overdose (OD) number, defined as the average number of patients treated at a dose above the true MTD per trial, is a safety index of a dose finding design. The in-trial dose allocation by three CRM designs used in Table 6.1 is shown in Table 6.2. Under Curves 1–3, the OD number is the average number of assignments to dose level 5; under Curve 4, the OD number is invariably zero because all doses are safe. As a benchmark, a balanced randomization scheme on average assigns 4 patients to each dose for a trial with $K = 5$ dose levels and sample size $N = 20$. In other words, we expect an OD number of 4 patients under Curves 1–3 by balanced randomization. The CRM using models (5.6) and (5.10) improve safety by having an OD number lower than 4 under Curves 1–3. In contrast, using model (5.17) may raise safety concerns as it could lead to a OD number higher than balanced randomization; this model may not be used on this ground.

6.2.3 Average Toxicity Number

Another summary index for safety is the average toxicity number (ATN) per trial, which is roughly equal to

$$\sum_{k=1}^K p_k \times \text{Number of assignments at dose } k.$$

A large ATN certainly causes safety concerns. On the other hand, a low ATN is not necessarily an accurate indication of safety. For example, under Curve 4 where the

Table 6.2 In-trial dose allocation and the average toxicity number (ATN) of the three designs used in Table 6.1

Model	Number allocated to dose					ATN
	1	2	3	4	5	
Curve 1:	.01	.05	.12	.25	.46	
(5.6)	2.3	3.0	5.0	6.5	3.3	4.0
(5.10)	2.2	2.8	5.7	7.2	2.1	3.7
(5.17)	2.3	2.7	4.7	6.2	4.2	4.3
Curve 2:	.03	.11	.21	.25	.28	
(5.6)	2.9	4.6	5.4	3.9	3.2	3.7
(5.10)	2.8	4.5	6.3	4.2	2.2	3.6
(5.17)	2.9	4.1	4.8	3.8	4.4	3.8
Curve 3:	.01	.02	.09	.25	.54	
(5.6)	2.1	2.4	4.7	7.5	3.3	4.2
(5.10)	2.1	2.3	5.0	8.3	2.3	3.8
(5.17)	2.1	2.3	4.5	7.0	4.0	4.4
Curve 4:	.03	.05	.08	.12	.25	
(5.6)	2.4	2.8	3.3	4.5	7.0	2.8
(5.10)	2.4	2.7	3.9	5.7	5.3	2.6
(5.17)	2.4	2.5	2.9	4.1	8.1	3.0

Note: Numbers associated with OD are in bold.

highest dose is the MTD, using model (5.17) causes the largest ATN per trial, but this simply reflects that it treats more patients at the MTD than the other two models. Therefore, a high ATN is desirable in this scenario. To set up a benchmark, consider the ideal situation where all patients are being treated at the true MTD. We would expect $N\theta$ toxic outcomes per trial. Therefore, a design that yields roughly $N\theta$ toxic outcomes per trial can be considered safe, as long as the OD number is kept low relative to balanced randomization.

The OD number and the ATN extract complementary safety information from the in-trial allocation of a design. In Table 6.3, Design A is the ideal design with all patients treated at the true MTD. If we use ATN as the only safety criterion, design B will be declared as safe as design A although it treats half the patients at an overdose. In contrast, if we use OD number only, design C appears safe while it is indeed overconservative by treating the majority at an underdose; such conservatism is indicated by a low ATN.

6.3 A Nonparametric Optimal Benchmark

In a trial, each patient is given a dose, with the toxicity outcome observed only at that dose. Only a *partial* toxicity profile is observable. For example, suppose a patient receives dose level 3 in a trial and has a nontoxic outcome. We can then infer by monotonicity that he would have had a nontoxic outcome at doses 1 and 2. However,

Table 6.3 In-trial allocation of three generic designs for a trial with $\theta = 0.25$, $K = 3$, and $N = 20$ under a curve with $p_1 = 0.01$, $p_2 = 0.25$, and $p_3 = 0.50$

Design	Number allocated to dose			ATN
	1	2	3	
p_k :	.01	.25	.50	
A	0	20	0	5
B	10	0	10	5
C	15	5	0	1.4

Note: Numbers associated with OD are in bold.

we will have no information whether the patient would have suffered a toxic outcome had he been given a dose that is above dose level 3.

In a computer-simulated trial, however, by generating the toxicity tolerance for each patient, we can observe the toxicity outcome at every test dose, that is, a *complete* toxicity profile. Table 6.4 gives the complete toxicity profiles of 20 simulated patients under Curve 1 in Table 6.1. Recall that u_i indicates the toxicity tolerance of patient i , so that the patient will have a toxic outcome if treated at dose k where $u_i \leq p_k$. Take patient 9, for example. He will have a nontoxic outcome at doses 1, 2, and 3 because $u_9 = 0.175 > p_k$ for $k = 1, 2, 3$, and a toxic outcome at doses 4 and 5 because $u_9 < p_4, p_5$.

Since the toxicity tolerance u_i is not observable in reality, the complete toxicity information in Table 6.4 is generally not available from a trial. However, since we do have the complete toxicity profiles in a simulated trial, we can estimate the toxicity probability p_k by the sample proportion \hat{p}_k of toxicity among the simulated patients at each dose, and select the MTD in accordance with these estimates. The bottom row in Table 6.4 gives the sample proportions of this simulated trial; based on these, we may choose dose level 4 as the MTD for a trial with a target $\theta = 0.25$ because \hat{p}_4 is closest to θ . (Here, we adopt the conservative convention that when more than one dose is equally close to θ , the largest dose with $\hat{p}_k \leq \theta$ will be selected.) In addition, this MTD estimate is optimal in the sense that \hat{p}_k is unbiased for p_k and achieves the Cramér-Rao lower bound. This design may thus be called a nonparametric optimal design.

While the nonparametric optimal design cannot be implemented in a real trial where complete toxicity profiles are not available, its performance can be evaluated over many simulation runs and be used as a benchmark for efficiency. In particular, the accuracy indices \mathcal{A}_{20} of this optimal benchmark using the absolute discrepancy are 0.63, 0.45, 0.77, and 0.81, respectively, under Curves 1–4 in Table 6.1. Thus, the optimal design is superior to the CRM under Curves 1, 3, and 4. This is expected because the former uses information that is not available to the CRM. On the other hand, some of the CRM designs get quite close to the optimal benchmark, and it is quite interesting to note that the CRM outperforms the nonparametric optimal design under Curve 2. While it appears to be impossible at first, this may be explained by the fact that the CRM utilizes parametric dose–toxicity assumption that is not used by

Table 6.4 The complete toxicity profile of 20 simulated patients under a dose–toxicity curve. A “0” indicates a nontoxic outcome at a dose, and “1” a toxic outcome

i	u_i	Toxicity probability p_k				
		.01	.05	.12	.25	.46
1	.571	0	0	0	0	0
2	.642	0	0	0	0	0
3	.466	0	0	0	0	0
4	.870	0	0	0	0	0
5	.634	0	0	0	0	0
6	.390	0	0	0	0	1
7	.524	0	0	0	0	0
8	.773	0	0	0	0	0
9	.175	0	0	0	1	1
10	.627	0	0	0	0	0
11	.321	0	0	0	0	1
12	.099	0	0	1	1	1
13	.383	0	0	0	0	1
14	.995	0	0	0	0	0
15	.628	0	0	0	0	0
16	.346	0	0	0	0	1
17	.919	0	0	0	0	0
18	.022	0	1	1	1	1
19	.647	0	0	0	0	0
20	.469	0	0	0	0	0
Sample proportion \hat{p}_k		.00	.05	.10	.15	.35

the *nonparametric* optimal design. It indicates the potential advantage of the CRM that allows borrowing strength from information across doses via a parametric dose–toxicity model. O’Quiley et al. [77] give additional numerical evidence that the CRM is indeed near the optimal benchmark.

6.4 Exercises and Further Results

Exercise 6.1. Prove or verify the following statements:

- There is a one-to-one correspondence between the PCS and the accuracy index using the classification error.
- The absolute discrepancy measure is a special case of the overdose error.

Part II

Design Calibration

Specifications of a CRM Design

7.1 Introduction

A CRM design is completely specified by a set of design parameters. The first is the set of *clinical parameters*. These parameters are specified in accordance with the clinical objective, experimental setup, and practical constraints of the trial. Clinical parameters usually require inputs from clinical investigators, and may also be called *clinician-input parameters*. Clinical parameters include the target toxicity probability θ , the number of test doses K , sample size N , the prior MTD v_0 , and the starting dose of the trial x_1 .

The statistical component of a CRM design is specified by the *model parameters*. These parameters directly specify the dose–toxicity model used in the CRM. Model parameters include the functional form of the dose–toxicity function $F(\cdot, \beta)$, the initial guesses p_{0k} s of toxicity probabilities associated with the test doses, and the prior distribution $G(\beta)$ of the model parameter β . In case a two-stage CRM is warranted, namely, when $x_1 < d_{v_0}$, one will also need to specify a dose sequence $\{x_{i,0}\}$ as the *initial design*.

This division is not the only way to classify the CRM design parameters nor is it to be used unalterably. While some design parameters clearly require clinician inputs (e.g., θ), other clinical parameters (e.g., N and v_0) may be specified with respect to the statistical properties of the design. Also, some model parameters may first appear to require clinician inputs (e.g., p_{0k} s). However, this classification of CRM parameters provides a pragmatic system under which the statistical component is defined subsequent to the clinical parameters, so that the CRM can be tuned to obtain good operating characteristics with respect to the clinical objective. The focus of the following chapters (and this book) is the specification of the statistical component, although brief guidelines on the choice of the clinical parameters will be provided in Section 7.2 of this chapter. Section 7.3 prescribes a roadmap for calibrating the model parameters and the initial design, and outlines the concrete steps that are to be elaborated in the next three chapters. Section 7.4 presents two real trial examples and describes a trial-and-error approach used to choose the design parameters in these trials.

7.2 Specifying the Clinical Parameters

7.2.1 Target Rate θ

For chemotherapy trials in cancer patients, a DLT is typically defined as a grade 3 or higher toxicity according to the National Cancer Institute Common Terminology Criteria for Adverse Events, and the convention for the target toxicity rate θ lies somewhere between 0.20 and 0.30. While the rationale for this convention is not entirely clear apart from its rough agreement with the 3+3 algorithm, a target value outside this range should be adequately justified.

For other disease areas, the target rate θ varies with the severity of the adverse effects and the potential benefits to that patient population; see Section 7.4.2 for an example. Generally, it is useful to direct the clinical investigators to the literature or their clinical experience by asking questions such as “what are the potential adverse effects of new drug?”, “are these events expected to occur more frequently with a higher dose?”, “what is the toxicity rate of the standard treatment when given in a comparable patient population?” and “will a higher toxicity rate be acceptable with the experimental treatment than with the standard treatment?”

7.2.2 Number of Test Doses K

The choice of test doses involves the determination of dose increments. In the many situations where prior information is available to identify the lowest and highest dose regimens, the specified dose increments will in turn determine the number K of test doses for the trial.

There has been a perception that the dose increments are set based on the modified Fibonacci sequence defined by the diminishing multiples 2.00, 1.67, 1.50, 1.33, 1.33, . . . , although in practice the sequence is not generally used [88]. In fact, there is no biological basis for the use of the Fibonacci sequence. Rather, the magnitude of increment is closely tied to the nature of the treatment. In cancer vaccine dose finding studies, each subsequent dose increases exponentially in terms of plaque-forming units [51]. In drug trials, dose increments are relatively constant. Also, pharmacokinetic data in a few pilot patients may be used to assess whether two doses are pharmacologically separated. While the plasma concentration level tends to increase with dose, there is considerable variation in the individual values. If two successive doses have comparable range of plasma concentration, they will likely demonstrate similar pharmacological properties, and may be too close to be differentiated in terms of therapeutic benefits. Thus, testing one of the two doses, or the average of both, may prove sufficient.

7.2.3 Sample Size N

In general, we expect the required sample size N to increase as the number of test doses increases. When planning a trial, we may make three pragmatic assumptions and considerations:

- The CRM will allocate the majority of the subjects roughly evenly between the MTD and a neighbor dose.
- The other $K - 2$ doses will receive about 3 subjects each.
- The number of subjects treated at the MTD should be greater than $1/\theta$ so that we will expect to observe at least one toxic observation at the MTD.

Based on these considerations, we suggest that the clinical parameters ought to be constrained so that

$$\frac{N - 3(K - 2)}{2} > \frac{1}{\theta}. \quad (7.1)$$

For example, if $N = 20$ subjects are available for a trial with a target $\theta = 0.25$, the inequality prescribes $K < 6$ and suggests that the resources are not sufficient to test 6 or more dose levels. The criterion (7.1) has been shown to be a reasonable rule of thumb based on extensive simulations and experience with the numerical performance of the one-stage Bayesian CRM.

Alternatively, it may be natural to view the criterion (7.1) as a constraint on the sample size N for given θ and K . For example, for a trial with $\theta = 0.25$ and $K = 5$ test doses, the inequality (7.1) implies $N > 17$, that is, the sample size is at least 18. This inequality is by no means sharp, and should not be used as a formal justification of the sample size. It may at first seem unusual to consider N as a clinical parameter rather than a statistical parameter. In phase I trial practice, however, there is often a rough idea of the feasible range of sample size based on patient and funding availability. Therefore, it appears to be a reasonable strategy to start planning with a sample size N along with θ and K based on the clinical investigators' inputs and use (7.1) to screen out underresourced studies. In this regard, the inequality (7.1) may be used as a sample size *constraint* instead of a sample size *formula*. A similar sample size constraint for the two-stage CRM can be found in Chapter 10.

After the model parameters and the initial design are specified, one could (and should) reevaluate the required sample size based on the statistical properties of the design by simulations (See [Figure 7.1](#)).

7.2.4 Prior MTD v_0 and Starting Dose x_1

The prior MTD v_0 is to be distinguished from the starting dose level of the trial, that is, the dose level given to the first patient enrolled in the trial. The choice of the starting dose x_1 is a purely clinical decision based on the current state of knowledge about the experimental drug. Ideally, the prior MTD v_0 is equal to the starting dose level, as originally prescribed by the CRM. However, it is not uncommon that the clinical investigators have reasons to believe that the MTD is among the higher test doses, and are not comfortable to start the trial at any dose that is higher than the conventional use of the drug in other patient populations, which often is the lowest dose d_1 . Should this be the case, a two-stage CRM will be an appropriate dose-escalation strategy. A practical recommendation is to set v_0 at the highest dose level

that is less than or equal to the median dose level. That is, set

$$v_0 = \begin{cases} K/2 & \text{for even } K \\ (K+1)/2 & \text{for odd } K \end{cases} \quad (7.2)$$

The convention (7.2) is based on the consideration of statistical properties of the CRM design; this point will be elaborated in Chapter 9. According to this viewpoint, the prior MTD v_0 is considered a model parameter rather than a clinical input, and is tuned to obtain good operating characteristics.

7.3 A Roadmap for Choosing the Statistical Component

As indicated by the number and dimension of the design parameters, specifying a CRM design for a dose finding study is a complex process. Also, since the method's performance depends on how these parameters are chosen, this is a crucial process for the success of a CRM application. This section outlines a calibration approach to obtain "good" choices for the statistical component, composed of the following design parameters:

Functional form of dose–toxicity curve. The first component of a CRM model is the functional form of the dose–toxicity curve $F(\cdot, \beta)$. Among the many possible choices of F , the two most commonly used functions in the CRM literature are the empiric function (3.2) and the logistic function (3.3) with a fixed intercept $a_0 = 3$. As the consistency and coherence properties of the CRM do not depend on the particular choice of F as long as they satisfy the regularity conditions in Section 4.3, the use of the empiric or the logistic functions shall prove sufficient in many situations. In the following chapters, we will see calibration techniques and results using the logistic function with intercept $a_0 = 3$. This is to complement the calibration results available for the empiric function in the literature [60].

Initial guesses of toxicity probabilities (skeleton). The second component of a CRM model is the set of labels d_{ks} associated with the test doses. Choosing d_{ks} in turn is equivalent to specifying the initial guesses p_{0ks} of toxicity probabilities associated with the doses, the so-called skeleton. The choice of the skeleton has a large impact on the operating characteristics of the CRM. Ideally, these initial guesses should be chosen as clinical parameters so as to reflect the clinician's beliefs regarding the test doses. Such information is seldom available in practice. Chapter 8 will present a pragmatic approach to determining a set of p_{0ks} that achieve a prespecified δ level of sensitivity (see Chapter 5, Section 5.3.1). This approach simplifies the calibration process as it reduces the number of model parameters from K to 1. As δ -sensitivity is an asymptotic concept, the choice of δ in finite sample settings will also be discussed.

Prior distribution. When the Bayesian CRM is used, we need to choose a prior distribution for the model parameter β . Since we will consider only the empiric (3.2) and the logistic (3.3) functions, for which the parameter is free to take any values on the real line, we can focus on any prior distribution that takes support on the real line. Specifically, Chapter 9 will illustrate the calibration techniques with a normal prior,

that is,

$$\beta \sim N(\hat{\beta}_0, \sigma_\beta).$$

In light of the fact that the theoretical (asymptotic) properties do not depend on the particular form of the prior distribution, there are two practical advantages for the normality choice. First, posterior computations using Gauss–Hermite quadrature [72] under parameterization (3.2) and (3.3) are accurate when a normal prior is used. This is partly because of good normality approximation of the posterior distribution even when the sample size is small. Second, the Bayesian CRM using model (4.1) is invariant to the mean of a prior that constitutes a location-scale family. This property allows us to set

$$\hat{\beta}_0 = 0$$

without loss of generality and further simplify the calibration process: the prior is completely specified by a single number, that is, the prior standard deviation σ_β .

Initial design. A two-stage CRM is an appropriate design strategy when the actual starting dose x_1 is lower than the model-based prior MTD d_{V_0} . In that case, we need to specify a dose sequence $\{x_{i,0}\}$ as the initial design. Chapter 10 will present some general characterization of dose sequences and discuss the calibration of $\{x_{i,0}\}$. In particular, while a two-stage CRM is motivated by the inclination to be conservative, we will see that an overconservative initial design is incompatible with the clinical objective θ . The initial design calibration will thus be guided by the compatibility criterion (see Definition 5.1).

Many of the model choices made above, such as the use of logistic F and a normal prior, are motivated by the practical needs for simplicity. The purpose is to streamline the CRM calibration process so that the study statistician will be able to come up with a reasonable design for a dose finding study in a *timely and reproducible* fashion. Figure 7.1 presents a diagrammatic scheme for the reduced specification problem of the following design parameters:

- Half-width δ of the indifference interval (Chapter 8)
- Standard deviation σ_β of the normal prior distribution (Chapter 9)
- Initial dose sequence $x_{1,0}, x_{2,0}, \dots, x_{N,0}$ (Chapter 10)

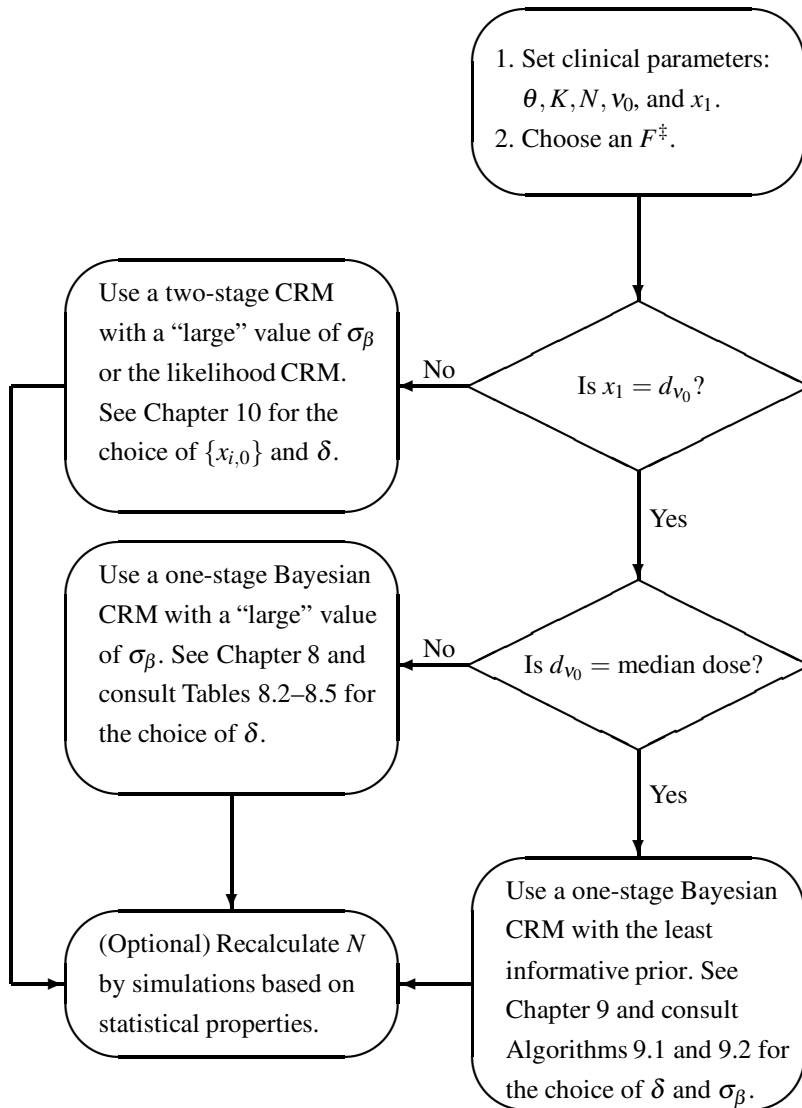
Each of these will be dealt with in detail in the next three chapters.

7.4 The Trial-and-Error Approach: Two Case Studies

7.4.1 The Bortezomib Trial

The bortezomib trial introduced in Chapter 1 evaluated $K = 5$ test dose schedules of bortezomib when administered in combination with the standard chemotherapy in lymphoma patients. The study used the standard definition of DLT as the study endpoint, and the trial objective was to identify a dose associated with a $\theta = 0.25$ toxicity probability.

The study design was a one-stage Bayesian time-to-event CRM, a variant of the CRM. In this section, as we focus on the calibration of the dose–toxicity model,



[‡]The subsequent chapters will use the logistic function (3.3) with $a_0 = 3$.

Figure 7.1 A flowchart for the specification of the model parameters and the initial design in the CRM.

we will disregard the time-to-event component, for which we will return in Chapter 11. The starting dose was level 3, that is, $x_1 = d_3$, which was the median dose; see Chapter 1, Table 1.1 for a detailed description of the dose schedule. Naturally, we specified the prior MTD $v_0 = 3$ for the CRM model. The sample size of the study was set at 18 (as a multiple of 3) based on budget and trial duration limitations. Using the criterion (7.1) with $K = 5$ and $\theta = 0.25$, we got

$$N > \frac{2}{\theta} + 3(K - 2) = 17.$$

Thus, a sample size of 18 seemed to be reasonable. In the actual trial, a total of 20 patients were enrolled in the phase I portion, and one patient dropped out from the study (see Figure 1.1).

At the planning stage of the study, we considered the empiric model (3.2) with a normal prior on β with mean 0 and variance $\sigma_\beta^2 = 1.34$. The major reasons for choosing this model were convenience and convention: the empiric function was the only dose–toxicity function we had programmed for at that time, and the prior distribution was determined based on an existing reference [80].

Next, we specified the initial guesses of DLT probabilities by $p_{03} = 0.25$ per the assumed $v_0 = 3$, along with $p_{01} = 0.05$, $p_{02} = 0.12$, $p_{04} = 0.40$, and $p_{05} = 0.55$. The overall indifference interval of this CRM model was 0.25 ± 0.07 . This choice was made by trial and error: an initial skeleton was chosen within the range seen in the literature [25], with a constraint that the model’s indifference interval was no wider than 20 percentage points, that is, the half-width $\delta \leq 0.10$. The CRM design with this initial skeleton was evaluated by simulations under the five dose–toxicity scenarios given in Table 7.1. Then the skeleton was iteratively modified in light of the simulated operating characteristics. This process was repeated several times in a haphazard manner, although at the end we felt the operating characteristics of the final CRM model with $N = 18$ were adequate. The numerical properties of this design will be presented in Chapter 8.

Table 7.1 The simulation scenarios used in the bortezomib trial protocol

Scenario	p_1	p_2	p_3	p_4	p_5
1	.25	.40	.45	.55	.60
2	.05	.25	.40	.45	.55
3	.05	.05	.25	.45	.55
4	.05	.05	.08	.25	.45
5	.05	.05	.08	.12	.25

Note: Numbers associated with the MTD are in bold.

7.4.2 NeuSTART

The NeuSTART introduced in Example 2.1 studied the use of high-dose lovastatin in acute ischemic stroke patients. In the trial, the primary safety event was defined as a

posttreatment elevation of peak liver or muscle enzyme levels. The trial objective was to identify the MTD associated with a $\theta = 0.10$ event probability. This decision was reached in collaboration with the clinical investigators based on a rough estimate of 3% event rate in the untreated study population; a target rate θ higher than the background was deemed acceptable after taking into account the considerations that the disability caused by an acute stroke could be severe, and that enzyme elevation would likely be reversed by withdrawing treatment. In the trial, treatment-related withdrawal would be counted as reaching the primary endpoint.

The NeuSTART used a two-stage time-to-event CRM for dose escalation at $K = 5$ test doses. As in the previous subsection, we disregard the time-to-event component here. The dose schedules used in the trial are shown in Table 7.2. The CRM model was specified by the empiric function (3.2) with skeleton $p_{01} = 0.02, p_{02} = 0.06, p_{03} = 0.10, p_{04} = 0.18, p_{05} = 0.30$. The overall indifference interval of this model was 0.10 ± 0.04 . The prior distribution was $\beta \sim N(0, 1.34)$. The model calibration process of the NeuSTART was by and large similar to that of the bortezomib trial, and was characterized by ad hoc iteration of the skeleton. The process also iterated over various sample sizes ranging from 24 to 36 (only multiples of 3 were considered), and the final sample size was set at $N = 33$.

Table 7.2 Dose schedules used in the NeuSTART and cohort sizes in case of no toxicity

Dose tier	Sample size	Lovastatin dose	
		0–72 hours	days 3–30
1	3	1 mg/kg/day	20 mg/day
2	3	3 mg/kg/day	20 mg/day
3	6	6 mg/kg/day	20 mg/day
4	9	8 mg/kg/day	20 mg/day
5	12	10 mg/kg/day	20 mg/day

This CRM model reflected the prior belief that the MTD was among the higher doses by specifying $v_0 = 3$, that is, 6 mg/kg/day for 3 days. However, the clinicians deemed it appropriate to start the trial at the standard lovastatin dose 1 mg/kg/day (dose tier 1) and escalate through the five dose tiers sequentially. Specifically, the initial design was defined by the dose sequence

$$x_{i,0}^{Neu} = \begin{cases} d_1 & \text{for } i = 1, 2, 3 \\ d_2 & \text{for } i = 4, 5, 6 \\ d_3 & \text{for } i = 7, \dots, 12 \\ d_4 & \text{for } i = 13, \dots, 21 \\ d_5 & \text{for } i = 22, \dots, 33 \end{cases} \quad (7.3)$$

and would be used to assign doses before any safety event occurred. In other words, in the case where there was no event throughout the trial, three patients would be treated at dose level 1, three at level 2, and so on, so that it would take 21 nonevents to reach the highest dose level. The determination of the dose initial sequence (7.3)

was also by trial and error. In fact, calibrating a two-stage CRM was much more complicated than calibrating a one-stage CRM by this trial-and-error approach, as the process would jointly iterate the initial dose sequence $\{x_{i,0}\}$ and the skeleton p_{0k} s. However, at the end, the initial design satisfied two criteria:

1. The clinical investigators of the study thought that the escalation speed of $x_{i,0}^{Neu}$ was appropriately cautious.
2. The initial design was not overconservative in that it was compatible; Chapter 10 provides details.

Furthermore, the model and the initial design that resulted from this process was found to perform well under the five dose–toxicity scenarios given in Table 7.3.

Table 7.3 The simulation scenarios used in the NeuSTART protocol

Scenario	p_1	p_2	p_3	p_4	p_5
1	.10	.25	.30	.35	.40
2	.04	.10	.25	.30	.35
3	.01	.04	.10	.25	.30
4	.01	.01	.04	.10	.25
5	.01	.01	.01	.04	.10

Note: Numbers associated with MTD are in bold.

7.4.3 The Case for an Automated Process

Overall, the design experience of the bortezomib trial and the NeuSTART was a mix of conventions, intuition, and validation by simulations. While the fine-tuning process was guided by objective criteria such as model sensitivity and compatibility, the use of these criteria was not systematic. This trial-and-error approach has some obvious shortcomings. First, the process is time consuming. Also, since the CRM system is quite complex, proper design tuning requires a high level of familiarity with the method. This prevents widespread use of the CRM. Second, the choice of the final design is quite arbitrary. The third problem, which is related to the second difficulty, is that the process is virtually irreproducible. The previous subsections give my best recount of what happened at the planning stage of the two trials; however it is likely to obtain a different CRM design should we go through the same trial-and-error process again. This is highly undesirable and confusing to our clinical colleagues—imagine 10 statisticians giving 10 different answers when asked to do a sample size calculation for a single-armed study!

As seen in Shen and O’Quigley [94] and Cheung and Chappell [21], a poorly calibrated CRM can lead to poor operating characteristics. Thus, it is important to choose the CRM design parameter carefully. This, together with the difficulties with the above trial-and-error approach, calls for an automated calibration process. The flowchart in Figure 7.1 outlines one such process. To illustrate how to use the flowchart, suppose the clinical investigators have chosen six dose levels for a trial,

and all but the lowest dose are higher than the standard use of the drug. If there is no external evidence a priori to suggest that the MTD is among the higher doses, it may be prudent to specify $v_0 = 1$ in the CRM model and to start the trial at $x_1 = d_1$. According to the flowchart in Figure 7.1, we may consult the techniques in Chapter 8 for model calibration.

The scheme in Figure 7.1 is not to be used as a substitute for communication between the study statistician and the clinical investigators, but rather to facilitate such communication. In this sense, the flowchart presents a semiautomated process instead of complete automation. For example, the choice of the sample size N should be revisited after all other design parameters are specified. Another useful approach is to calibrate a CRM design over a range of feasible sample sizes, so that the trade-offs between accuracy and cost can be presented to the clinical investigators for a final decision.

Also, the presented approach is a pragmatic choice and is by no means optimal nor the only logical automated approach. While the flowchart can be used as a “quick start menu” for statisticians new to the CRM, it may also serve as practical guidelines for an informed trial-and-error approach done by statisticians with experience in the CRM.

Initial Guesses of Toxicity Probabilities

8.1 Introduction

This chapter discusses the calibration of the initial guesses of toxicity probabilities associated with the test doses, or the skeleton. Section 8.2 presents an approach by which the initial guesses are specified via a single design parameter, namely, the half-width δ of the indifference intervals of the CRM model. Section 8.3 describes how δ can be chosen for a CRM application, and gives some numerical recommendations for the logistic model. Section 8.4 applies the techniques to redesign the bortezomib trial as a case study.

8.2 Half-width (δ) of Indifferent Interval

As mentioned in the earlier chapters, the dose label d_k in a CRM model is obtained via backward substitution by solving $p_{0k} = F(d_k, \hat{\beta}_0)$ where $\hat{\beta}_0$ is the prior mean of the model parameter β and p_{0k} is the initial guess of the toxicity probability associated with dose level k . Ideally, the parameter p_{0k} is a clinical choice to reflect the prior belief about dose level k . In practice, such information is seldom available from the clinical investigators, except perhaps at the prior MTD d_{v_0} for which we may set $p_{0v_0} = \theta$ so that

$$\theta = F(d_{v_0}, \hat{\beta}_0). \quad (8.1)$$

Even if the clinical investigators attempt to provide their best guess on the skeleton, the resulted CRM model may fail to meet the regularity conditions in Section 4.3; see Exercise 8.1. Therefore, it is pragmatic to treat each p_{0k} as a model parameter calibrated so that the CRM will yield good operating characteristics. To determine p_{0k} for a dose other than v_0 , Lee and Cheung [60] take an approach that is based on the sensitivity of a CRM model. Recall that the indifference interval of a CRM model under a dose–toxicity configuration with MTD $v \in \{2, \dots, K-1\}$ is equal to $(F(d_{v-1}, b_v), F(d_{v+1}, b_{v+1}))$, where b_{js} are the limits of the home sets defined in Chapter 5. Let

$$\delta_v = \frac{F(d_{v+1}, b_{v+1}) - F(d_{v-1}, b_v)}{2}$$

denote the half-width of the interval. For a given v , a CRM design \mathcal{D} with a small δ_v will converge to a dose close to the true MTD v on the probability scale, because

$$P_{\mathcal{D}}\{x_n \in \mathcal{I}_{\pi}(p - \delta_v, p + \delta_v) \text{ for all } n \geq N\} = 1$$

for some $N > 0$; see displayed equation (5.3). Suppose we can prespecify a common half-width δ that we would like to achieve for all v , that is, $\delta_v \equiv \delta$, so that the indifference intervals are invariably $\theta \pm \delta$. Then we can compute $\{d_1, \dots, d_K\}$ recursively by solving

$$F(d_{v-1}, b_v) = \theta - \delta \text{ and } F(d_{v+1}, b_{v+1}) = \theta + \delta \quad (8.2)$$

for $v = 2, \dots, K-1$, along with b_v s in accordance with Definition (5.5):

$$F(d_{v-1}, b_v) + F(d_v, b_v) = 2\theta. \quad (8.3)$$

Simplifying (8.2) and (8.3) gives the following algorithm:

Algorithm 8.1. Initial guesses of toxicity probabilities via δ

1. Calculate d_{v_0} by solving (8.1).
2. If $v_0 = 1$, skip this step; otherwise, iterate v from v_0 down to 2:
 - (a) Calculate b_v by solving $F(d_v, b_v) = \theta + \delta$;
 - (b) Calculate d_{v-1} by solving $F(d_{v-1}, b_v) = \theta - \delta$.
3. If $v_0 = K$, skip this step; otherwise, iterate v from $v_0 + 1$ up to K :
 - (a) Calculate b_v by solving $F(d_{v-1}, b_v) = \theta - \delta$;
 - (b) Calculate d_v by solving $F(d_v, b_v) = \theta + \delta$.
4. Set $p_{0k} = F(d_k, \hat{\beta}_0)$ for $k = 1, \dots, K$.

Consider the logistic model (3.3) with $a_0 = 3$ and a prior mean 0 for β , that is, $\hat{\beta}_0 = 0$, for a trial with $\theta = 0.25$, $K = 5$, and $v_0 = 3$. If we set $\delta = 0.07$, then Step 1 of Algorithm 8.1 yields $d_3 = -4.10$ by solving (8.1):

$$0.25 = \frac{\exp\{3 + \exp(0)d_3\}}{1 + \exp\{3 + \exp(0)d_3\}}.$$

Iterating v in Step 2 of the algorithm further gives, for $v = v_0 = 3$,

$$b_3 = \log\{\{\text{logit}(0.25 + 0.07) - 3\}/d_3\} = -0.088$$

and

$$d_2 = \{\text{logit}(0.25 - 0.07) - 3\}/\exp(b_3) = -4.93;$$

and, $b_2 = -0.273$ and $d_1 = -5.93$ for $v = 2$.

Likewise, we obtain $b_4 = 0.097$, $b_5 = 0.282$, $d_4 = -3.41$, and $d_5 = -2.83$ by iterating $v_0 = 4, 5$ in Step 3. Finally, applying Step 4 gives $p_{01} = 0.05$, $p_{02} = 0.13$, $p_{03} = 0.25$, $p_{04} = 0.40$, and $p_{05} = 0.54$.

Two technical points are noteworthy. First, since the indifference interval is an asymptotic criterion, the prior distribution does not play a role in the calculation of δ . Second, Algorithm 8.1 is invariant to the parameterization of the dose–toxicity model, that is, the function $c(\beta)$ in model (4.1), and will give identical outputs for all ψ -equivalent models.

Algorithm 8.1 is implemented for the empiric (3.2) and the logistic (3.3) models by the R function `getprior` in the package ‘`dfcrm`’. The initial guesses in the above CRM model can be obtained by the following R codes:

```
> # Get initial guesses with half-width del = 0.07
> library(dfcrm)
> del <- 0.07
> theta <- 0.25
> nu0 <- 3
> p0 <- getprior(del,theta,nu0,K,model="logistic")
> round(p0,digits=2)
[1] 0.05 0.13 0.25 0.40 0.54
>
```

Choosing the initial guesses based on a prespecified sensitivity δ provides a theoretical basis for the calibration process. At the same time, this approach makes the calibration process much more manageable by reducing the number of model parameters from K (p_{0k} s) to 1 (δ). The specification of the skeleton also implicitly involves the choice of the prior MTD v_0 , which is assumed to be chosen either as a clinical parameter or in accordance with (7.2) in the previous chapter. Naturally, the remaining question is how δ should be chosen.

8.3 Calibration of δ

8.3.1 Effects of δ on the Accuracy Index

A δ -sensitive CRM will eventually select a dose with toxicity probability falling between $\theta \pm \delta$. Therefore, a small value δ is desirable from an asymptotic viewpoint. On the other hand, in the finite sample setting, the CRM with a smaller δ does not necessarily perform better, especially when the true MTD is close to the prior v_0 . Let B_j^δ denote the home set for dose j evaluated in accordance with Algorithm 8.1 for a given δ —the superscript δ in B_j^δ makes it explicit that the home sets depends on the half-width δ .

Theorem 8.1. *Assume that $F_k(\beta)$ is decreasing in β for all k . Then for $\delta_1 < \delta_2$ and $v_0 = 2, \dots, K - 1$,*

$$B_1^{\delta_1} \supset B_1^{\delta_2}, B_{v_0}^{\delta_1} \subset B_{v_0}^{\delta_2} \text{ and } B_K^{\delta_1} \supset B_K^{\delta_2}.$$

The theorem states that a small value of δ corresponds to a narrow $B_{v_0}^\delta$. And, by the definition (5.4) of home sets, the Bayesian CRM will select dose level j for the next patient, if $\hat{\beta}_n \in B_j^\delta$. Therefore, the theorem suggests that the CRM with a small δ will select v_0 less often than that with a large δ . In other words, if we expect that v_0

is a correct guess of the MTD, it is preferable to use a large value of δ . Using similar arguments, we can infer that a small δ should be used if the true MTD is expected to be among the extreme ends of the dose range.

8.3.2 The Calibration Approach

The approach in this section is to calibrate the half-width δ for a one-stage Bayesian CRM, for given clinical parameters (θ, K, v_0, N) , model functional form F , and prior distribution of β , so that the CRM design will maintain a reasonably high accuracy index. As suggested by Theorem 8.1, the ideal choice of δ depends on the unknown MTD v , as well as the underlying dose–toxicity curve π . Therefore, we will need to simplify the calibration criteria by adopting the notion of risk-adjusted average [?] with respect to a calibration set \mathcal{P} of toxicity configurations and a corresponding weight vector W_π such that $\sum_{\pi \in \mathcal{P}} W_\pi = 1$.

Therefore, let $\mathcal{A}_N(\delta, \pi)$ denote the accuracy index of a CRM model specified by a half-width δ under dose–toxicity curve π . To focus on the algorithm, we consider here the accuracy index (6.1) with respect to the absolute discrepancy $\rho_k = |p_k - \theta|$. The risk-adjusted average of accuracy is defined as

$$\bar{\mathcal{A}}_N(\delta) = \sum_{\pi \in \mathcal{P}} W_\pi \mathcal{A}_N(\delta, \pi).$$

By iterating δ from 0.01 to $\theta - 0.01$ on a discrete domain having grid width 0.01, the half-width δ with the largest value of $\bar{\mathcal{A}}_N(\delta)$ will be selected as the optimal choice.

While there are many ways to choose the calibration set, a reasonable way to encompass a wide range of scenarios with minimal number of configurations is to use the plateau set $\mathcal{P}^* = \{\pi_l^* : l = 1, \dots, K\}$ where each $\pi_l^* = (p_1^*, \dots, p_K^*)^T$ with

$$p_l^* = \theta, \quad \frac{p_l^*/(1-p_l^*)}{p_k^*/(1-p_k^*)} = 2 = \frac{p_{k'}^*/(1-p_{k'}^*)}{p_l^*/(1-p_l^*)} \text{ for } k < l < k', \quad (8.4)$$

and a uniform weight $W_{\pi_l^*} = 1/K$ for all l . The number of calibration configurations in the plateau set \mathcal{P}^* is equal to the number of test doses, and each configuration has an unambiguous MTD v so that they together account for all possible v . The plateau configurations prescribe difficult scenarios for the CRM to identify the true MTD, because the performance of the CRM is expected to improve with steeper dose–toxicity curves. In addition, the toxicity probabilities at doses other than the MTD are specified so that the toxicity odds ratio is 2 when compared to the MTD: this prescribes modest jump sizes. Table 8.1 lists the plateau configurations for the scenario with $\theta = 0.25$ and $K = 5$: the jump size of toxicity probability from doses below the MTD is 11 percentage points, and above 15 percentage points; these jump sizes are modest in light of the typical sample size (20–40) used in phase I trials. Lee and Cheung [60] verify that accuracy will improve if the calibration set is defined by an odds ratio greater than 2 as in (8.4). The calibration algorithm based on this plateau \mathcal{P}^* will thus produce a design with reasonable accuracy under difficult scenarios, so as to give a lower bound of performance for steeper dose–toxicity curves. After a CRM model is chosen, we may further evaluate the design’s operating characteristics under a set of validation scenarios; see Section 8.4.

Table 8.1 The plateau calibration configurations for $\theta = 0.25$ and $K = 5$

curve	p_1^*	p_2^*	p_3^*	p_4^*	p_5^*	v
π_1^*	0.25	0.40	0.40	0.40	0.40	1
π_2^*	0.14	0.25	0.40	0.40	0.40	2
π_3^*	0.14	0.14	0.25	0.40	0.40	3
π_4^*	0.14	0.14	0.14	0.25	0.40	4
π_5^*	0.14	0.14	0.14	0.14	0.40	5

8.3.3 Optimal δ for the Logistic Model

This section presents the numerical values of the optimal δ for the one-stage Bayesian CRM using the logistic function (3.3) with $a_0 = 3$ and $\beta \sim N(0, 1.34)$ under a variety of commonly used clinical parameters. The calibration is done with respect to the plateau configurations (8.4).

To illustrate how the calibration process operates, the calibration results for $\theta = 0.25, K = 5$, and $v_0 = 3$ are plotted in Figure 8.1 for $N = 20, 30, 40$, based on 5000 simulated trials for each N . While the accuracy index depends on δ in a nonsmooth manner due to discreteness in the design parameters, some trends are quite clear:

- Accuracy improves with a small value of δ when the MTD is among the extreme ends of the dose range (i.e., MTD = 1 or 5).
- There is a reverse trend when the true MTD is equal to the prior $v_0 = 3$: accuracy is best with large δ .
- Accuracy improves as the sample size N increases, except with large values of δ .

The first two trends are consistent with what Theorem 8.1 implies, and reveal that no δ is uniformly best. The last bullet point is in line with asymptotic reasonings: the CRM with a large δ is likely inconsistent under the plateau configurations (8.4); therefore, an increase in sample size may not improve accuracy. In light of this, we will limit our search of δ by iterating from 0.01 to 0.8θ .

While there is no uniformly best δ , the lower-right panel of Figure 8.1 shows that the risk-adjusted average is a unimodal function of δ . This being the case, the average accuracy around the maximum point is quite flat. For $N = 20$, the average accuracy reaches maximum at $\delta = 0.07$ with $\bar{\mathcal{A}}_{20}(0.07) = 0.359$, but the index $\bar{\mathcal{A}}_{20}(\delta)$ coming within 0.01 of the maximum for δ ranging from 0.05 to 0.11. For $N = 40$, the highest average accuracy 0.514 is attained when $\delta = 0.05$ with $\delta = 0.04 - 0.07$ coming within 0.01 unit of the index. Therefore, we should view the “optimal” δ obtained from the calibration algorithm as a typical value among the good δ s, rather than a unique optimal choice. More importantly, the CRM model with the “optimal” δ (indicated by a solid dot in Figure 8.1) yields competitive accuracy under the five individual configurations. In the absence of a uniformly optimal solution, we deem this a reasonable design approach.

It is interesting to note that the range of δ that yields high average accuracy in Figure 8.1 gradually shifts to the left as N increases. The intuition is that as sample

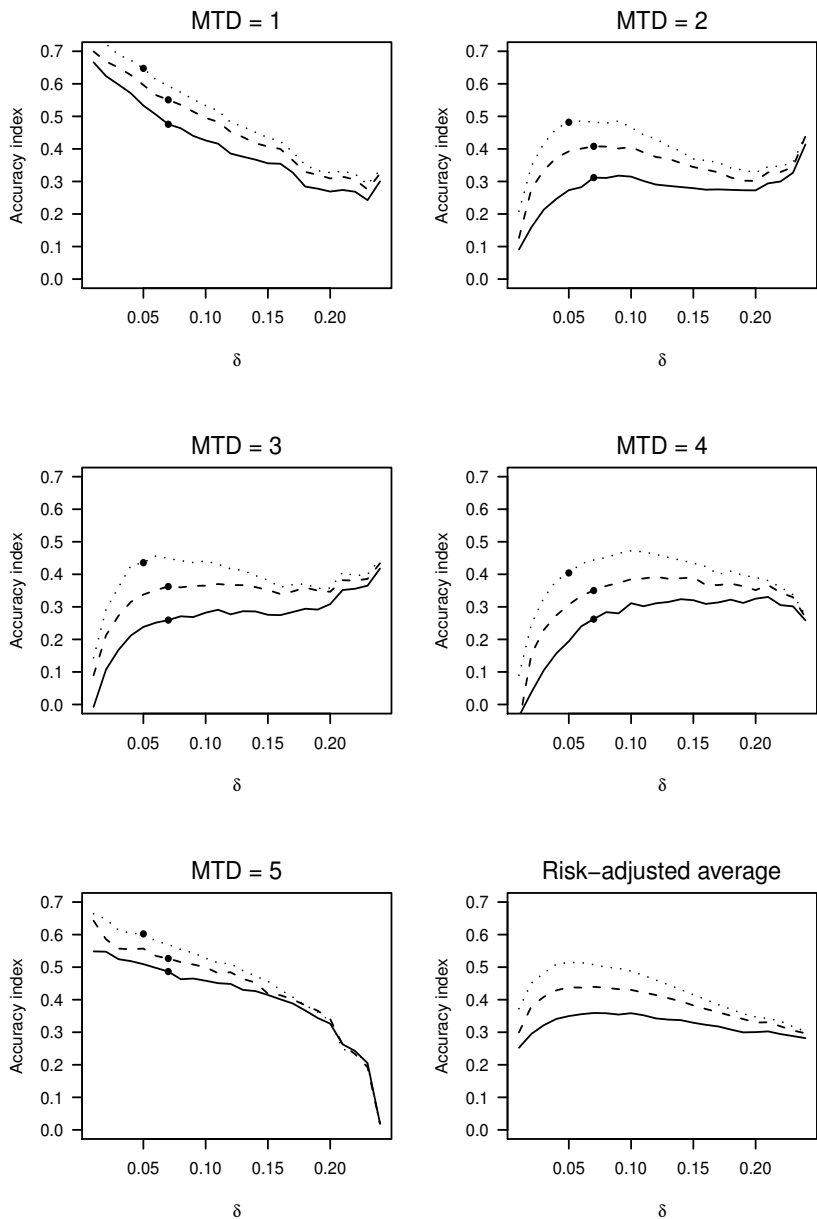


Figure 8.1 Accuracy index versus δ under the plateau configurations (8.4) with MTD ranging from 1 to 5, for clinical settings $\theta = 0.25, K = 5, v_0 = 3$, and $N = 20$ (solid), 30 (dashed), and 40 (dotted). The dot on each line indicates the value of δ that maximizes the risk-adjusted average of the accuracy index.

size grows, an asymptotic consideration becomes increasingly relevant, and thus a small δ (which is asymptotically desirable) is preferred.

We choose the accuracy index (6.1) with the absolute discrepancy as a summary index of the performance of the CRM and a numerical basis for comparison in the calibration process. However, the PCS is likely a more intuitive and direct concept than an accuracy index. Figure 8.2 plots the PCS under the plateau configurations and the risk-adjusted average PCS against δ . The patterns are extremely similar to those in Figure 8.1. In addition, the optimal δ obtained on the basis of the accuracy index, indicated by solid dots in Figure 8.2, yields competitive operating characteristics in terms of PCS and the risk-adjusted average of PCS. In fact, the calibration results are generally insensitive to the choice of the discrepancy measure used in the accuracy index, when the calibration is done with respect to the plateau calibration set \mathcal{P}^* .

Tables 8.2–8.5 show the optimal δ along with the corresponding $\bar{\mathcal{A}}_N(\delta)$ for all possible combinations of the following clinical parameters, such that the sample size constraint (7.1) is met:

- $\theta = 0.10, 0.20, 0.25, 0.33$
- $K = 4, 5, 6, 7$
- $v_0 = 1, 2, \dots, \lceil K/2 \rceil$
- $N = 20, 25, 30, 35, 40$

These results are obtained on a discrete domain of δ with grid width 0.01, ranging from 0.01 to 0.8θ . For example, for $\theta = 0.25$, the values of δ are iterated over the grid $0.01, 0.02, 0.03, \dots, 0.19, 0.20$. These results pertain to the use of the logistic function (3.3) with $a_0 = 3$ and $\beta \sim N(0, 1.34)$. Analogous results for the use of empiric model are reported in Lee and Cheung [60], who use the average PCS as the basis of comparison. Some general trends are observed:

1. The risk-adjusted average accuracy $\bar{\mathcal{A}}_N(\delta)$ of the optimal design increases with the sample size N , but decreases with the number K of test doses.
2. The range of optimal δ shifts gradually toward small values as N increases, but does not seem to depend on K .
3. There is no apparent effect of v_0 on the optimal δ and the accuracy $\bar{\mathcal{A}}_N(\delta)$, except for $\theta = 0.33$ when a higher starting dose v_0 corresponds to a larger optimal δ and higher accuracy $\bar{\mathcal{A}}_N(\delta)$.
4. The optimal δ range is 0.02–0.04 for $\theta = 0.10$; 0.04–0.08 for $\theta = 0.20, 0.25$; and 0.04–0.10 for $\theta = 0.33$.

8.4 Case Study: The Bortezomib Trial

In this section, we will discuss the calibration of the skeleton in the context of the bortezomib trial. In particular we will consider the empiric and logistic models with their respective optimal values of δ , and compare them with the original study model presented in Chapter 7, Section 7.4.1.

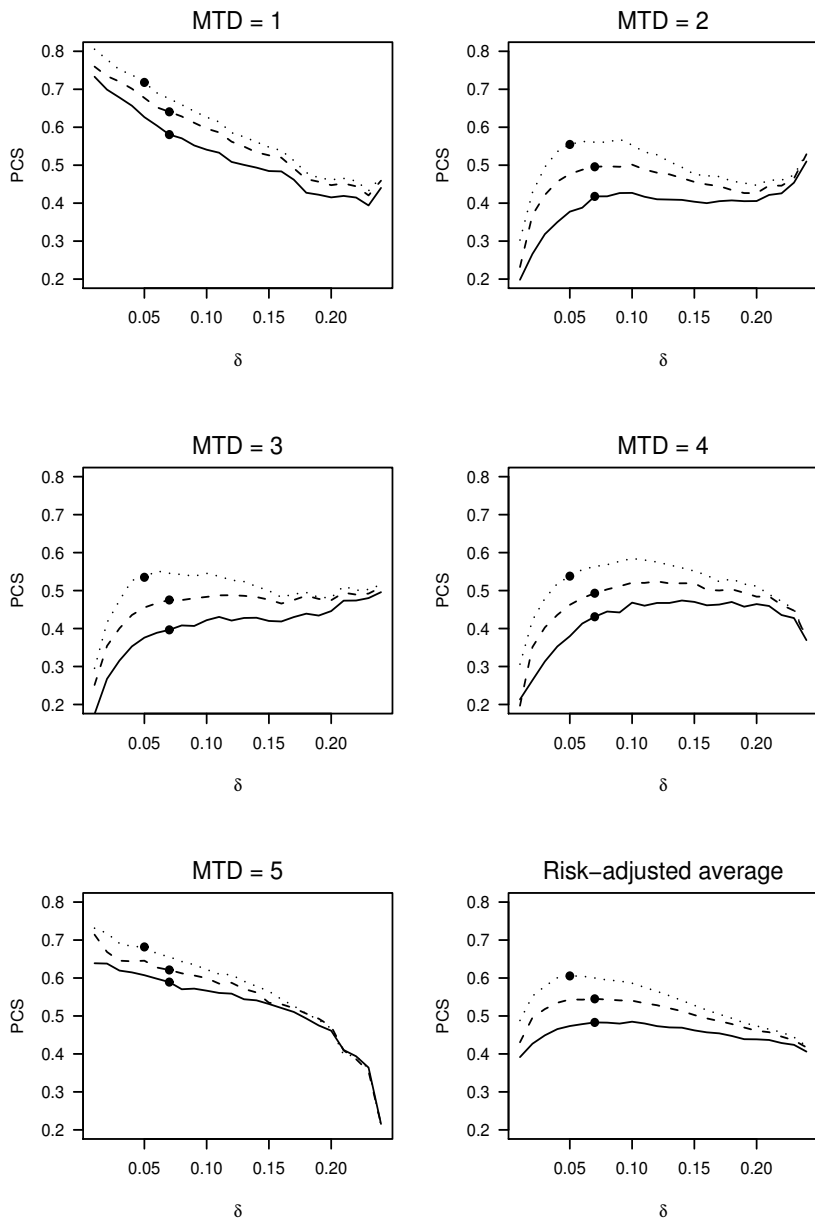


Figure 8.2 PCS versus δ under the plateau configurations (8.4) with MTD ranging from 1 to 5, for clinical settings $\theta = 0.25, K = 5, v_0 = 3$, and $N = 20$ (solid), 30 (dashed), and 40 (dotted). The dot on each line indicates the value of δ that maximizes the risk-adjusted average of the accuracy index.

Table 8.2 Optimal δ and the corresponding $\mathcal{A}_N(\delta)$ for the logistic model (3.3) with $a_0 = 3$ and $\beta \sim N(0, 1.34)$ given $\theta = 0.10$, number of dose level (K), prior MTD (v_0), and sample size (N)

K	v_0	N				
		20	25	30	35	40
4	1			(.04, .32)	(.03, .35)	(.03, .38)
	2			(.04, .32)	(.04, .36)	(.03, .38)
5	1			(.03, .29)	(.03, .31)	(.03, .34)
	2			(.03, .29)	(.03, .32)	(.03, .34)
	3			(.03, .28)	(.02, .31)	(.02, .34)
6	1				(.03, .29)	(.03, .31)
	2				(.02, .29)	(.03, .31)
	3				(.02, .29)	(.02, .31)
7	1				(.03, .27)	(.03, .30)
	2				(.03, .27)	(.02, .29)
	3				(.02, .27)	(.02, .29)
	4				(.02, .26)	(.02, .29)
Range of δ				.03-.04	.02-.04	.02-.03

Table 8.3 Optimal δ and the corresponding $\mathcal{A}_N(\delta)$ for the logistic model (3.3) with $a_0 = 3$ and $\beta \sim N(0, 1.34)$ given $\theta = 0.20$, number of dose level (K), prior MTD (v_0), and sample size (N)

K	v_0	N				
		20	25	30	35	40
4	1	(.06, .35)	(.07, .39)	(.05, .43)	(.06, .46)	(.05, .50)
	2	(.08, .36)	(.06, .40)	(.06, .44)	(.06, .48)	(.06, .51)
5	1	(.06, .32)	(.07, .36)	(.04, .40)	(.05, .43)	(.04, .46)
	2	(.07, .33)	(.07, .37)	(.06, .41)	(.06, .44)	(.06, .48)
	3	(.07, .33)	(.06, .36)	(.06, .40)	(.05, .44)	(.05, .47)
6	1		(.06, .33)	(.04, .37)	(.05, .41)	(.04, .43)
	2		(.07, .35)	(.06, .38)	(.06, .42)	(.04, .45)
	3		(.05, .35)	(.05, .38)	(.05, .42)	(.05, .45)
7	1		(.07, .32)	(.05, .35)	(.05, .38)	(.04, .42)
	2		(.06, .33)	(.06, .37)	(.06, .40)	(.04, .43)
	3		(.07, .33)	(.06, .37)	(.05, .40)	(.04, .43)
	4		(.07, .32)	(.05, .36)	(.04, .39)	(.05, .42)
Range of δ		.06-.08	.05-.07	.04-.06	.04-.06	.04-.06

Table 8.4 Optimal δ and the corresponding $\varphi_N(\delta)$ for the logistic model (3.3) with $a_0 = 3$ and $\beta \sim N(0, 1.34)$ given $\theta = 0.25$, number of dose level (K), prior MTD (v_0), and sample size (N)

K	v_0	N				
		20	25	30	35	40
4	1	(.06, .38)	(.05, .41)	(.06, .46)	(.06, .50)	(.05, .54)
	2	(.07, .39)	(.08, .44)	(.07, .47)	(.07, .51)	(.06, .55)
5	1	(.07, .35)	(.06, .39)	(.06, .43)	(.05, .46)	(.05, .50)
	2	(.06, .36)	(.08, .40)	(.07, .45)	(.07, .48)	(.06, .51)
	3	(.07, .36)	(.08, .42)	(.07, .44)	(.06, .48)	(.05, .51)
6	1	(.07, .32)	(.05, .36)	(.05, .41)	(.05, .44)	(.04, .47)
	2	(.07, .34)	(.06, .38)	(.07, .42)	(.06, .45)	(.06, .49)
	3	(.07, .34)	(.06, .38)	(.06, .42)	(.06, .46)	(.05, .49)
7	1		(.05, .34)	(.05, .39)	(.05, .41)	(.04, .44)
	2		(.06, .36)	(.05, .40)	(.05, .43)	(.06, .46)
	3		(.06, .36)	(.06, .41)	(.06, .44)	(.05, .47)
	4		(.07, .36)	(.07, .40)	(.07, .43)	(.05, .46)
Range of δ		.06-.08	.05-.08	.05-.07	.05-.07	.04-.06

Table 8.5 Optimal δ and the corresponding $\varphi_N(\delta)$ for the logistic model (3.3) with $a_0 = 3$ and $\beta \sim N(0, 1.34)$ given $\theta = 0.33$, number of dose level (K), prior MTD (v_0), and sample size (N)

K	v_0	N				
		20	25	30	35	40
4	1	(.07, .40)	(.07, .45)	(.07, .49)	(.07, .53)	(.07, .57)
	2	(.10, .43)	(.09, .47)	(.09, .51)	(.07, .55)	(.06, .58)
5	1	(.07, .37)	(.07, .41)	(.06, .45)	(.05, .49)	(.05, .53)
	2	(.07, .39)	(.07, .43)	(.07, .47)	(.07, .52)	(.06, .55)
	3	(.10, .40)	(.09, .44)	(.09, .48)	(.07, .51)	(.06, .55)
6	1	(.07, .35)	(.07, .38)	(.05, .42)	(.05, .46)	(.04, .50)
	2	(.07, .37)	(.07, .41)	(.06, .44)	(.07, .48)	(.06, .52)
	3	(.09, .38)	(.09, .42)	(.07, .45)	(.06, .49)	(.06, .53)
7	1		(.06, .36)	(.05, .40)	(.05, .43)	(.04, .47)
	2		(.07, .38)	(.06, .42)	(.06, .45)	(.05, .49)
	3		(.09, .40)	(.06, .43)	(.06, .47)	(.06, .50)
	4		(.09, .40)	(.07, .43)	(.07, .47)	(.06, .50)
Range of δ		.07-.10	.06-.09	.05-.09	.05-.07	.04-.07

The optimal empiric model. Lee and Cheung [60] propose a calibration algorithm for δ similar to that in Section 8.3.2 for the empiric model used in the bortezomib study, and recommend $\delta = 0.10$. Using the function `getprior` in the ‘dfcrm’ library

```
> p0 <- getprior(0.10, 0.25, 3, 5, model="empiric")
```

we get $p_{01} = 0.01$, $p_{02} = 0.08$, $p_{03} = 0.25$, $p_{04} = 0.46$, and $p_{05} = 0.65$. This skeleton represents a slightly steeper prior than that of the bortezomib study model.

The optimal logistic model. To use the logistic model (3.3) with $a_0 = 3$, one may refer to Table 8.4. For $K = 5$, $v_0 = 3$, and $N = 20$ (closest to 18), the optimal value of δ is 0.07. Using the function `getprior`

```
> p0 <- getprior(0.07, 0.25, 3, 5, model="logistic")
```

we get $p_{01} = 0.05$, $p_{02} = 0.12$, $p_{03} = 0.25$, $p_{04} = 0.40$, and $p_{05} = 0.54$. It seems to be a coincidence that this set of initial guesses is almost identical to that of the study model and that both have similar indifference intervals.

While the calibration of the optimal empiric and logistic models are done with respect to the plateau configurations, it is important to validate the CRM’s accuracy under realistic dose–toxicity scenarios. These validation scenarios should be chosen in collaboration with the clinical investigators. The validation step is particularly crucial for the bortezomib trial because the plateau configurations do not reflect the general expectation of dose–toxicity relationship of an anticancer agent. Table 8.6 compares the performance of the study model and the two optimal models under the dose–toxicity configurations used in the original study protocols (cf. Table 7.1). Overall, the three designs have comparable accuracy with about 60% or higher PCS under all scenarios: the average PCS over the five scenarios is 0.65, 0.63, and 0.64, respectively, for the study model, the optimal empiric, and the optimal logistic model; the average accuracy index \mathcal{A}_{20} is identical to 0.63 for all three models. The accuracy of the nonparametric optimal design is also provided in the table as a reference. The nonparametric optimal design is generally superior to the CRM designs with average PCS 0.76 and average $\mathcal{A}_{20} = 0.75$; but the CRM designs are quite efficient. For example, the efficiency of the study model in terms of average PCS can be computed as

$$\frac{0.65}{0.76} \times 100\% \approx 86\%.$$

Similarly, the efficiencies of the three CRM designs are approximately 82% to 86% of the nonparametric optimal design.

From a safety viewpoint, the optimal empiric model is preferred because the overdose numbers are below that of balanced randomization in all five scenarios. In contrast, the study model and the optimal logistic model treat on average over 4 patients at dose level 5 (an overdose) under Scenario 4.

The variability of the PCS across the scenarios is also noteworthy. Since the true dose–toxicity curve is unknown, it will be reassuring to look at the range of the PCS and verify that the method’s performance is adequate across all scenarios. In this regard, the study model is *slightly* superior: the ranges of the PCS are 10, 13, and 13 percentage points, respectively, for the study model, optimal empiric, and optimal logistic; and, the respective standard deviations of PCS are 0.04, 0.06, and

Table 8.6 The distribution of dose selection, average toxicity number (ATN), and overdose number (OD) of three CRM models for the bortezomib trial with $\theta = 0.25$, $K = 5$, $v_0 = 3$, and $N = 20$

Model	Probability of selecting dose					\mathcal{A}_{20}	ATN	OD
Scenario 1:	.25	.40	.45	.55	.60			
NP Optimal	.78	.16	.05	.01	.00	.82	—	—
The study model	.67	.25	.08	.01	.00	.72	6.9	9.5
Optimal empiric	.58	.31	.10	.01	.00	.65	7.2	11.3
Optimal logistic	.69	.23	.07	.01	.00	.74	6.7	8.5
Scenario 2:	.05	.25	.40	.45	.55			
NP Optimal	.07	.71	.16	.05	.01	.70	—	—
The study model	.12	.58	.25	.05	.00	.57	5.8	8.4
Optimal empiric	.08	.60	.26	.06	.00	.60	6.0	8.7
Optimal logistic	.12	.57	.24	.06	.01	.56	5.6	8.2
Scenario 3:	.05	.05	.25	.45	.55			
NP Optimal	.00	.07	.78	.13	.01	.76	—	—
The study model	.00	.15	.68	.16	.01	.64	5.2	5.1
Optimal empiric	.00	.12	.71	.16	.01	.68	5.2	4.8
Optimal logistic	.00	.16	.64	.18	.01	.59	5.1	5.4
Scenario 4:	.05	.05	.08	.25	.45			
NP Optimal	.00	.00	.12	.74	.14	.68	—	—
The study model	.00	.01	.19	.64	.16	.57	4.6	4.2
Optimal empiric	.00	.01	.20	.67	.12	.61	4.5	3.1
Optimal logistic	.00	.02	.19	.61	.18	.53	4.8	4.8
Scenario 5:	.05	.05	.08	.12	.25			
NP Optimal	.00	.00	.01	.19	.79	.80	—	—
The study model	.00	.01	.05	.28	.66	.67	3.6	0.0
Optimal empiric	.00	.01	.06	.34	.60	.60	3.4	0.0
Optimal logistic	.00	.01	.04	.24	.70	.71	3.7	0.0

Note: The results for the nonparametric optimal design (NP optimal) are given as references. Numbers associated with the MTD are in bold.

0.05. It should be noted, however, that comparing variability of PCS is meaningful only when the methods are comparable in terms of average PCS. This aspect will be explored further in Chapter 9. Scenario-by-scenario, the optimal empiric model, having a larger half-width δ than the other two models, is correct more often than the other two when the true MTD is equal to the prior MTD $v_0 = 3$. The converse holds when the true MTD is on the extreme ends of the dose range. However, the overall performance of the three models is very comparable.

Based on the accumulative simulation experience from a number of CRM trials, the CRM models obtained by the calibration process generally yield competitive operating characteristics against the corresponding study models. This should be read in light of the convenience the automated calibration approach offers. The automated

calibration process is much less labor-intensive than the trial-and-error approach, and can be programmed in a few lines of R code with help from the ‘dfcrm’ package. The process is mildly computer-intensive and may take a day or two for the computer to run. For the most common clinical settings, the δ values presented in Section 8.3.3 usually give reasonable designs.

8.5 Exercises and Further Results

Exercise 8.1. For a trial with $K = 5$ dose levels and a target $\theta = 0.10$, a clinical investigator believes that the highest dose is the MTD a priori and that the low doses are virtually nontoxic. Thus, he provides $p_{01} = p_{02} = 0.01$, $p_{03} = p_{04} = 0.05$, and $p_{05} = 0.10$. Evaluate the dose labels for the five levels if the empiric function (3.2) with $\hat{\beta}_0 = 0$. Comment on this skeleton.

Exercise 8.2. Prove the following properties about skeletons obtained by applying Algorithm 8.1 with a half-width δ :

- a. The skeleton is strictly increasing, that is, $p_{01} < \dots < p_{0K}$.
- b. $|p_{0k} - \theta|$ is strictly increasing in δ (Lee and Cheung [60], Theorem 1).

The result in part b implies that a large δ corresponds to a steep skeleton, and a small δ to a shallow one. As such, a large δ may in turn be interpreted as a strong prior input, because the prior MTD v_0 is quite distinct from the other doses in that initial guess p_{0k} is much different from $p_{0,v_0} = \theta$ for $k \neq v_0$. Likewise, a small δ (shallow skeleton) indicates little difference between the prior MTD v_0 and the other doses in terms of initial guesses of toxicity probabilities, and may be viewed as a weak prior input.

Least Informative Normal Prior

9.1 Introduction

This chapter discusses the specification of the prior standard deviation σ_β for the model parameter β in a Bayesian CRM. Section 9.2 introduces the notion of least informative prior in the context of dose finding, and summarizes the guidelines on the specification of prior distribution in terms of its variability. We consider two calibration approaches: Section 9.3 presents an approach informed by the use of the least informative prior for a given sensitivity δ , and Section 9.4 optimizes the least informative model with respect to δ . Section 9.5 compares these approaches in the context of the bortezomib trial.

9.2 Least Informative Prior

9.2.1 Definitions

As discussed in Chapter 4, a Bayesian CRM using model (4.1) with $c(\beta) = \exp(\beta)$ is invariant to the mean of a normal prior distribution. Thus, we can consider a mean zero normal prior

$$\beta \sim N(0, \sigma_\beta^2) \tag{9.1}$$

without loss of generality. As a result, the prior distribution is completely specified by the standard deviation σ_β . In conventional Bayesian inference, a large value of σ_β corresponds to a vague prior, whereas a small σ_β conveys strong prior beliefs. In the context of dose finding, however, the notion of a noninformative prior is to be defined with reference to our belief about the MTD, rather than that about the model parameter β .

Let μ_j denote the probability that dose j is the MTD computed under the prior distribution, that is, $\mu_j \equiv \Pr(\beta \in B_j) = \Pr(v_\beta = j)$, where v_β is the model-based MTD when β is the true parameter value, so that $v_\beta = j$ if and only if β belongs to the home set B_j . Moreover, since the home sets are disjoint intervals defined by (5.5), the prior mass $\mu_j = G(b_{j+1}) - G(b_j)$ for $j = 1, \dots, K$, where $b_1 \equiv -\infty$ and $b_{K+1} \equiv \infty$. For the mean zero normal prior (9.1), we have

$$\mu_j = \Phi(b_{j+1}/\sigma_\beta) - \Phi(b_j/\sigma_\beta)$$

for $j = 1, \dots, K$, where Φ is the cumulative distribution function of a standard normal.

In addition, the first two central moments of v_β are

$$E(v_\beta) \equiv e_v = \sum_{j=1}^K j\mu_j \text{ and } \text{var}(v_\beta) \equiv \sigma_v^2 = \sum_{j=1}^K j^2\mu_j - e_v^2. \quad (9.2)$$

Table 9.1 shows the distribution of v_β for the logistic model (3.3) with $a_0 = 3$, $\delta = 0.07$, and various values of σ_β for a trial with $\theta = 0.25$, $K = 5$, and $v_0 = 3$. The distribution of v_β is unimodal with peak at $v_0 = 3$ when $\sigma_\beta = 0.20$, and gradually turns U-shaped as σ_β increases. For example, when $\sigma_\beta = 1.16$ (which is the value used in the optimal logistic model in Chapter 8, Section 8.4), the prior implies that either dose 1 or dose 5 is the MTD with over 80% probability. Therefore, a prior with a large value of σ_β is by no means vague. In contrast, when $\sigma_\beta = 0.33$ or 0.50 , the distribution of v_β is quite flat, and represents ignorance about the true MTD a priori. Thus, a noninformative prior in terms of v_β corresponds to the uniform distribution, that is, $\mu_j = 1/K$ for all j , and the vagueness of a prior can be determined by how close μ_j s is to a uniform distribution.

Table 9.1 The distribution of v_β for the logistic model (3.3) with $a_0 = 3$ and $\delta = 0.07$ for a trial with target $\theta = 0.25$, $K = 5$, and $v_0 = 3$

j	(b_j, b_{j+1})	μ_j with $\sigma_\beta =$			
		0.20	0.33	0.50	1.16
1	$(-\infty, -0.27)$	0.09	0.21	0.29	0.41
2	$(-0.27, -0.09)$	0.24	0.19	0.14	0.06
3	$(-0.09, 0.10)$	0.36	0.22	0.15	0.06
4	$(0.10, 0.28)$	0.23	0.19	0.14	0.06
5	$(0.28, \infty)$	0.08	0.20	0.29	0.40
e_v		2.98	2.98	2.99	2.99
σ_v		1.07	1.41	1.61	1.84

Among the many ways to measure the distance between two distributions, we focus on an approach that matches the standard deviations of two distributions.

Definition 9.1 (least informative normal prior). For any given CRM model $F_k(\beta)$, the normal prior (9.1) is said to be least informative if the corresponding σ_v as defined in (9.2) equals the standard deviation of a uniform distribution on $\{1, \dots, K\}$, that is, $\sigma_v = \sqrt{(K^2 - 1)/12}$. The standard deviation of the least informative prior of β is denoted by σ_β^{LI} .

For example, for $K = 5$, a least informative prior distribution should correspond to $\sigma_v = \sqrt{2} = 1.41$. For the logistic model in Table 9.1, we see that $\sigma_\beta^{LI} = 0.33$ and the corresponding $\mu_j \approx 0.20$ for all j .

A natural alternative to compare μ_j s to the uniform distribution is based on the

Kullback-Leibler divergence criterion:

$$D_{KL}(\mu \mid unif) = \sum_{j=1}^K \mu_j \log(K\mu_j). \tag{9.3}$$

A noninformative normal prior can be obtained by choosing σ_β that minimizes the “distance” $D_{KL}(\mu \mid unif)$, or equivalently, maximizes the *entropy* of the distribution μ_j s: $-\sum_{j=1}^K \mu_j \log \mu_j$. The value of σ_β chosen in accordance with this criterion will be called the maximum entropy standard deviation. See Berger [9] (pages 90–94).

Figure 9.1 shows σ_v (left) and D_{KL} (right) against σ_β for various CRM model. In all the three models shown in the figure, the least informative and maximum entropy standard deviations turn out to be identical. In general, the least informative prior per Definition 9.1 is very similar to that by maximizing entropy. Also, it is clear that σ_v increases as σ_β increases; thus, the least informative σ_β^I can be easily calculated by binary search. For these reasons, we will focus on the least informative prior from now on.

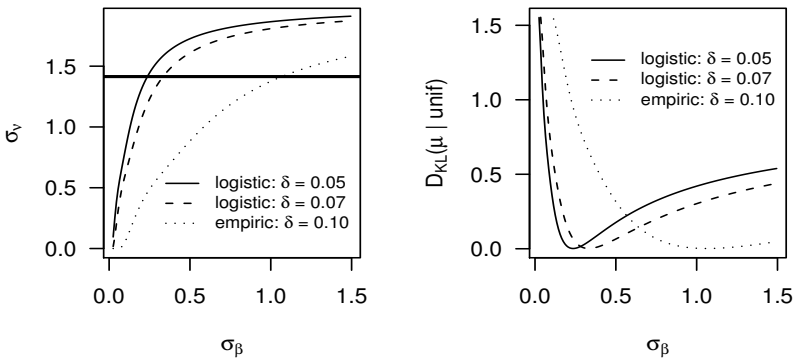


Figure 9.1 Left: σ_v versus σ_β for the logistic model (3.3) with $a_0 = 3$ and $\delta = 0.05, 0.07$ and the empiric model (3.2) with $\delta = 0.10$ under the setting with $\theta = 0.25, K = 5$, and $v_0 = 3$. The least informative standard deviation is the σ_β value that gives $\sigma_v = 1.41$, as indicated by the dark solid line. Right: The Kullback-Leibler divergence against σ_β for the same model specifications. The maximum entropy standard deviation is the σ_β value that minimizes the Kullback-Leibler divergence criterion (9.3).

9.2.2 Rules of Thumb

From Figure 9.1, we observe that the value of least informative standard deviation varies with the CRM models. This underlines the fact that the parameter β lives on different scales under different models. Therefore, a given value of σ_β may convey

different degrees of informativeness in different CRM models. The notion of least informative standard deviation hence provides some useful guidelines on the notion of “vagueness”:

Definition 9.2. *The normal prior distribution (9.1) is said to be vague for a given CRM model if $\sigma_\beta \approx \sigma_\beta^{LI}$.*

Definition 9.3. *The normal prior distribution (9.1) is said to have large variation for a given CRM model if $\sigma_\beta \gg \sigma_\beta^{LI}$.*

Definition 9.4. *The normal prior distribution (9.1) is said to be informative favoring v_0 for a given CRM model if $\sigma_\beta \ll \sigma_\beta^{LI}$.*

Table 9.2 shows the least informative σ_β^{LI} for the logistic model (3.3) with $a_0 = 3$ for $\theta = 0.25$ under some common design parameters. For this logistic model, the least informative σ_β^{LI} ranges from 0.169 to 0.835. And as a case in point, the optimal logistic model presented in Section 8.4 is used with $\sigma_\beta = 1.16$, which is clearly a prior with large variation.

As a rule of thumb, when the trial has a low starting dose level (e.g., $x_1 = d_1$) and the specified prior MTD v_0 is equal to the starting dose, a large variation prior may be used in order to ensure flexibility for the CRM to reach the highest dose level within a reasonable sample size (cf. Table 9.1). On the other hand, if the starting dose is close to the median dose level (7.2), a vague prior will likely be appropriate. An informative prior favoring v_0 should almost never be used unless there is very strong prior evidence that v_0 is the true MTD. These guidelines are to be elaborated next.

Table 9.2 Least informative standard deviation σ_β^{LI} for the logistic model (3.3) with $a_0 = 3$ for given $\theta = 0.25$, number of dose level (K), prior MTD (v_0), and half-width (δ) of the indifference intervals

K	$\sigma_{uniform}$	v_0	δ				
			0.04	0.05	0.06	0.07	0.08
4	1.118	1	0.248	0.312	0.376	0.441	0.507
		2	0.169	0.212	0.255	0.299	0.344
5	1.414	1	0.301	0.378	0.456	0.534	0.614
		2	0.223	0.279	0.337	0.395	0.454
		3	0.189	0.237	0.285	0.334	0.383
6	1.708	1	0.355	0.446	0.537	0.630	0.724
		2	0.278	0.349	0.420	0.492	0.566
		3	0.229	0.288	0.346	0.406	0.466
7	2.000	1	0.410	0.514	0.620	0.727	0.835
		2	0.333	0.418	0.504	0.590	0.679
		3	0.277	0.347	0.418	0.491	0.564
		4	0.256	0.321	0.387	0.453	0.521

9.3 Calibration of σ_β

9.3.1 Calibration Criteria

One may adopt a similar calibration approach to that in Chapter 8, Section 8.3.2: fix all other design parameters ($\theta, K, v_0, N, F, p_{0ks}$), and iterate σ_β over a fine grid so as to obtain a value that yields high accuracy. More precisely, let $\mathcal{A}_N(\sigma_\beta, \pi)$ denote the accuracy index of a one-stage Bayesian CRM assuming σ_β under dose–toxicity curve π and define the risk-adjusted average accuracy as

$$\bar{\mathcal{A}}_N(\sigma_\beta) = \sum_{\pi \in \mathcal{P}^*} W_\pi \mathcal{A}_N(\sigma_\beta, \pi) \quad (9.4)$$

where \mathcal{P}^* is the plateau configurations (8.4) and $W_\pi \equiv 1/K$.

The left panel of Figure 9.2 plots the risk-adjusted average accuracy (9.4) against σ_β for the logistic model (3.3) with $a_0 = 3$ and the empiric model (3.2) for the clinical setting $\theta = 0.25, K = 5, v_0 = 1, 2, 3$, and $N = 20$. The calibration is done by iterating σ_β from 0.10 to 10.00 on a discrete grid with width 0.05, but the plots include only a partial range with prominent features so as to highlight the essential trend visually: The average accuracy index becomes quite flat and remains close to the maximum once σ_β is large enough. For example, for the logistic model with $\delta = 0.07$ and $v_0 = 3$, the average accuracy ranges from 0.35 to 0.39 for $\sigma_\beta \geq 0.30$. In this regard, the risk-adjusted average criterion (9.4) does not appear to be discriminatory enough for choosing σ_β .

The right panel of Figure 9.2 plots the standard deviation $\mathcal{S}_N(\sigma_\beta)$ of the accuracy index against σ_β , where

$$\mathcal{S}_N(\sigma_\beta) = \sum_{\pi \in \mathcal{P}^*} W_\pi \left\{ \mathcal{A}_N(\sigma_\beta, \pi) - \bar{\mathcal{A}}_N(\sigma_\beta) \right\}^2. \quad (9.5)$$

It is noteworthy that the variability $\mathcal{S}_N(\sigma_\beta)$ demonstrates a roughly reverse trend in its dependence on σ_β when compared to $\bar{\mathcal{A}}_N(\sigma_\beta)$. Provided that the accuracy is near maximum, a small variability is desirable because it implies the method performs uniformly across dose–toxicity scenarios with the true MTD ranging from 1 to K . Hence, an alternative calibration criterion is to choose σ_β that minimizes $\mathcal{S}_N(\sigma_\beta)$. The criterion based on (9.5) appears to be more sensitive to σ_β at the large values than that based on (9.4). For the logistic model with $\delta = 0.07$ and $v_0 = 3$, for instance, the standard deviation $\mathcal{S}_N(\sigma_\beta)$ ranges from 0.017 to 0.107, a six-fold difference, for $0.3 \leq \sigma_\beta \leq 1.0$.

9.3.2 An Application to the Choice of v_0

Figure 9.2 suggests that setting v_0 as the median dose level reduces variability. For the empiric model, the standard deviation $\mathcal{S}_N(\sigma_\beta) \geq 0.15$ for all σ_β with $v_0 = 1$, whereas $\mathcal{S}_N(\sigma_\beta)$ is uniformly below 0.10 for sufficiently large σ_β with $v_0 = 3$. In light of the fact that the average accuracies are comparable for different v_0 , it seems to be advantageous to specify the median dose as the prior MTD in the CRM model.

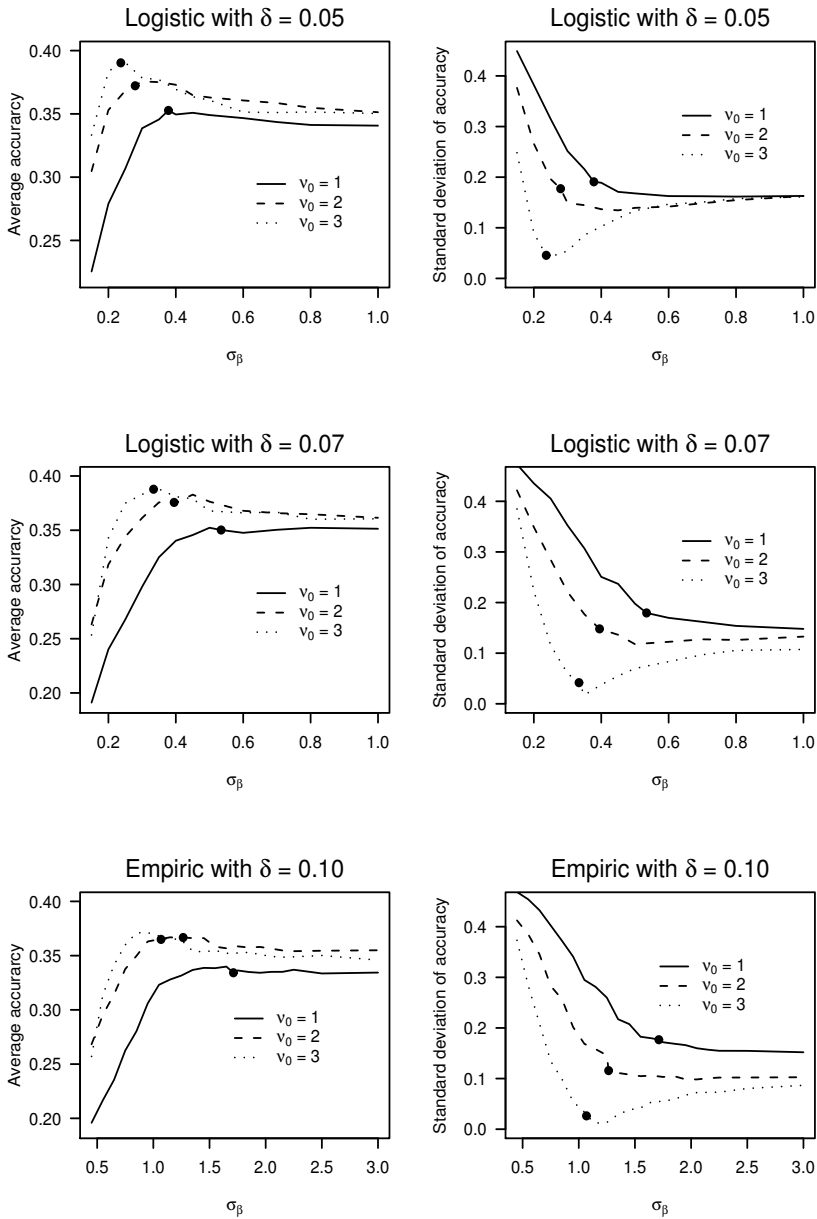


Figure 9.2 Left: Average accuracy $\bar{\mathcal{A}}_N(\sigma_\beta)$ versus σ_β under the plateau configurations in Table 8.1 for $\theta = 0.25, K = 5, \nu_0 = 1, 2, 3$, and $N = 20$. Right: $\mathcal{S}_N(\sigma_\beta)$ versus σ_β . The solid dot on each line indicates the value at the least informative σ_β^{LI} .

This is in accordance with the original recommendation by O’Quigley et al. [78]. In summary, a pragmatic approach is to take v_0 as the highest dose level that is less than or equal to the median dose level according to (7.2); this approach in effect takes the view that v_0 is a model parameter (that is to be tuned) rather than a clinical parameter (that is to be set with respect to the clinical opinions and constraints).

9.3.3 Optimality Near σ_β^{LI}

Interestingly, Figure 9.2 shows that the CRM with a *vague* prior is very competitive in terms of both the average (9.4) and the variability (9.5) criteria.

To offer an explanation, the left panel of Figure 9.3 plots the individual accuracy index $\mathcal{A}_N(\sigma_\beta, \pi_k^*)$ of the CRM using the logistic model with $\delta = 0.07$ and $v_0 = 3$ under the plateau configurations (Table 8.1). When σ_β is small, the prior of v_β is unimodal at $v_0 = 3$ (cf. Table 9.1), and, therefore, favors dose level 3 a priori and tends to choose dose 3 a posteriori. Thus, the CRM performs well under π_3^* , where the true MTD is dose 3. For the same reason, however, the CRM shows little flexibility to choose the extreme doses under π_1^* and π_5^* , thus lowering the average accuracy and increasing the variability of the accuracy index. On the other hand, with a prior with large variation, the CRM does better under π_1^* and π_5^* than under π_3^* . When σ_β is around 0.3–0.4, the prior of v_β is roughly uniform and does not favor any dose. As a result, the CRM performs uniformly across all scenarios and demonstrates small variability $\mathcal{S}_N(\sigma_\beta)$ of the accuracy index.

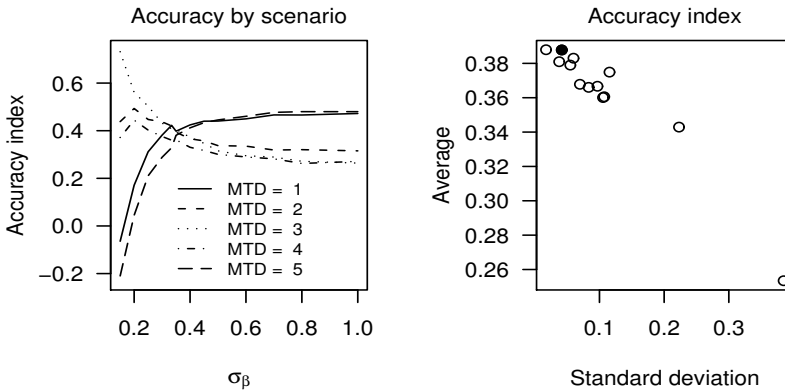


Figure 9.3 Left: $\mathcal{A}_N(\sigma_\beta, \pi_k^*)$ versus σ_β under the plateau configurations (Table 8.1) for the logistic model with $a_0 = 3$ and $\delta = 0.07$ under $\theta = 0.25, K = 5, v_0 = 3$, and $N = 20$. Right: $\bar{\mathcal{A}}_N(\sigma_\beta)$ versus $\mathcal{S}_N(\sigma_\beta)$; the solid dot indicates the least informative prior.

The above reasoning suggests that a *locally* optimal value of σ_β among the vague priors may be a reasonable choice. This approach solves the difficulties that the range of σ_β in theory increases without bound and that the range of good σ_β (i.e., large \mathcal{A}_N

and small \mathcal{S}_N) varies with the other design parameters such as the dose–toxicity function. The right panel of Figure 9.3 shows a strong negative association between $\mathcal{S}_N(\sigma_\beta)$ and $\tilde{\mathcal{S}}_N(\sigma_\beta)$, and confirms that maximizing average (9.4) and minimizing variability (9.5) are concordant objectives. Therefore, in the search for a local optimum of σ_β , we will consider both criteria, although the variability is expected to be more discriminatory than the average criterion.

Algorithm 9.1. Locally optimal σ_β near the least informative prior

1. For given clinical parameters (θ, K, v_0, N) and CRM model $F_k(\beta)$, evaluate the least informative σ_β^{LI} .
2. Iterate σ_β in the neighborhood of σ_β^{LI} from $\lambda_1 \sigma_\beta^{LI}$ to $\lambda_2 \sigma_\beta^{LI}$ on a discrete domain having grid width 0.05 for some $\lambda_1 < 1 < \lambda_2$. We suggest $\lambda_1 = 0.8$ and $\lambda_2 = 1.5$.
3. Choose the standard deviation σ_β that maximizes the average accuracy (9.4) or minimizes the variability of accuracy (9.5) with respect to the plateau calibration configurations (8.4).

Based on the observed trends about σ_β^{LI} in Table 9.2 and the search domain as defined in Algorithm 9.1, we expect the range of locally optimal σ_β to shift to the right as

- The number K of test doses increases, or
- The prior MTD v_0 decreases, or
- The half-width δ of the model’s indifference interval increases.

Intuitively, as the number of test doses increases, large variability a priori is required to allow sufficient flexibility so that the experimentation can move quickly across doses in response to the data during the trial. This intuition also applies to trials with the prior MTD v_0 being among an extreme end. Finally, a large value of σ_β is needed to counterbalance the strong prior input due to a large half-width δ (see Chapter 8, Exercise 8.2).

Example 9.1. Consider the logistic model with $a_0 = 3$, $\delta = 0.07$, $K = 5$, and $v_0 = 3$. From Table 9.2, we obtain $\sigma_\beta^{LI} = 0.334$. Then applying Algorithm 9.1 with $\lambda_1 \approx 0.8$ and $\lambda_2 \approx 1.5$, we iterate σ_β on the grid $\{0.25, 0.30, 0.35, 0.40, 0.45, 0.50\}$ and obtain 0.35 as the locally optimal standard deviation based on both criteria (9.4) and (9.5).

Example 9.2. Consider the empiric model with $\delta = 0.10$, $K = 5$, and $v_0 = 3$. The least informative $\sigma_\beta^{LI} = 1.068$ (cf. Figure 9.1). By a grid search from 0.85 to 1.60 with width 0.05, the average accuracy (9.4) is maximized when $\sigma_\beta = 0.95$ and the variability (9.5) is minimized when $\sigma_\beta = 1.15$. The latter is almost identical to the empiric model in Section 8.4 that assumes a normal prior with $\sigma_\beta^2 = 1.34$.

In both examples, the average and variability criteria yield similar optimal

choices that are close to the least informative σ_β^{LI} ; for the logistic model in particular, Algorithm 9.1 reaches the same local optimum regardless of the criterion used.

9.4 Optimal Least Informative Model

It is in principle straightforward to extend the calibration process to account for both the standard deviation σ_β and the sensitivity δ by iterating these two parameters on a discrete domain over a two-dimensional grid. The computational burdens may, however, impose an impractical time constraint.

As illustrated in Example 9.1 and 9.2, the least informative prior is quite close to the (locally) optimal standard deviation. Therefore, it seems reasonable to choose the δ that gives the best least informative CRM model:

Algorithm 9.2. Optimal δ among the least informative CRM models

1. For given clinical parameters (θ, K, v_0, N) and a dose–toxicity function F , iterate δ from 0.01 to 0.8θ on a discrete domain having grid width 0.01.
2. For each δ and hence a completely specified CRM model $F_k(\beta)$, evaluate the least informative standard deviation, $\sigma_\beta^{LI}(\delta)$.
3. Choose the pair $(\delta, \sigma_\beta^{LI}(\delta))$ that maximizes $\bar{\mathcal{A}}_N(\sigma_\beta^{LI}(\delta))$ as defined in (9.4) with respect to the calibration configurations (8.4).

The calibration results using Algorithm 9.2 are plotted in Figure 9.4 for the empiric model (3.2) and the logistic model (3.3) with $a_0 = 3$ under the setting $\theta = 0.25, K = 5, v_0 = 3$, and $N = 20$. The dependence of the accuracy index on δ (top panel of the figure) demonstrates a different pattern from that in Figure 8.1, where a fixed and large standard deviation is used. For example, under π_5^* where the true MTD is dose 5, the performance of CRM improves as δ increases in Figure 9.4, but worsens in Figure 8.1. This is not unexpected. Because the least informative standard deviation is an increasing function of δ (cf. lower left corner of Figure 9.4), when δ is small, a prior concentrated around $v_0 = 3$ is used, and thus prevents the CRM from testing and choosing an extreme dose. In contrast, the accuracy index depends on the standard deviation (middle panel of Figure 9.4) in the same manner as in Figure 9.3.

The lower right corner of Figure 9.4 summarizes the calibration results. In brief, the optimal least informative model is specified by $\delta = .06$ with $\sigma_\beta^{LI} = 0.285$ for the logistic, and $\delta = .04$ with $\sigma_\beta^{LI} = 0.418$ for the empiric. Both models have the same average accuracy $\bar{\mathcal{A}}_N = 0.39$. This suggests the choice of the dose–toxicity function is not important as long as the other model parameters, namely δ and σ_β , are chosen properly. In addition, the average criterion (9.4) seems to be able to distinguish a range of good δ from the bad values: the average accuracies of the least informative models are quite flat when δ ranges from 0.04 to 0.08, and are very comparable for the two models. Therefore, the optimal δ obtained per Algorithm 9.2 should be viewed as a typical value among the good δ s.

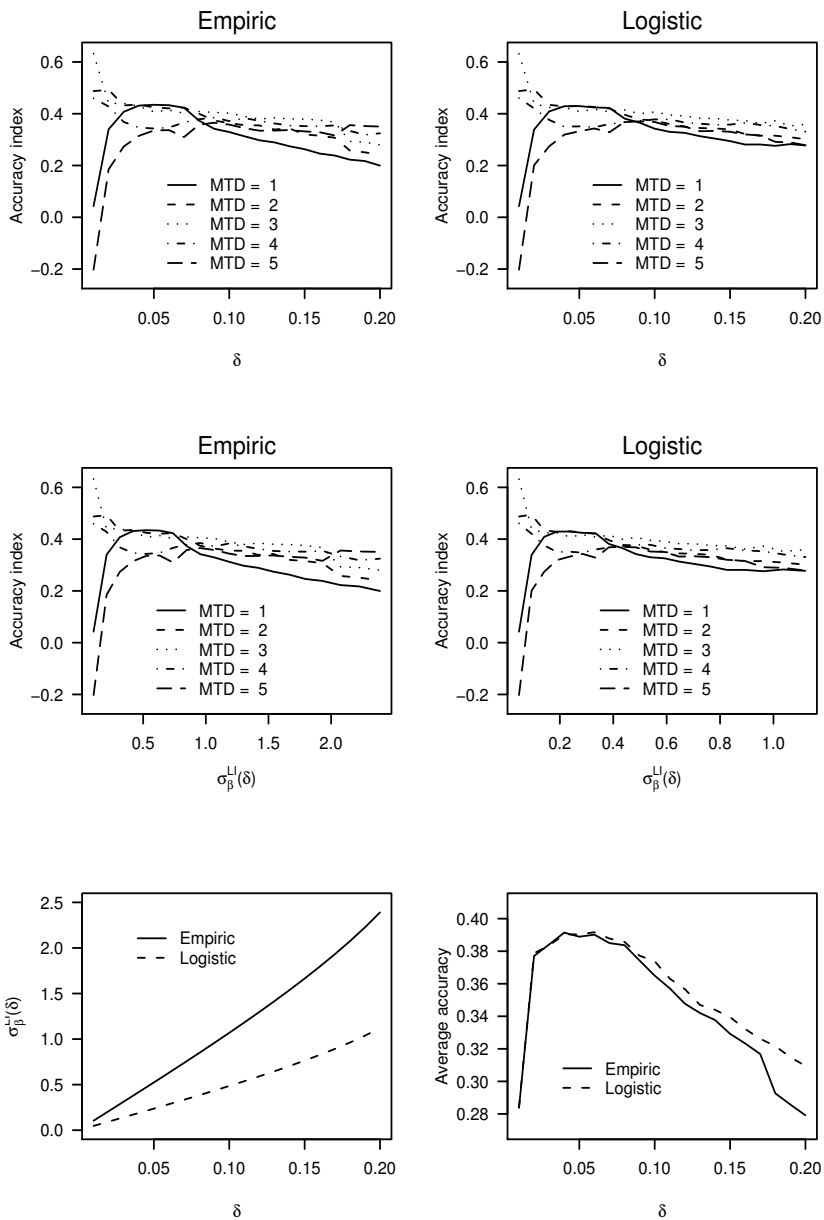


Figure 9.4 Top: Accuracy index versus δ under the five plateau configurations in Table 8.1 for the corresponding least informative empiric and logistic models under the clinical setting $\theta = 0.25, K = 5, v_0 = 3$, and $N = 20$. Middle: Accuracy index versus $\sigma_{\beta}^{LI}(\delta)$. Lower left: $\sigma_{\beta}^{LI}(\delta)$ versus δ . Lower right: Average accuracy index versus δ .

Table 9.3 shows the optimal least informative specification $(\delta, \sigma_{\beta}^{LI}(\delta))$ for the logistic model (3.3) with $a_0 = 3$ and a mean zero normal prior (9.1) under some common clinical parameters. The dependence of the optimal least informative δ on the sample size N and the number K of test doses is less clear than that shown in Chapter 8 where the standard deviation σ_{β} is fixed and large. However, the range of good δ is quite similar to, albeit slightly narrower than, those given in Tables 8.2–8.5. To summarize:

- For $\theta = 0.10$, the optimal least informative δ range is 0.02–0.03.
- For $\theta = 0.20$, the optimal least informative δ range is 0.03–0.06.
- For $\theta = 0.25, 0.33$, the optimal least informative δ range is 0.04–0.06.

Table 9.3 Optimal least informative specification $(\delta, \sigma_{\beta}^{LI}(\delta))$ of the logistic model (3.3) with $a_0 = 3$ and the normal prior (9.1) for given θ, K, N

θ	K	N				
		20	25	30	35	40
0.10	4			(.03, .21)	(.02, .14)	(.03, .21)
	5			(.02, .16)	(.02, .16)	(.02, .16)
	6				(.02, .19)	(.02, .19)
	7				(.02, .21)	(.02, .21)
0.20	4	(.05, .23)	(.05, .23)	(.04, .19)	(.06, .28)	(.05, .23)
	5	(.05, .26)	(.04, .21)	(.03, .16)	(.04, .21)	(.05, .26)
	6		(.06, .38)	(.04, .25)	(.05, .32)	(.05, .32)
	7		(.04, .28)	(.04, .28)	(.04, .28)	(.04, .28)
0.25	4	(.04, .17)	(.05, .21)	(.05, .21)	(.06, .26)	(.06, .26)
	5	(.06, .29)	(.06, .29)	(.06, .29)	(.04, .19)	(.04, .19)
	6	(.06, .35)	(.05, .29)	(.05, .29)	(.05, .29)	(.05, .29)
	7		(.05, .32)	(.04, .26)	(.05, .32)	(.04, .26)
0.33	4	(.09, .36)	(.06, .24)	(.06, .24)	(.06, .24)	(.06, .24)
	5	(.06, .27)	(.05, .22)	(.04, .18)	(.06, .27)	(.05, .22)
	6	(.06, .32)	(.06, .32)	(.06, .32)	(.06, .32)	(.06, .32)
	7		(.05, .30)	(.05, .30)	(.05, .30)	(.05, .30)

Note: The prior MTD v_0 is assumed to be the median dose level, that is, $v_0 = 2, 3, 3, 4$ respectively for $K = 4, 5, 6, 7$.

9.5 Revisiting the Bortezomib Trial

This section compares the performance of various models by Algorithms 9.1 and 9.2 under the validation scenarios in Table 8.6 for the bortezomib trial. The comparison includes the following models:

- The study model, the optimal empiric and logistic models in Section 8.4, denoted, respectively, as models (S), (E0), and (L0).

- The logistic model with locally optimal $\sigma_\beta = 0.35$ for $\delta = 0.07$ in Example 9.1, denoted as model (L1).
- The empiric model with locally optimal $\sigma_\beta = 0.95$ for $\delta = 0.10$ with respect to the average accuracy (9.4) criterion in Example 9.2, denoted as model (E1). The locally optimal σ_β with respect to the variability (9.5) criterion is not considered because it is very similar to model (E0).
- The optimal least informative empiric and logistic models in Section 9.4, denoted as models (E2) and (L2) respectively.

The numerical results are shown in Table 9.4, which also gives the results for the nonparametric optimal design as references. The model (E1) has large variability in PCS; there is no reason to expect otherwise since the calibration criterion is the average accuracy (9.4). Using this model, the CRM does very well when the MTD is dose 3—with a 0.75 PCS when compared to a 0.78 PCS for the nonparametric optimal design. However, this design is correct only 54% of the time when the MTD is dose 1. This variation in accuracy can be quite unsettling as one does not have control nor knowledge about the true MTD. By contrast, model (L1) improves upon (L0) in terms of variability; this may be due to the fact that (L1) explicitly minimizes the variability criterion (9.5). In addition, this reduction is accomplished without sacrificing the average accuracy. The optimal least informative models (E2) and (L2) yield lower average PCS than that of the study model (E0), but the difference is small. Also, both models bring down the variability in PCS to a comparable range to that of the nonparametric optimal design. The advantage is clear especially for model (L2): the individual PCS under all five scenarios are uniformly above 60%, thus providing some assurance of accuracy.

Table 9.4 Operating characteristics of the nonparametric optimal design (NP) and CRM with various model specifications for the bortezomib trial with $\theta = 0.25$, $K = 5$, $v_0 = 3$, and $N = 20$ under the validation configurations in Table 8.6

Model	δ	σ_β	PCS under Scenario					Ave	Range ^a
			1	2	3	4	5		
(NP)	—	—	.78	.71	.78	.74	.79	.76	.08
(S)	—	1.16	.67	.58	.68	.64	.66	.65	.10
(E0)	0.10	1.16	.58	.60	.71	.67	.60	.63	.13
(E1)	0.10	0.95	.54	.61	.75	.69	.57	.63	.21
(E2)	0.04	0.42	.61	.61	.68	.63	.58	.62	.10
(L0)	0.07	1.16	.69	.57	.64	.61	.70	.64	.13
(L1)	0.07	0.35	.62	.60	.69	.66	.61	.64	.09
(L2)	0.06	0.29	.61	.61	.69	.65	.61	.63	.08

^aRange = maximum PCS – minimum PCS.

In this simulation, the best performing models are (L1) and (L2). In light of the high variation of model (E1) in Table 9.4, we may use the variability criterion (9.5) with Algorithm 9.1. Finally, the optimal least informative models reduce variability

in the accuracy from the carefully calibrated study model (E0) while maintaining comparable PCS. Thus, the convenience and reproducibility of Algorithm 9.2 shall prove a reasonable design approach; and the optimal least informative specification for the logistic model is readily available in Table 9.3 for most common settings.

Initial Design

10.1 Introduction

In situations where the clinical investigators believe the MTD is among the higher doses but are not comfortable treating patients at a high dose without first testing its lower doses, one may adopt a two-stage CRM that starts a trial according to an initial sequence $\{x_{i,0}\}$ with $x_{1,0} = d_1$. This chapter discusses the characterization and the choice of the initial sequence. To be specific, Section 10.2 introduces some basic concepts about the ordering of dose sequences. Section 10.3 presents the calibration approach for $\{x_{i,0}\}$, along with numerical recommendations in some common clinical scenarios. Practical issues such as sample size considerations related to the use of a two-stage CRM will be addressed in Section 10.4. The calibration techniques will be illustrated in Section 10.5 in the context of the NeuSTART.

10.2 Ordering of Dose Sequences

A two-stage CRM assigns doses in accordance with a predetermined, nondecreasing dose sequence $\{x_{i,0}\}$ until the first toxicity is seen, that is, $x_i = x_{i,0}$ for $i \leq \min\{j : Y_j = 1\}$ where

$$x_{i,0} \leq x_{i+1,0}. \quad (10.1)$$

The primary appeal of a two-stage design is that it allows a low starting dose, which is often the lowest dose level, that is, $x_{1,0} = d_1$. On the other hand, the initial dose sequence should avoid treating many patients at low and inefficacious doses, for which the 3+3 algorithm is criticized. A reasonable strategy is to escalate quickly initially and slow down as the trial moves to higher doses [98]. Mathematically, it implies that the initial cohort sizes are monotone, that is,

$$m_{01} \leq m_{02} \leq \cdots \leq m_{0,K} \quad (10.2)$$

where

$$m_{0k} = \sum_{i=1}^N I(x_{i,0} = d_k)$$

is the *initial cohort size* at dose d_k according to the sequence $\{x_{i,0}\}$. Under (10.1), the initial cohort sizes m_{0k} s uniquely define the initial design. In this chapter, we will focus on initial dose sequences that satisfy (10.1) and (10.2).

A fast dose escalation scheme is often perceived to be aggressive, and a slow one conservative. Thus, a notion of relative dose escalation speed is in order:

Definition 10.1 (partial ordering). Let $\mathbf{x} = \{x_i\}$ and $\tilde{\mathbf{x}} = \{\tilde{x}_i\}$ denote two sequences of doses. Dose escalation in the sequence \mathbf{x} is said to be faster or greater than that in the sequence $\tilde{\mathbf{x}}$, denoted by $\mathbf{x} > \tilde{\mathbf{x}}$ if $x_i \geq \tilde{x}_i$ for all i , with the inequality being strict for some i .

To illustrate, let \mathbf{x}_0 denote the initial design used in a two-stage CRM \mathcal{D}_2 , and $\tilde{\mathbf{x}}$ denote a dose sequence generated by this design, that is, $\tilde{x}_i = \mathcal{D}_2(H_i)$. Hence, \mathbf{x}_0 is a predetermined fixed sequence, whereas $\tilde{\mathbf{x}}$ is a random sequence that depends on the toxicity outcomes in the trial. The notion of compatibility (Definition 5.1) can then be expressed as the condition under which $\mathbf{x}_0 > \tilde{\mathbf{x}}$ with probability one. In other words, compatibility requires that the initial design represents the fastest (or, so to speak, most aggressive) dose escalation plan that is permissible with respect to the trial objective θ . This is a sensible requirement: the initial design is the dose escalation plan when no toxicity is observed; and when there are any observed toxic outcomes, it is undesirable that we should escalate faster than when there is none.

Another application of Definition 10.1 is to compare two initial dose sequences used in a two-stage CRM. To illustrate, the dose escalation plan for the NeuSTART was designed using a two-stage CRM with $N = 33$ patients and $K = 5$ test doses. The study initial design is given by (7.3), which can be equivalently represented by $m_{01} = m_{02} = 3$, $m_{03} = 6$, $m_{04} = 9$, and $m_{05} = 12$. Table 10.1 shows this NeuSTART initial sequence (bottom row) along with three other sequences that may be used for the trial. Since larger initial cohort sizes m_{0k} s correspond to slower escalation, it is obvious that the group-of-three design in Table 10.1 is faster than the group-of-four design and the NeuSTART sequence $\mathbf{x}_0^{\text{Neu}}$.

Table 10.1 Four examples of initial designs for a trial with $K = 5$ and $N = 33$

Design	m_{01}	m_{02}	m_{03}	m_{04}	m_{05}	$m_{+,4}$
Group-of-three, $\mathbf{x}_0^{(3)}$	3	3	3	3	21	12
Group-of-four, $\mathbf{x}_0^{(4)}$	4	4	4	4	17	16
Increasing cohort size, $\mathbf{x}_0^{\text{Inc}}$	2	4	6	8	13	20
NeuSTART, $\mathbf{x}_0^{\text{Neu}}$	3	3	6	9	12	21

Furthermore, we can verify that

$$\mathbf{x}_0^{\text{Inc}} > \mathbf{x}_0^{\text{Neu}}$$

because

$$x_{i,0}^{\text{Inc}} \geq x_{i,0}^{\text{Neu}} \text{ for all } i$$

with the inequality being strict when $i = 3, 21$; cf., Figure 10.1. It is equally interesting to note that although the group-of-three escalation is faster than group-of-four, there exists a sequence (i.e., $\mathbf{x}_0^{\text{Inc}}$) that does not show clear ordering with neither

of these. In other words, strict ordering does not necessarily exist between two sequences, and as such, the initial dose sequences constitute a partially ordered set.

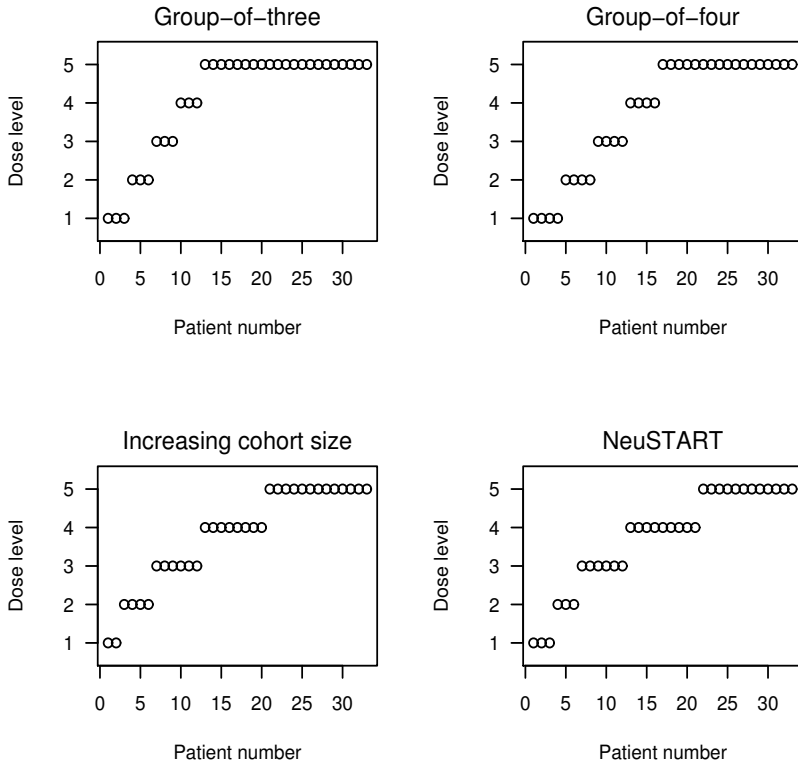


Figure 10.1 Dose level by patient number according to the initial designs in Table 10.1.

A simple index to indicate conservatism of an initial sequence is by the number of nontoxic observations required to reach the highest dose, that is,

$$m_{+,K-1} = \sum_{k=1}^{K-1} m_{0k}.$$

A conservative initial design is associated with a large value of $m_{+,K-1}$. In general, let the number of nontoxic observations required by an initial design to reach dose $j + 1$ be denoted as

$$m_{+,j} = \sum_{k=1}^j m_{0k}. \tag{10.3}$$

Definition 10.2 (total ordering). Let $\mathbf{x} = \{x_i\}$ and $\tilde{\mathbf{x}} = \{\tilde{x}_i\}$ denote two sequences of doses that satisfy (10.1) and (10.2), and \mathbf{m} and $\tilde{\mathbf{m}}$ their respective cohort sizes. The sequence \mathbf{x} is said to be more aggressive than $\tilde{\mathbf{x}}$, denoted by $\mathbf{x} \succ \tilde{\mathbf{x}}$, if $m_{+,J_0} < \tilde{m}_{+,J_0}$ where $J_0 = \max\{j : m_{+,j} \neq \tilde{m}_{+,j}\}$.

A practical advantage of the relation “ \succ ” is that it is a total order. As a result, we can use it to rank dose sequences according to their aggressiveness. For example, in Table 10.1, the sequence $\mathbf{x}_0^{\text{Neu}}$ is the least aggressive as it takes $m_{+,4} = 21$ consecutive nontoxic outcomes to escalate to dose level 5, whereas the group-of-three is the most aggressive with $m_{+,4} = 12$. Overall, we have

$$\mathbf{x}_0^{(3)} \succ \mathbf{x}_0^{(4)} \succ \mathbf{x}_0^{\text{Inc}} \succ \mathbf{x}_0^{\text{Neu}}.$$

10.3 Building Reference Initial Designs

10.3.1 Coherence-Based Criterion

As stated in Chapter 5, Section 5.2, a one-stage CRM is coherent, but a two-stage CRM is not necessarily so. A logical deduction is that the incoherence problem arises from improper choice of the initial design. Theorem 5.2 precisely indicates that a cause of incoherence is the use of an incompatible initial design. To see how incoherence may occur, consider an outcome sequence by a two-stage CRM with $x_{1,0} = \dots = x_{m',0} = d_1$ and $Y_1 = \dots = Y_{m'-1} = 0$ and $Y_{m'} = 1$ such that $\mathcal{D}_2(H_{m'+1}) \geq d_2$. This sequence is incoherent because an escalation takes place for patient $m' + 1$ after a toxic outcome is observed in patient m' . Indeed, this outcome sequence is made possible only when m_{01} is set to be sufficiently large, that is, $m_{01} \geq m'$, and can be avoided by choosing a small m_{01} . This illustrates that an overconservative initial design causes incoherence and incompatibility.

To appeal to our intuition, take another instance with a target $\theta = 0.25$. It is quite clear then that an initial design that escalates after every 10 nontoxic outcomes, that is, $m_{0k} \equiv 10$, is overconservative. Likewise, an initial sequence with $m_{0k} \equiv 8$ also appears conservative, though not as much as $m_{0k} \equiv 10$. As we decrease m_{0k} , the initial design seems to be increasingly reflective of the objective $\theta = 0.25$. The natural question is precisely how small m_{0k} should be so that the initial design is appropriate. This is the question the coherence (compatibility) criterion may address.

Example 10.1. Consider the two-stage CRM used in Chapter 3, Figure 3.1. The right panel of the figure shows that the group-of-three initial rule is incompatible. In Chapter 5, Section 5.2.2 subsequently illustrates that the group-of-two initial design is compatible; and so is the group-of-one rule. That is, among all initial designs with constant initial cohort sizes, that is, $m_{0k} \equiv m_0$, the only compatible choices are $m_{0k} = 1$ or 2. In view of the inclination to be conservative, one may adopt $m_{0k} = 2$ for $k = 1, \dots, 4$ in conjunction with this CRM model.

10.3.2 *Calibrating Compatible Dose Sequences*

As compatibility and the inclination to be conservative are two opposing criteria, it is reasonable to calibrate the initial design to be the “most conservative” compatible design. There are then two possible approaches to implement the calibration process. The first approach starts with a compatible initial design, iterates to a less aggressive dose sequence at each step, and stops iteration once an incompatible initial design is attained. This calibration approach can be implemented by the following algorithm:

Algorithm 10.1. Base- b compatible benchmark

1. For given θ, K, v_0 , CRM model $F_k(\beta)$, and the prior on β , set $j = v_0$ and specify a base (b) initial design $m_{0k} = 0$ for $k < j$, and $= b$ for $k \geq j$.
2. Check the compatibility of the initial design $\{m_{01}, \dots, m_{0,K-1}, m_{0K}\}^\dagger$.
3. If compatibility holds:

- (a) Record this initial design
- (b) Iterate

$$j = \begin{cases} j - 1 & \text{if } j > 1 \\ K - 1 & \text{if } j = 1 \end{cases}$$

and then update m_{0j} with $m_{0j} + b$, and repeat Step (2).

4. If compatibility fails to hold in Step 2, stop iteration and choose the last recorded initial design in Step 3a.

[†] The cohort size m_{0K} of the highest dose does not affect whether the initial design is compatible, and can be arbitrarily specified and updated in the algorithm as long as (10.2) is satisfied.

At each iteration, Step 3b moves to test a more conservative initial design. It can be proved that once an incompatible sequence is reached in Step 4, all subsequent dose sequences will be incompatible; see Theorem 10.1 in Section 10.4. Thus, the last recorded design in Step 3a may be viewed as a conservative benchmark for compatible dose sequences, and shall be called a base- b benchmark. It is easy to see any base- b benchmark satisfies the constraint (10.2).

Example 10.2. For the CRM model in Example 10.1, if we apply Algorithm 10.1 with $b = 1$, we will test dose sequences in the following order:

Iteration	m_{01}	m_{02}	m_{03}	m_{04}	Compatible
1	0	0	1	1	Yes
2	0	1	1	1	Yes
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
9	2	2	3	3	Yes
10	2	3	3	3	No

We stop at the tenth iteration when we reach an incompatible initial design. The dose

sequence in the ninth iteration, $m_{01} = m_{02} = 2$ and $m_{03} = m_{04} = 3$, is the base-1 benchmark.

At each iteration, compatibility can be verified using the function `cohere`. For instance, for the initial design at the ninth iteration,

```
> theta <- 0.25
> K <- 5
> p0 <- c(0.05,0.12, 0.25, 0.40, 0.55)
> m0 <- c(2,2,3,3,3) # set m0[5]=3 arbitrarily
> x0 <- rep(1:5, m0)
> foo <- cohere(p0,theta,x0,model="logistic")
> foo$message
[1] "Coherent"
>
```

The base-1 benchmark in Example 10.2 is more conservative than the group-of-two design obtained in Example 10.1, and may agree more with clinicians' expectations. Ideally, one should seek a compatible initial sequence $\tilde{\mathbf{x}}_0$ such that $\mathbf{x}_0 > \tilde{\mathbf{x}}_0$ for all compatible \mathbf{x}_0 . Unfortunately, it is unclear whether such a sequence exists in general, mainly because of the fact that dose sequences form a partially ordered set. Table 10.2 displays all possible base- b benchmarks for this CRM model; when $b \geq 9$, there is no compatible initial sequence that starts at $x_{1,0} \leq d_{v_0} = d_3$. To summarize, the base-2 design is more conservative than the base-1, but is neither faster nor slower than the base-8. This confirms the conjecture that the slowest escalating compatible sequence may not exist due to partial ordering. However, the base-8 design is the most conservative sequence among all, according to the relation \succ . Since the starting dose is a clinical decision, there may be further restrictions on $x_{1,0}$. In particular, it is common to adhere to sequences with $x_{1,0} = d_1$, which implies $m_{01} > 0$. In this case, we will choose between the base-1 and base-2 sequences as the initial design for this given CRM model. As a two-stage CRM is motivated by conservatism, it is natural to choose base-2 benchmark in this example.

Table 10.2 Base- b initial sequences for the CRM model in Example 10.1

b	m_{01}	m_{02}	m_{03}	m_{04}	$m_{+,2}$	$m_{+,3}$	$m_{+,4}$
1	2	2	3	3	4	7	10
2	2	2	4	4	4	8	12
3	0	3	3	3	3	6	9
4	0	4	4	4	4	8	12
5	0	0	5	5	0	5	10
6	0	0	6	6	0	6	12
7	0	0	7	7	0	7	14
8	0	0	8	8	0	8	16

A second calibration approach proceeds in the opposite direction: start with an

incompatible initial design, and increase escalation speed at each iteration until a compatible initial design is obtained. An analog of Algorithm 10.1 can be used to implement this calibration approach. However, it can be verified using Theorem 10.1 that this approach will yield the same base- b benchmark given by Algorithm 10.1.

10.3.3 Reference Initial Designs for the Logistic Model

This subsection presents some conservative references of initial designs that are to be used with the logistic model (3.3):

$$F(d_k, \beta) = \frac{\exp\{3 + \exp(\beta)d_k\}}{1 + \exp\{3 + \exp(\beta)d_k\}} \quad (10.4)$$

in a two-stage Bayesian CRM, where $\beta \sim N(0, 1.34)$ a priori and the dose labels d_k s are determined via a specified half-width δ by Algorithm 8.1. We take the following calibration approach: for a given CRM model, apply Algorithm 10.1 to obtain the base- b benchmark for each possible b with $m_{01} > 0$, and among them select the most conservative sequence according to the total order “ \succ ”. (This is essentially what we did in Section 10.3.2.) The results are shown in Table 10.3 for the clinical parameters with

- $\theta = 0.10, 0.20, 0.25, 0.33$
- $K = 4, 5, 6, 7$.

The prior MTD v_0 is set at the median dose level (7.2). The ranges of δ in the table are chosen in accordance with the optimal δ given in Tables 8.2–8.5 for cases with $N \geq 30$ and v_0 at the median dose; the results here are, therefore, to be used in conjunction with those tables.

Table 10.3 demonstrates some general trends regarding the escalation speed of the reference initial designs. Especially, the initial dose sequence will need to be more aggressive as

- A higher target θ is used, or,
- A larger number K of doses are tested, or,
- The CRM model is specified by a smaller half-width δ .

The first trend is expected; the second trend is also quite clear. The third trend is comparatively nuanced, but provides an additional consideration for the choice of δ in the context of a two-stage CRM. For example, one may be inclined to choose a δ that corresponds to a more conservative reference initial design.

10.4 Practical Issues

10.4.1 Sample Size Constraint

Algorithm 10.1 prescribes a base- b benchmark without regard for the initial cohort size m_{0K} at the highest dose. In practice, it may be prudent to set aside an adequate number m^* of observations at the highest level in case of no toxicity, that is,

$$m_{0K} \geq m^*.$$

Table 10.3 Reference initial designs for model (10.4)

θ	K	δ	m_{01}	m_{02}	m_{03}	m_{04}	m_{05}	m_{06}	b	
0.10	4	.03	9	9	18	—	—	—	9	
		.04	9	18	18	—	—	—	9	
	5	.02	5	5	5	10	—	—	5	
		.03	5	5	10	10	—	—	5	
	6	.02	5	5	5	10	10	—	5	
	7	.02	3	3	3	6	6	6	3	
	0.20	4	.06	4	8	8	—	—	—	4
.05			2	2	4	4	—	—	2	
5		.06	3	3	6	6	—	—	3	
		.05	2	2	4	4	4	—	2	
6		.04	2	2	2	4	4	4	2	
7		.04	2	2	2	2	2	2	2	
		.05	1	2	2	2	2	2	1	
0.25	4	.06	4	4	8	—	—	—	4	
		.07	4	6	6	—	—	—	2	
	5	.05	2	2	4	4	—	—	2	
		.06	2	2	4	4	—	—	2	
		.07	2	2	4	4	—	—	2	
	6	.05	2	2	4	4	4	—	2	
		.06	2	2	4	4	4	—	2	
	7	.05	1	2	2	2	2	2	1	
		.07	1	1	2	2	2	2	1	
	0.33	4	.06	3	3	6	—	—	—	3
			.07	3	3	6	—	—	—	3
			.09	3	3	6	—	—	—	3
		5	.06	1	2	2	2	—	—	1
			.07	1	2	2	2	—	—	1
.09			1	2	2	2	—	—	1	
6		.06	1	2	2	2	2	—	1	
		.07	1	2	2	2	2	—	1	
7		.06	1	1	1	2	2	2	1	
		.07	1	1	2	2	2	2	1	

The choice of m^* can be partly informed by the target rate θ . For example, if we set $m^* = 1/\theta$, we will expect to observe one toxic outcome at the highest dose if no toxic outcome is observed at the lower doses. To account for the possibility that the actual escalation is slower and more patients will be treated at the lower doses, we may choose $m^* \geq \lambda/\theta$ for some $\lambda > 1$. As a practical guideline, we recommend choosing λ to be somewhere between 1.5 and 2.0. This recommendation serves as a lower limit of the required sample size; but the final choice should be evaluated based on the operating characteristics by simulations. At any rate, the sample size N

is determined by the choice of m^* and the initial design:

$$N \geq m_{+,K-1} + m^*. \tag{10.5}$$

For example, if we use the reference initial design corresponding to $\theta = 0.25$, $K = 5$, and $\delta = 0.07$ in Table 10.3, and set $m^* = 1.5/0.25 = 6$, then we may use $N \geq 18$. That is, the constraint (10.5) can be used to give a quick assessment as to whether the sample size N is adequate; cf., constraint (7.1) for the one-stage CRM.

In many practical situations, however, the sample size N is limited by the time and resources available; and a base- b benchmark may require more patients than the available resources can provide. Consider the NeuSTART where $N = 33$, for instance. If we use the reference initial design corresponding to $\theta = 0.10$, $K = 5$, and $\delta = 0.03$ in Table 10.3, the trial will have enrolled 30 subjects before reaching the highest dose and left $m_{05} = 3$ for the dose—if no toxicity occurred in the first 30 subjects; the trial may not reach dose level 5 with this initial design even with one observed toxic outcome. Also, this initial design—with $m_{05} = 3$ —violates constraint (10.2). In other words, a fixed and small sample size N prevents the initial escalation from being too slow. A pruned base- b benchmark may then be used:

Algorithm 10.2. Pruning a compatible sequence for given N and m^*

1. For given $\theta, K, v_0, F_k(\beta)$, the prior on β , specify a compatible initial design $\{m_{01}, \dots, m_{0,K-1}\}$ and set $m_{0K} = N - m_{+,K-1}$.
2. If $m_{0K} < m^*$, set $j = 0$ and proceed with the following steps:

(a) Iterate

$$j = \begin{cases} j + 1 & \text{if } j < K - 1 \\ 1 & \text{if } j = K - 1 \end{cases}$$

and then update m_{0j} with $m_{0j} - 1$ and m_{0K} with $m_{0K} + 1$.

(b) Stop iteration when $m_{0K} = m^*$.

Example 10.3. Consider $N = 33$ and $m^* = 12$ in the context of the NeuSTART. The initial design $m_{01} = m_{02} = 5, m_{03} = m_{04} = 10$ in Table 10.3 may be pruned according to Algorithm 10.2 in the following order:

Iteration	m_{01}	m_{02}	m_{03}	m_{04}	m_{05}
0	5	5	10	10	3
1	4	5	10	10	4
2	4	4	10	10	5
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
8	3	3	8	8	11
9	2	3	8	8	12

The algorithm stops at the ninth iteration, where a compatible initial design with $m_{05} = 12$ is attained.

Step 2a of Algorithm 10.2 reverses Algorithm 10.1 and yields a faster initial sequence at each iteration; and hence the pruned sequence will be less conservative than the base- b benchmark. As a consequence of Theorem 10.1 shown later, the pruned sequence will also be compatible.

Pruning may sometimes lead to an unacceptably aggressive initial design. Should this be the case, an appropriate approach is to increase the sample size N . That is, we are to use (10.5) as a sample size constraint. In practice, it may be useful to produce several pruned compatible designs for a range of N s and a given m^* , solicit the most feasible design from a clinical perspective, and confirm the design's performance by simulation. In Example 10.3, if $N = 42$ is a feasible sample size, the unpruned initial design with $m_{01} = m_{02} = 5, m_{03} = m_{04} = 10$, and $m_{05} = 12$ may also be presented along with the pruned sequence, so that the clinical investigator may weigh in with the tangible trade-off between sample size and conservatism.

10.4.2 Dose Insertion

Dose insertion in phase I trials is a common idea among practitioners. Take the NeuSTART, for example. The investigators had suggested inserting an intermediate dose at 7 mg/kg/day between dose tier 3 (6 mg/kg/day) and dose tier 4 (8 mg/kg/day) in case toxicities are observed at dose tier 4. The motivation was to test a dose higher than 6 mg/kg/day (dose tier 3) if a deescalation from dose tier 4 was needed as a result. (This suggestion was not implemented in the study from a pharmacological perspective; cf., Section 7.2.2.)

From a statistical viewpoint, the cleanest way to handle dose insertion is to avoid any post hoc addition. That is, the suggested intermediate dose or doses should be included for potential testing in the planning stage. In the NeuSTART, if the dose 7 mg/kg/day is to be used, then one should plan a trial with $K = 6$ dose levels; the revised dose tiers will then be

Dose tier, k	1	2	3	4	5	6
Lovastatin dose (mg/kg/day)	1	3	6	7	8	10

To reflect that 7 mg/kg/day (dose tier 4) is inserted, we may prescribe an initial design that skips the dose in a two-stage CRM, that is, the initial cohort size $m_{04} = 0$. The calibration of the initial design can go through the similar process via Algorithms 10.1 and 10.2 otherwise.

A simpler alternative is to use the reference initial design in Table 10.3. For the revised dose tiers for NeuSTART, we may use the initial design corresponding to the CRM model for $\theta = 0.10, K = 6$, and $\delta = 0.02$ from the table, but set $m_{04} = 0$, that is,

$$m_{01} = m_{02} = m_{03} = 5, m_{04} = 0, m_{05} = 10. \quad (10.6)$$

Theorem 10.1. *Let $\mathbf{m}_0 = \{m_{0k}\}$ and $\tilde{\mathbf{m}}_0 = \{\tilde{m}_{0k}\}$ denote the initial cohort sizes of two sequences \mathbf{x}_0 and $\tilde{\mathbf{x}}_0$, respectively, where $m_{0k} \leq \tilde{m}_{0k}$ for all k and (10.1) and (10.2) are satisfied. If $\tilde{\mathbf{x}}_0$ is a compatible initial design for a given CRM model, then \mathbf{x}_0 is also compatible.*

Theorem 10.1 guarantees that (10.6) is a compatible initial design because the reference design obtained from Table 10.3 is compatible. Furthermore, suppose we impose the sample size constraint $N = 33$ and $m^* = 12$, the initial design (10.6) can be pruned according to Algorithm 10.2:

$$m_{01} = m_{02} = m_{03} = 4, m_{04} = 0, m_{05} = 9, m_{06} = 12. \tag{10.7}$$

Also as a consequence of Theorem 10.1, the pruned sequence (10.7) is compatible.

10.5 Case Study: NeuSTART

The NeuSTART adopted a two-stage CRM with $\theta = 0.10, K = 5$, and $v_0 = 3$. The empiric (3.2) function

$$F(d_k, \beta) = d_k^{\exp(\beta)} \tag{10.8}$$

was used to model the toxicity probability at dose level k with labels $d_1 = 0.02, d_2 = 0.06, d_3 = 0.10, d_4 = 0.18$, and $d_5 = 0.30$, where $\beta \sim N(0, 1.34)$ a priori. The initial dose sequence is given in (7.3). The calibration process of this CRM design is described in Section 7.4.2, and is characterized as a trial-and-error approach. In the following, we calibrate the initial design by Algorithms 10.1 and 10.2 for this study model (10.8). For comparison purposes, we set $N = 33$ and $m^* = 12$ (although in retrospect, we should have set $m^* = 15$). First, Table 10.4 displays the base- b benchmarks for all possible b so that $m_{01} > 0$. Of all sequences, the base-1 design is most conservative according to the total order \succ . Next, pruning the base-1 design with $N = 33$ and $m^* = 12$, we obtain

$$m_{01} = 4, m_{02} = 5, m_{03} = m_{04} = 6, \text{ and } m_{05} = 12. \tag{10.9}$$

It can be easily verified that the pruned design (10.9) is slower than the original initial sequence (7.3).

Table 10.4 Base- b initial sequences for the NeuSTART model

b	m_{01}	m_{02}	m_{03}	m_{04}	$m_{+,2}$	$m_{+,3}$	$m_{+,4}$
1	7	7	8	8	14	22	30
2	6	6	8	8	12	20	28
3	6	6	9	9	12	21	30
4	4	8	8	8	12	20	28
5	5	5	10	10	10	20	30
6	6	6	6	12	12	18	30
7	7	7	7	7	14	21	28

Suppose we opt to adopt the logistic model (10.4) instead of the empiric model in NeuSTART. We may choose $\delta = 0.02$ or 0.03 as recommended in Table 10.3. Further suppose we use $\delta = 0.03$. Then the pruned initial sequence is given in Example 10.3, as follows:

$$m_{01} = 2, m_{02} = 3, m_{03} = m_{04} = 8, \text{ and } m_{05} = 12. \tag{10.10}$$

This initial dose sequence is faster than (10.9), but is more conservative than the original (7.3), that is, $(7.3) \succ (10.10) \succ (10.9)$. For comparison, Table 10.5 shows the simulation results of two-stage CRM using these three initial dose sequences, along with that of the nonparametric optimal design. The three CRM designs have comparable accuracy. The logistic model with (10.10) has slightly larger average PCS over the five scenarios—0.65 versus 0.63 for the other two—but it also has more varied PCS than the other two. In comparison, the average PCS of the nonparametric optimal design is 0.76.

Table 10.5 The distribution of dose selection, average toxicity number (ATN), and overdose number (OD) of three two-stage CRM for the NeuSTART with $\theta = 0.25$, $K = 5$, $v_0 = 3$, $N = 33$, and $m^* = 12$

Model	\mathbf{x}_0	Probability of selecting dose				ATN	OD
Scenario 1:		.10	.25	.30	.35	.40	
NP optimal	—	.92	.07	.01	.00	.00	—
Empiric	(7.3)	.88	.11	.01	.00	.00	4.6
Empiric	(10.9)	.88	.11	.01	.00	.00	4.5
Logistic	(10.10)	.87	.12	.01	.00	.00	4.7
Scenario 2:		.04	.10	.25	.30	.35	
NP optimal	—	.23	.70	.07	.01	.00	—
Empiric	(7.3)	.32	.53	.14	.01	.00	3.6
Empiric	(10.9)	.32	.52	.15	.01	.00	3.5
Logistic	(10.10)	.32	.54	.13	.02	.00	3.7
Scenario 3:		.01	.04	.10	.25	.30	
NP optimal	—	.01	.22	.70	.07	.01	—
Empiric	(7.3)	.02	.27	.56	.14	.01	3.1
Empiric	(10.9)	.02	.26	.56	.13	.02	3.0
Logistic	(10.10)	.02	.26	.55	.14	.03	3.3
Scenario 4:		.01	.01	.04	.10	.25	
NP optimal	—	.00	.01	.22	.70	.08	—
Empiric	(7.3)	.00	.03	.25	.56	.16	2.4
Empiric	(10.9)	.00	.03	.26	.51	.19	2.3
Logistic	(10.10)	.00	.03	.25	.49	.23	2.6
Scenario 5:		.01	.01	.01	.04	.10	
NP optimal	—	.00	.00	.01	.22	.77	—
Empiric	(7.3)	.00	.01	.05	.28	.66	1.4
Empiric	(10.9)	.00	.01	.06	.26	.67	1.3
Logistic	(10.10)	.00	.01	.03	.18	.78	1.5

Note: The results for the nonparametric optimal design (NP optimal) are given as references. Numbers associated with the MTD are in bold.

The logistic model causes more toxicity and overdose than the empiric, but the difference is small. All three CRM designs lead to much reduced OD numbers when

compared to balanced randomization. This illustrates the conservative nature of the two-stage CRM design.

As we have emphasized in the previous chapters, the CRM designs obtained by the automated calibration approach give comparable operating characteristics to the original labor-intensive NeuSTART design. In particular, Table 10.3 provides a quick start to calibrate and prune a reasonable initial design for a two-stage CRM.

10.6 Exercises and Further Results

Exercise 10.1. Verify that $\mathbf{x} > \tilde{\mathbf{x}}$ implies $\mathbf{x} \succ \tilde{\mathbf{x}}$.

Exercise 10.2. Redesign the NeuSTART as in Section 10.5. For the logistic model (10.4) with $\delta = 0.02$, give an initial sequence under the constraint $N = 33$ and $m^* = 12$. Use simulation to compare this two-stage CRM with the designs in Table 10.5.

Exercise 10.3. Using Theorem 10.1, verify that pruning using Algorithm 10.2 will result in a compatible sequence.

Part III

CRM and Beyond

The Time-to-Event CRM

11.1 Introduction

In the bortezomib trial, each patient would receive up to six 21-day cycles of chemotherapy, and would thus be at risk of toxicity throughout the entire treatment period. If each patient was to be followed for 126 days before a dose decision was made for the next patient, the trial duration would be impractically long. On the other hand, the common phase I practice would usually count toxicity occurring only in the first cycle, and would underestimate the incidence of adverse events as toxicity might occur at a later cycle. To deal with the prospect of *late toxicity*, the bortezomib trial adopted a variant of the CRM, called the time-to-event CRM (TITE-CRM) that would allow continual accrual throughout the trial while using the 126-day toxicity endpoint as basis of dose escalation.

The basic idea of the TITE-CRM (read “tight C-R-M”) is to estimate the dose-toxicity curve based on the current toxicity status of all patients, including those who have not received a full six-cycle treatment. Section 11.2 makes this idea concrete and presents the mathematical framework of the method. Section 11.3 works out a numerical example to illustrate the implementation of the method, and Section 11.4 discusses the practical issue of enrollment scheduling. The theoretical aspects of the method are outlined in Section 11.5. Two-stage design with the TITE-CRM will be described in Section 11.6, and applied to the Poly E trial introduced in Example 2.2. The chapter is concluded in Section 11.7 with a brief review of the extension of the TITE-CRM in the literature, and alternatives to deal with late toxicities.

11.2 The Basic Approach

11.2.1 A Weighted Likelihood

A one-stage Bayesian TITE-CRM adopts the same dose assignment strategy as the regular CRM, and treats a newly enrolled patient at the model-based MTD estimate. The main difference is that the TITE-CRM uses information from partially followed patients in addition to completely followed patients. For example, suppose that a patient has been on treatment in the bortezomib trial for five cycles without any sign of toxicity. Although we have not observed the patient’s outcome in the last cycle, a five-cycle toxicity-free follow-up speaks for the safety of the dose that the patient receives, and in some sense may be counted as five-sixths of a nontoxic outcome.

This idea is made concrete through the use of a weighted likelihood of the model parameter, defined as

$$L_n(\beta; \mathbf{w}) = \prod_{i=1}^n \{w_{i,n+1} F(x_i, \beta)\}^{Y_{i,n+1}} \{1 - w_{i,n+1} F(x_i, \beta)\}^{1-Y_{i,n+1}} \quad (11.1)$$

where $Y_{i,n+1}$ and $w_{i,n+1}$ are, respectively, the toxicity indicator for patient i and the weight assigned to this observation just prior to the entry of patient $n+1$. Note that $Y_{i,n+1}$ is an increasing process in n and $Y_{i,n} = 1$ implies $Y_{i,n'} = 1$ for every $n' > n$. Generally, the eventual toxicity status $Y_{i,N+1}$ is the same as Y_i in the regular CRM. (Recall N denotes the sample size of the trial.) As patient $n+1$ becomes available for the trial, the parameter β is estimated by

$$\hat{\beta}_n^w = \frac{\int_{-\infty}^{\infty} \beta L_n(\beta; \mathbf{w}) dG(\beta)}{\int_{-\infty}^{\infty} L_n(\beta; \mathbf{w}) dG(\beta)}$$

so that the patient will receive the model-based MTD estimate

$$x_{n+1} = \arg \min_{d_k} |F(d_k, \hat{\beta}_n^w) - \theta|. \quad (11.2)$$

Alternatively, one may estimate β by maximizing the weighted likelihood. Let

$$\tilde{\beta}_n^w = \arg \max_{\beta} L_n(\beta; \mathbf{w})$$

denote the maximum weighted likelihood estimate, which exists if $Y_{j,n+1} = 0$ and $Y_{j',n+1} = 1$ for some $j, j' \leq n$. A two-stage strategy is needed when the maximum likelihood TITE-CRM is used, as is the case for the regular CRM. Once there is heterogeneity in the toxicity outcomes among patients, the dose assignment for patient $n+1$ will be:

$$x_{n+1} = \arg \min_{d_k} |F(d_k, \tilde{\beta}_n^w) - \theta|. \quad (11.3)$$

Before there is any toxic outcome in the trial, one will either treat the next patient according to a predetermined dose sequence $\{x_{i,0}\}$, or defer enrollment. See Section 11.6 for further discussion on the specification of two-stage TITE-CRM.

11.2.2 Weight Functions

The TITE-CRM incorporates the times-to-event of the patients through the weights in the weighted likelihood. Intuitively, the weight $w_{i,n+1}$ indicates the amount of information contributed to the likelihood by patient i , and is to be defined such that it increases with the patient's length of follow-up and that $w_{i,n+1} = 1$ if the patient has a complete follow-up.

In the case where a new patient is enrolled only when all current patients have been completely followed, the weighted likelihood reduces to the regular binomial likelihood, and the TITE-CRM to the regular CRM. More generally, if the weights, $w_{i,n+1}$ s, are determined in a manner that is independent of β , the weights associated

with the toxic outcomes can be factored out as a proportional constant in the weighted likelihood and do not affect the estimation of β , and hence dose assignment for the next patient. Thus, we may set $w_{i,n+1} = 1$ if patient i has experienced a toxic outcome. It is indeed intuitive to regard an observed toxic outcome as a complete observation irrespective of the length of follow-up.

To study the weighting scheme more rigorously, let T_i denote the time-to-toxicity of patient i , and T be the maximum length of follow-up for each patient, called the observation window. The observation window should be prespecified in the clinical protocol, so that a patient is said to experience a toxicity (i.e., $Y_i = 1$) if and only if $T_i \leq T$. For example, the bortezomib trial has an observation window of $T = 126$ days. Then by the definition of conditional probability, for $t \leq T$, we can write

$$\Pr(T_i \leq t \mid x_i) = \Pr(T_i \leq t \mid T_i \leq T, x_i) \Pr(T_i \leq T \mid x_i). \quad (11.4)$$

In the regular CRM, the marginal distribution $\Pr(T_i \leq T \mid x_i)$ is modeled through the dose-toxicity function $F(x_i, \beta)$. In the TITE-CRM, the conditional distribution $\Pr(T_i \leq t \mid T_i \leq T, x_i)$ is approximated by weight function $w(t; T, x_i)$. Therefore, the likelihood function based on current toxicity status can be written as

$$\begin{aligned} & \prod_{i=1}^n \{ \Pr(T_i \leq C_{i,n+1} \mid x_i) \}^{Y_{i,n+1}} \{ 1 - \Pr(T_i \leq C_{i,n+1} \mid x_i) \}^{1-Y_{i,n+1}} \\ & = \prod_{i=1}^n \{ w(C_{i,n+1}; T, x_i) F(x_i, \beta) \}^{Y_{i,n+1}} \{ 1 - w(C_{i,n+1}; T, x_i) F(x_i, \beta) \}^{1-Y_{i,n+1}} \end{aligned} \quad (11.5)$$

where $C_{i,n+1}$ is the length of follow-up of patient i just prior to the entry time of patient $n+1$. Technically, this likelihood is valid if the length of follow-up $C_{i,n+1}$ is independent of the time-to-toxicity T_i for all i . This assumption is realistic in phase I trials, where the censoring times are essentially administrative and are determined by the arrival times of new patients. Now, by comparing the weighted likelihood (11.1) and the current status likelihood (11.5), we see that the weight $w_{i,n+1}$ in the former may be viewed as an estimate of the weight function $w(C_{i,n+1}; T, x_i)$ in the latter. The weighted likelihood may thus be viewed as an approximate likelihood based on the current status data.

A practical implication of the conditioning argument (11.4) is that dose-toxicity modeling and time-to-toxicity modeling can be done separately. Thus, the calibration approach for $F_k(\beta)$ and the other model parameters can be done as outlined in the previous chapters. To model the conditional time-to-toxicity distribution, Cheung and Chappell [20] consider a simple linear weighting scheme that assigns weight as proportional to the length of follow-up. With respect to the constraints due to the probabilistic interpretation of $w(t; T, x_i)$, that is,

$$w(0; T, x_i) = 0 \text{ and } w(t; T, x_i) = 1 \text{ for } t \geq T,$$

the linear weight function is defined as

$$w(t; T, x_i) = \min\left(\frac{t}{T}, 1\right). \quad (11.6)$$

The linear weight (11.6) appears to be an oversimplistic estimate for the conditional distribution $\Pr(T_i \leq t \mid T_i \leq T, x_i)$ as it does not use data to approximate the shape of the distribution and assumes identical shape for all dose levels. Alternatively, one may use an adaptive weighting scheme in which the shape of the weight function is altered based on the accrued observations. In particular, with n enrolled patients, Cheung and Chappell [20] consider

$$w_n(t; T, x_i) = \frac{\kappa(t)}{z+1} + \frac{1}{z+1} \left\{ \frac{t - t_{(\kappa(t))}}{t_{(\kappa(t)+1)} - t_{(\kappa(t))}} \right\} \text{ for } t \leq T, \quad (11.7)$$

where z is the total number of toxic outcomes, the times-to-toxicity $t_{(j)}$ s are ordered such that $0 \equiv t_{(0)} < t_{(1)} \leq \dots \leq t_{(z)} < t_{(z+1)} \equiv T$, and $\kappa(t) = \max\{j \in [0, z] : t \geq t_{(j)}\}$. When there is no toxic outcome, that is, $z = 0$, the adaptive weight function (11.7) is identical to the linear weight function.

The functional form of w_n in (11.7) is adaptive to the data. However, it prescribes the same weight for patients at different doses as long as they have the same length of follow-up. Cheung [17] studies dose-adjusted weight that formally estimates the distribution of T_i given x_i under the assumption that toxicity occurs more rapidly at higher doses, that is, stochastic ordering of the time-to-toxicity distributions, that is,

$$\Pr(T_i \leq t \mid x_i) \leq \Pr(T_{i'} \leq t \mid x_{i'}) \text{ for } x_i < x_{i'}.$$

This regression model is fairly general. On the other hand, the computation of the dose-adjusted weights is complicated and requires the iterative pool-adjacent violator algorithm [3]. In addition, a moderate number of observations at each dose is needed to precisely estimate the dose-adjusted conditional distribution $\Pr(T_i \leq t \mid T_i \leq T, x_i)$. To be more parsimonious, one may estimate the conditional distribution of T_i for a given dose by parametric regression such as the accelerated failure time model with a Weibull distribution [20] and the cure model [12].

To summarize, the linear weight (11.6) is found to be adequate in most situations. Using an adaptive weighting scheme generally does not improve the accuracy of the TITE-CRM, but rather affects the in-trial allocation in some specific situations. When the late-onset tendency of toxicity is prominent (e.g., radiation therapy), using adaptive weights such as (11.7) will reduce the exposure of patients to overly toxic doses. Additionally, an adaptive weight scheme should be used in situations where we expect nonnegligible amount of patient dropout; see Theorems 11.2 and 11.3 in Section 11.5.

11.2.3 Individual Toxicity Risks

A derived concept from the time-to-toxicity distribution is a patient's risk for toxicity. Precisely, define the toxicity risk for patient i at time t , for $t < T$, as

$$r_i(t) = \Pr(T_i \leq T \mid T_i > t, x_i). \quad (11.8)$$

The toxicity risk (11.8) measures the likelihood of toxicity for patient i in the remaining treatment period, provided that the patient is toxicity-free up to time t . Some

algebraic manipulation will further give

$$r_i(t) = \frac{\{1 - \Pr(T_i \leq t \mid T_i \leq T, x_i)\} \Pr(T_i \leq T \mid x_i)}{1 - \Pr(T_i \leq t \mid x_i)},$$

from which it is clear that $r_i(0) = \Pr(T_i \leq T \mid x_i)$, and $r_i(t) \rightarrow 0$ as $t \rightarrow T$. Furthermore, it can be shown that $r_i(t)$ is a decreasing function of t . That is, the longer a patient has been followed without observing toxicity, the lower risk for toxicity the patient has.

As in (11.4), we can model $r_i(t)$ through the marginal model $F(x_i, \beta)$ and the weight function $w(t; T, x_i)$. In particular, just prior to the entry of patient $n + 1$, the risk of patient i , $r_i(C_{i,n+1})$, can be approximated by

$$r_{i,n+1} = \frac{(1 - w_{i,n+1})F(x_i, \hat{\beta}_{comp})}{1 - w_{i,n+1}F(x_i, \hat{\beta}_{comp})} \tag{11.9}$$

where $\hat{\beta}_{comp}$ is the posterior mean of β given the complete data, which is denoted by $H_{comp} = \{(x_j, Y_{j,n+1}) : w_{j,n+1} = 1\}$. In words, the complete data history H_{comp} contains the observations from patients either with complete follow-up or with a toxic outcome. The risk estimate (11.9) is applied to patients who are toxicity-free in the interim. If a toxic outcome has already been observed in patient i , risk becomes a certainty, and we will set $r_{i,n+1} = 1$.

11.3 Numerical Illustration

11.3.1 The Bortezomib Trial

To recall, the bortezomib trial used the TITE-CRM for dose escalation. The dose-toxicity relationship was modeled via the empiric function (3.2) with dose labels $d_1 = 0.05, d_2 = 0.12, d_3 = 0.25, d_4 = 0.40$, and $d_5 = 0.55$ and $\beta \sim N(0, 1.34)$. The calibration details can be found in Chapter 7, Section 7.4.1.

To incorporate the observed times-to-toxicity into the estimation of β , the linear weight function (11.6) was used to compute the weighted likelihood. For example, Table 11.1 displays the current status of the first 4 patients in the trial on the entry date of patient 5. Since the observation window in the trial is 126 days, patient 1 with 73-day follow-up has a weight of $73/126 = 0.58$; the weights of the other patients can be computed in the same manner. Since there has been no toxic outcome and all patients receive dose level 3, the weighted likelihood is

$$\prod_{i=1}^4 \{1 - w_{i,5} \times 0.25^{\exp(\beta)}\}$$

where $w_{i,5}$ s are given in Table 11.1. As a result, the posterior mean $\hat{\beta}_5^w = 0.49$, and the recommended dose by the TITE-CRM at this time is dose level 4. See the next subsection for the R code that produces these numerical results.

The toxicity risks for the patients are estimated using (11.9). Consider patient 1

who receives dose 3 and has an assigned weight 0.58. His risk for toxicity in the remaining treatment period is approximated by

$$r_{1,5} = \frac{(1 - w_{1,5}) F(d_3, \hat{\beta}_{comp})}{1 - w_{1,5} F(d_3, \hat{\beta}_{comp})} = \frac{(1 - 0.58) \times 0.25}{1 - 0.58 \times 0.25} = 0.123,$$

where $\hat{\beta}_{comp} = 0$; the prior mean of β is used because no patient has been completely followed at this time and H_{comp} is an empty set.

Table 11.1 The current status of the first 4 patients in the bortezomib trial on the entry date of patient 5 (April 16, 2004), and the estimates of their individual risks of toxicity

i	x_i	Entry date	$Y_{i,5}$	$C_{i,5}$	$w_{i,5}$	$r_{i,5}$
1	3	February 3, 2004	0	73	0.58	0.123
2	3	February 10, 2004	0	66	0.52	0.137
3	3	March 12, 2004	0	35	0.28	0.194
4	3	March 19, 2004	0	28	0.22	0.206

11.3.2 Implementation in R

The TITE-CRM is implemented by the R function `titecrm` in the ‘`dfrm`’ package:

```
> ### Determine dose for patient 5 in the bortezomib trial
> p0 <- c(0.05,0.12,0.25,0.40,0.55) # initial guesses
> theta <- 0.25
> y <- c(0,0,0,0) # current toxicity status
> x <- c(3,3,3,3) # dose level
> cc <- c(73,66,35,28) # lengths of follow-up (in days)
> wt <- cc / 126 # linear weights
> foo <- titecrm(p0,theta,y,x,weights=wt)
> foo$estimate # posterior mean of beta
[1] 0.4907791
> round(foo$ptox,digits=2) # posterior toxicity probability
[1] 0.01 0.05 0.14 0.27 0.42
> foo$mtd # model-based MTD
[1] 4
>
```

There are various ways to enter the weight information to the function `titecrm`. The most straightforward approach is to use the argument `weight` as in the above code. Alternatively, one could enter the patients’ lengths of follow-up via the argument `followup`, along with the observation window via the argument `obswin`. Then the weights of the patients will be computed according to the linear weight (11.6) or the adaptive weight (11.7).

```
> ### The following gives an equivalent object to ‘foo’
```

```
> foo2 <- titecrm(p0,theta,y,x, followup=cc, obswin=126)
> foo2$weights
[1] 0.5793651 0.5238095 0.2777778 0.2222222
> foo2$estimate      # posterior mean of beta
[1] 0.4907791
>
```

The default weighting scheme used by `titecrm` is the linear weight (11.6). When the argument `scheme` is specified as “adaptive,” the adaptive weight function (11.7) will be used. In this example, since there is no toxic outcome, the adaptive weights will be identical to the linear weights.

In addition, the R function `titesim` can be used to simulate TITE-CRM trials under a given dose–toxicity curve. Exercise 11.1 demonstrates the basic usage of the function as an example.

11.4 Enrollment Scheduling

11.4.1 Patient Accrual

The data in Table 11.1 led to the first escalation in the bortezomib trial after the first 4 patients had been followed for a total of 202 patient-days without any sign of toxicity; cf., Figure 1.1. To compare, a CRM with the same dose–toxicity model and prior would escalate dose after a complete nontoxic outcome (i.e., 126 patient-days). This illustrates that the TITE-CRM is a reasonably cautious dose escalation scheme and can be applied with a fully sequential enrollment plan without causing delays in enrollment.

On the other hand, one will need to exercise caution if the patient accrual rate is rapid. Figure 11.1 shows the operating characteristics of the TITE-CRM design in the bortezomib trial under the five dose–toxicity scenarios in Table 7.1 with various patient accrual rate ξ . In each simulated trial, the interpatient arrival times were fixed so that ξ patients would be enrolled per observation window; for example, for $\xi = 3$ with an observation window of 126 days, patient 1 would be enrolled on day 42, patient 2 day 84, patient 3 day 126, and so on. The regular CRM was a special case corresponding to $\xi = 1$; that is, a new patient would arrive only after all previous patients were completely followed. The conditional uniform model was used to generate the times-to-toxicity in the simulations [20]: we first determined if a patient had a toxic response; if so, we would generate a number on the interval $(0,126)$ uniformly as the time-to-toxicity.

The figure shows general trends that the PCS declines as ξ increases, while the toxicity and overdose numbers increase. These trends indicate the potential problem of erroneous escalations due to inadequate follow-up of individual patients: if many patients are toxicity-free after very limited follow-up, their aggregate effects on the weighted likelihood may occasion an escalation. Thus, the TITE-CRM will expose an increased number of patients at toxic doses before any toxicity is observed.

However, the impacts on the operating characteristics are gradual when compared the reduction in trial duration shown in Figure 11.1d. The bortezomib trial took 406

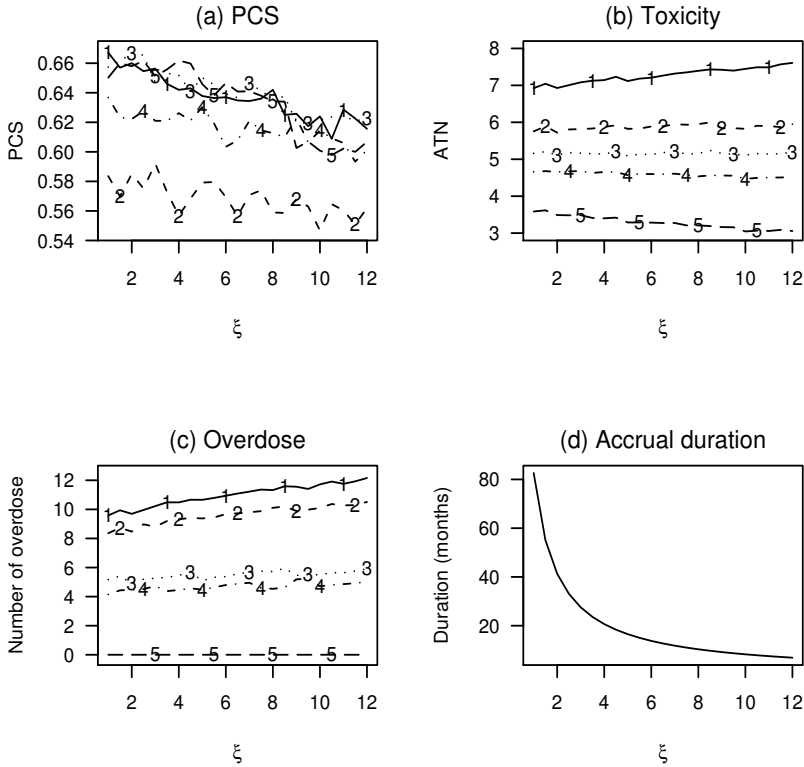


Figure 11.1 Operating characteristics of the TITE-CRM against patient accrual rate ξ under the dose–toxicity scenarios in Table 7.1. The numbers in figures (a), (b), and (c) indicate the simulation scenarios.

days to enroll 20 patients (i.e., roughly 6.2 patients per 126 days). Thus, we may retrospectively compare the TITE-CRM with $\xi = 6.5$ and the CRM ($\xi = 1$): the PCS drops by 1–3 percentage points, the increase in the average number of toxicities is no greater than 0.4 per trial, and the number of overdose under scenarios 1–4 are up by 0.7–1.5 per trial. In comparison, there is a sixfold difference in the accrual duration: the TITE-CRM reduces the accrual period from 7 years to 14 months as shown in Figure 11.1d. Also, if a group accrual CRM with size 3 is used, the accrual duration for $N = 20$ will be roughly 3 years if we assume patients are immediately available. While this reduction is substantial, a 3-year long phase I trial with a small sample size may prove infeasible in many clinical settings where there are other extraneous and practical considerations.

11.4.2 *Interim Suspensions*

When patient availability is quick relative to the observation window, we may close the study temporarily as a safety measure, so that dose decisions for future patients can be based on adequate follow-up of the current patients. In practice, the suspension decisions may be made through the mechanism with a data safety and monitoring committee (DSMC). However, it will be reassuring if the investigators can state the monitoring rules in the study protocol. This section presents an interim monitoring approach used in conjunction with the TITE-CRM. In brief, the basic idea is to invoke accrual suspension if there is a nonnegligible likelihood that the TITE-CRM will treat the next patient at an overdose as a result of not having the full follow-up data. This process entails the enumeration of all possible outcome sequences should all the patients be completely followed, and the quantification of the likelihood of each sequence. The concrete steps are carried out as follows:

Algorithm 11.1. Interim risk monitoring in the TITE-CRM

At an interim with n enrolled patients, let x_{n+1} denote the recommended dose by the TITE-CRM based on the observed data. Then the monitoring objective is to assess the risk of x_{n+1} using the following steps:

1. Enumerate all possible eventual outcome configurations given the observed data. Each configuration, denoted by $\mathbf{y}_{n+1}^f = (y_{1,n+1}^f, \dots, y_{n,n+1}^f)$, is called a full outcome set.
2. For each full outcome set, evaluate the recommended dose, denoted as x_{n+1}^f .
3. Approximate the *risk likelihood* of occurrence of each full outcome set based on the complete data H_{comp} .
4. Accumulate the risk likelihoods of all full outcome sets with $x_{n+1}^f < x_{n+1}$. This accumulation measures the likelihood of selecting a lower dose than x_{n+1} when all patients are completely followed. Thus, suspend the trial if this likelihood is “large” such as when ≥ 0.50 .

To illustrate these steps, consider the bortezomib data in Table 11.1. As 4 patients were partially followed with no toxic outcome on April 16, 2004, there were $2^4 = 16$ possible full outcome sets. Table 11.2 lists each of the 16 outcome sets along with their recommended dose x_5^f for patient 5. This enumeration thus completes Steps 1 and 2 of Algorithm 11.1.

To calculate the likelihood of each full outcome (Step 3), we first determine the individual risks of toxicity as shown in Table 11.1. With these individual risks, the likelihood of each full outcome is computed as

$$\prod_{i=1}^n (r_{i,n+1})^{y_{i,n+1}^f} (1 - r_{i,n+1})^{1-y_{i,n+1}^f}.$$

For example, for the outcome $y_{1,5}^f = y_{2,5}^f = y_{3,5}^f = y_{4,5}^f = 0$ (row 1 in Table 11.2), the

Table 11.2 All possible full outcome sets in the first 4 patients in the bortezomib trial as of April 16, 2004 (see Table 11.1), and their corresponding recommended doses and likelihoods of occurrence

$y_{1,5}^f$	$y_{2,5}^f$	$y_{3,5}^f$	$y_{4,5}^f$	x_5^f	Likelihood
0	0	0	0	5	0.484
1	0	0	0	3	0.068
0	1	0	0	3	0.077
1	1	0	0	1	0.011
0	0	1	0	3	0.117
1	0	1	0	1	0.016
0	1	1	0	1	0.019
1	1	1	0	1	0.003
0	0	0	1	3	0.126
1	0	0	1	1	0.018
0	1	0	1	1	0.020
1	1	0	1	1	0.003
0	0	1	1	1	0.030
1	0	1	1	1	0.004
0	1	1	1	1	0.005
1	1	1	1	1	0.001

likelihood is

$$(1 - 0.123) \times (1 - 0.137) \times (1 - 0.194) \times (1 - 0.206) = 0.484.$$

The recommended x_5^f based on this full outcome set is dose level 5. Now, except for this particular outcome, the recommended doses based on the other full data are lower than what the TITE-CRM would choose (i.e., dose level 4) at this interim. It thus seems quite likely that the TITE-CRM is recommending a dose higher than the MTD estimate should we wait long enough for all patients to be completely followed. The likelihood of this event is $0.516 = 1 - 0.484$, on which basis we may invoke a temporary accrual suspension.

While we did not implement Algorithm 11.1 or other monitoring rules in the bortezomib trial, this type of risk calculation could have been easily incorporated in the protocol. Generally, the calculations can be carried out on an ad hoc basis, or when the TITE-CRM recommends an escalation from the current patients. For the DSMC review purposes, it will be useful to submit the risk likelihood as auxiliary safety information, along with the method calculation. Furthermore, prespecifying a suspension threshold for the risk likelihood in Step 4 of the algorithm can simplify decisions during the trial.

Bekele et al. [8] propose a trial monitoring approach that uses Bayesian predictive risks of toxicity at prospective doses for future patients. The basic idea is similar to Algorithm 11.1: accrual will be suspended if the predicted risk is unacceptable high. Obtaining Bayesian predicted risks, however, requires special programming such as

Markov Chain Monte Carlo. Algorithm 11.1 is comparatively less computer intensive and can be easily programmed using the functions in the ‘dfcrm’ package.

11.5 Theoretical Properties

11.5.1 Real-Time Formulation

To study the theoretical properties of the TITE-CRM, it is useful to introduce notation to reconcile calendar time, t , and an individual patient’s time, s , where $s = 0$ indicates the time the patient enters the trial. Let e_i denote the calendar time when patient i enters the trial. For any calendar time t , the patient’s length of follow-up is $C_i(t) = \min\{\max(t - e_i, 0), T\}$. The weighted likelihood (11.1) can then be rewritten as

$$L(\beta, t) = \prod_{i=1}^{N(t)} \{w_i(t)F(x_i, \beta)\}^{Y_i(t)} \{1 - w_i(t)F(x_i, \beta)\}^{1-Y_i(t)} \quad (11.10)$$

where $Y_i(t) = I\{T_i \leq C_i(t)\}$ and $w_i(t) = w\{C_i(t); T, x_i\}$ are respectively the current toxicity status and the weight of patient i , and $N(t)$ is the number of enrollments just prior to calendar time t .

11.5.2 Real-Time Coherence

In summary, a coherent design does not allow escalation right after a toxic outcome, and deescalation after a nontoxic outcome. Real-time version of coherence is needed for the situations where toxicity may occur randomly over an extended observation window while patients are enrolled continually in a staggered fashion, such as that in the bortezomib trial. Naturally, a real-time version of coherence in deescalation does not allow deescalation from a time point t to another time point $t + h$ if no new toxicity is reported during the time period $[t, t + h)$. Mathematically, a dose finding design \mathcal{D} is coherent in deescalation if, with probability one,

$$P_{\mathcal{D}} \{U(t, h) < 0 \mid Y_i(t + h) - Y_i(t) = 0 \text{ for all } i\} = 0 \quad (11.11)$$

where $U(t, h) = \mathcal{D}(H_{t+h}) - \mathcal{D}(H_t)$, and $H_t \equiv \{(x_i, Y_i(t), T_i \wedge C_i(t)) : i \leq N(t)\}$ is the entire history of observations just prior to time t .

A design is called real-time coherent in escalation if it does not escalate at the instant immediately after a toxic observation: with probability one,

$$\lim_{h \rightarrow 0} P_{\mathcal{D}} \{U(t, h) > 0 \mid Y_i(t + h) - Y_i(t) = 1 \text{ for some } i\} = 0. \quad (11.12)$$

The conditions given by (11.11) and (11.12) imply (5.2) and (5.1), respectively, when the toxicity outcomes can be evaluated instantaneously. Although it may not be easy to interpret (11.12), the notion of real-time coherence provides an objective criterion to evaluate a real-time dose finding method.

Theorem 11.1. *If the weight function is defined such that $w_i(t)$ is continuous and nondecreasing in t for all i , then the one-stage Bayesian TITE-CRM is real-time coherent.*

It is easy to verify that the linear weight (11.6) satisfies the coherence condition in the theorem. For other adaptive weight functions such as (11.7), one can easily apply a restriction that each patient's weight is not to decrease at the next evaluation. The proof of Theorem 11.1 is given in Cheung [16].

11.5.3 Consistency

Cheung and Chappell [20] and Cheung [17] study the asymptotic properties of the maximum likelihood TITE-CRM under the regularity conditions on the dose–toxicity model stated in Chapter 4, Section 4.3, with additional assumptions about the weight function and the patient accrual process. These asymptotic results can be modified to suit the Bayesian TITE-CRM as in Chapter 5, Section 5.4 and are stated in the following:

Theorem 11.2. *Suppose that the conditions in Theorem 5.3 hold. If*

$$\sum_{i=1}^{N(t)} I\{C_i(t) < T\} / N(t) \rightarrow 0 \quad (11.13)$$

as $t \rightarrow \infty$ and the weight function does not depend on β , then the TITE-CRM will select dose d_l eventually with probability one, and $l = v$ is the true MTD.

Theorem 11.2 requires the same set of assumptions in regard to the dose–toxicity model $F_k(\beta)$ as in the CRM. A practical implication is that calibration of the dose–toxicity model for the TITE-CRM may be done using the same techniques for the CRM. For time-to-event modeling, Theorem 11.2 allows a broad spectrum of weight functions, provided that the number of incomplete observations is small when compared to the number of enrolled patients. The assumption (11.13) may be violated in situations where we expect a certain level of patient dropout, which can be a legitimate concern for trials with an extended follow-up period.

Theorem 11.3. *Suppose that the conditions in Theorem 5.3 hold. If the process $C_i(t)$ of patient i is conditionally independent of the eventual outcome Y_i (which may not be observed) given $Y_i(t)$, and that*

$$\sup_{s \in [0, T]} |w_{N(t)}(s; T, x_i) - \Pr(T_i \leq s \mid T_i \leq T, x_i)| \rightarrow 0 \quad (11.14)$$

as $t \rightarrow \infty$, the TITE-CRM will select dose d_l eventually with probability one, and $l = v$ is the true MTD.

The assumption in Theorem 11.3 amounts to a missing-at-random (MAR) dropout pattern [91]. While MAR is an untestable assumption, it has been implicitly used in dose finding trials. The assumption is reasonable albeit debatable—under one condition: if dose withdrawal is due to adverse events caused by the drug, this event should be defined as dose limiting even though the adverse events may not

have reached the severity level of a dose-limiting toxicity. By allowing a nonnegligible amount of patient dropout, the consistency requirement on the weight function in Theorem 11.3 becomes stricter than that in Theorem 11.2. In particular, the linear weight (11.6) does not generally meet the condition in Theorem 11.3, nor does the adaptive weight function (11.7) when the conditional distribution of T_i depends on dose x_i . However, with the typical small-to-moderate sample size in dose finding studies, the data may prove to be too sparse for any dose-adjusted weighting scheme. Extensive simulation shows that the use of adaptive function (11.7) is comparable to a dose-adjusted scheme when there is a moderate level (30%) of dropout, even though the assumption (11.14) is not met by the function (11.7) generally [17].

11.6 Two-Stage Design

11.6.1 Waiting Window

Like a two-stage CRM, a two-stage TITE-CRM is implemented in accordance with a prespecified dose sequence $x_{1,0} \leq \dots \leq x_{N,0}$ before any toxicity is seen. In addition, the initial design includes the minimum length of waiting time, called the waiting window τ , between cohorts before an escalation may take place in the initial stage. Thus, a two-stage TITE-CRM avoids aggressive dose assignments in two ways. First, it starts the trial at a low dose. Second, the waiting window allows time to pass before escalation.

Figure 11.2 shows a simulated trial of size 30 using a two-stage TITE-CRM with a group-of-three initial escalation rule starting at the lowest dose. There are $K = 3$ dose levels, and the trial objective is to identify a dose with a $\theta = 0.25$ toxicity probability during a 6-month period. In the simulation, the enrollment schedule is fixed at a rate of one patient per month (i.e., $\xi = 6$ per observation window), except during stage 1 where a 4-month waiting window is imposed between patients if an escalation is to take place (i.e., $\tau = 4$). For example, patient 4 in the figure is enrolled at 4 months after patient 3. There is another 4-month delay after the enrollment of patient 6. This delay allows the observation of the toxic outcome in patient 5 before enrolling patient 7; this toxic observation then triggers the transition to the model-based TITE-CRM, and patients are enrolled at the rate of 1 per month without further delay. As a result of the waiting window in the initial stage, it takes 36 months to recruit 30 subjects. In comparison, an uninterrupted accrual schedule will take 30 months to enroll.

This example illustrates the trade-off between timeliness and patient safety through accrual suspensions. As the risk calculations in Algorithm 11.1 use solely prior information before any toxicity is seen, the waiting window has a similar role in a two-stage design. However, the window τ is to be chosen explicitly as a clinician-input parameter that represents a duration in which toxicity is plausible and relatively probable.

Theorem 11.4. *Suppose the initial design $x_{1,0} \leq \dots \leq x_{N,0}$ are chosen such that a two-stage CRM is coherent per Theorem 5.1. A two-stage TITE-CRM using the same*




Figure 11.2 A simulation of a two-stage TITE-CRM trial using model (11.15) described in Section 11.6.2. Each number indicates a patient: A number marks the patient's entry time; a circled number indicates the time when a toxicity occurs. Vertical positions of some patients are jittered for visual clarity.

dose-toxicity model and initial design sequence is real-time coherent regardless of the waiting window τ .

Theorem 11.4 implies that we may specify the initial dose sequence $\{x_{i,0}\}$ using the guidelines in Chapter 10. The sketch of the proof of Theorem 11.4 is outlined in Exercise 11.3.

11.6.2 Case Study: The Poly E Trial

Chemoprevention involves the use of specific agents to reverse, suppress, or prevent carcinogenic progression to invasive cancer [65]. Chemoprevention is intended for many diverse populations: healthy individuals, individual high-risk for cancer, or cancer survivors who are at risk for recurrence or new primary tumors. These target populations are generally healthier than cancer patients. Also, since carcinogenesis is a long process, chemoprevention will be given over an extended period so that mild toxicities may become intolerable, and long-term effects concerning. Therefore, it is important to allow an adequate follow-up when evaluating a chemopreventive agent. These concerns are exemplified in the Poly E trial introduced in Example 2.2

in Chapter 2. This dose-escalation study tested $K = 3$ doses of Polyphenon E, 400, 600, or 800 mg taken orally per day over a 6-month period, in $N = 30$ women with a history of breast cancer. The lowest dose level 400 mg was chosen based on previous preclinical and clinical safety data, and would be the starting dose in the trial, that is, $x_1 = d_1$. The objective was to identify a dose causing $\theta = 0.25$ dose-limiting toxicity that was to occur during the 6-month treatment period. In this trial, a grade 2 or higher grade toxicity would be counted as dose limiting; this was in contrast to the commonly used grade 3 dose-limiting threshold in chemotherapy intervention. The lower dose-limiting threshold was used for feasibility reasons: mild toxicities such as headache over an extended period might prove intolerable to women who were at risk but lived without cancer, and thus prevent the subjects from achieving full compliance with the treatment.

The TITE-CRM was used to determine dose escalation. The linear weight (11.6) was used so that the weight of an observation was proportional to the length of follow-up. For dose–toxicity modeling, the empiric function

$$F(d_k, \beta) = d_k^{\exp(\beta)} \quad (11.15)$$

was used with $\beta \sim N(0, 1.34)$ and dose labels $d_1 = 0.15$, $d_2 = 0.25$, and $d_3 = 0.35$. The choice of this dose–toxicity model was made by trial-and-error, done in a similar manner as in the bortezomib trial and the NeuSTART; see Section 7.4. According to this model, the prior MTD $v_0 = 2$ and $x_1 < d_{v_0}$; thus, a two-stage design was adopted with a group-of-five initial design, that is,

$$x_{1,0} = \cdots = x_{5,0} = d_1, \quad x_{6,0} = \cdots = x_{10,0} = d_2, \quad x_{11,0} = \cdots = x_{30,0} = d_3.$$

This initial design can be shown to be incompatible with model (11.15) and may cause incoherent escalation: The concept of coherence had not been fully developed at the time the Poly E trial was planned; rather, it was decided by an ad hoc safety committee. However, there was a restriction that no escalation would be allowed if the previous patient had a toxic outcome. In other words, no coherent escalation was allowed.

In this section, we will focus on the use of model (11.15) in the context of the Poly E trial, and examine the effects of the initial design—the dose sequence $\{x_{i,0}\}$ and the waiting window τ —on the operating characteristics of the two-stage TITE-CRM. Now, applying Algorithm 10.1, we obtain the group-of-three initial dose sequence as a conservative base- b benchmark, that is,

$$x_{1,0} = x_{2,0} = x_{3,0} = d_1, \quad x_{4,0} = x_{5,0} = x_{6,0} = d_2, \quad x_{7,0} = \cdots = x_{30,0} = d_3.$$

Therefore, for the initial design, we will consider

- Initial dose sequence $x_{i,0}$ s: group-of-three versus group-of-five
- Waiting window τ : 1, 2, ..., 6 months.

For the two-stage design with group-of-five initial sequence, we enforce coherence by restriction.

Figure 11.3 shows the PCS and the safety indices of the TITE-CRM under the

scenario with $p_1 = 0.10$, $p_2 = 0.25$, and $p_3 = 0.40$, that is, dose level 2 is the MTD, and level 3 an overdose. The results are based on a fixed enrollment schedule with $\xi = 6$ patients per observation window (6 months). Overall, the impacts of the initial dose sequence and the waiting window τ are quite small, although some trends may be noted. First, the group-of-three initial design induces larger numbers of toxicity and overdose than the group-of-five design; this is expected because the former escalates faster than the latter. Second, the toxicity indices decline as τ increases; this is also expected because erroneous escalations may be avoided by making dose decisions based on longer follow-up data. However, all design specifications in the figure represent improved safety upon balanced randomization, which will treat about 10 patients at an overdose under this dose-toxicity scenario resulting in an average of 7.5 toxic outcomes. This being said, if the group-of-three sequence is used, it may be prudent to be cautious by setting $\tau = 3$ so as to dampen the overdose number; this will add about 3 months to the trial duration.

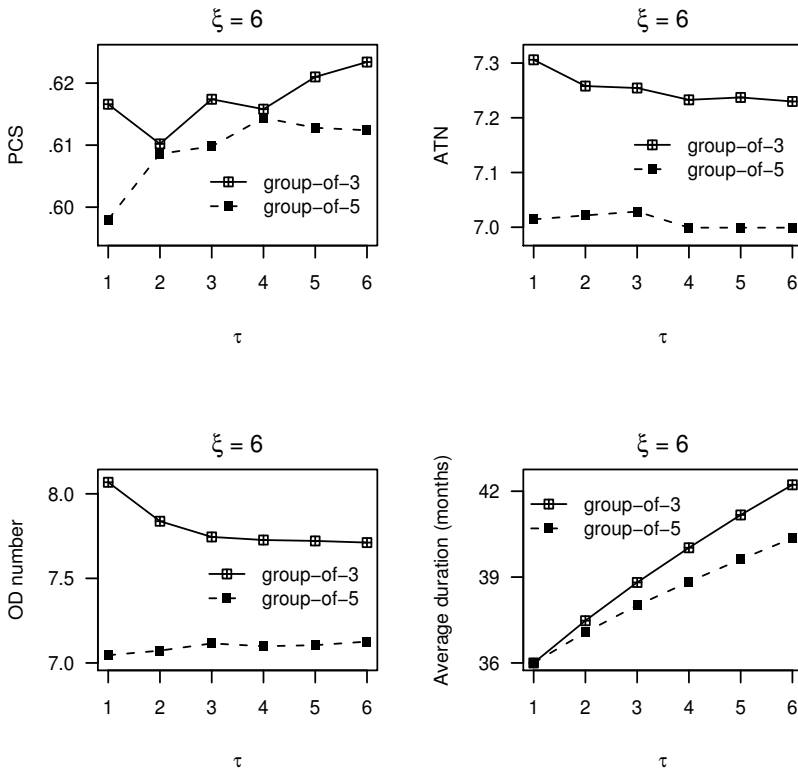


Figure 11.3 Operating characteristics of the two-stage TITE-CRM using model (11.15) and various initial design under the dose-toxicity curve: $p_1 = 0.10$, $p_2 = 0.25$, and $p_3 = 0.40$.

Figure 11.3 also shows that the group-of-three initial design leads to higher PCS but longer duration than the group-of-five design; but the difference is negligible. To give a sense about the impacts of the initial design, Figure 11.4 shows the PCS and duration of the TITE-CRM designs under another curve with $p_1 = p_2 = 0.10$ and $p_3 = 0.25$, that is, dose level 3 is the MTD, under two patient accrual rates at $\xi = 6, 12$ patients per observation window. While we still see that the group-of-three design is slightly superior in terms of PCS and inferior in terms of duration, the impacts due to patient accrual rate ξ are more noticeable. Also, for the group-of-five design, the impact of τ is much more pronounced with a fast accrual rate $\xi = 12$.

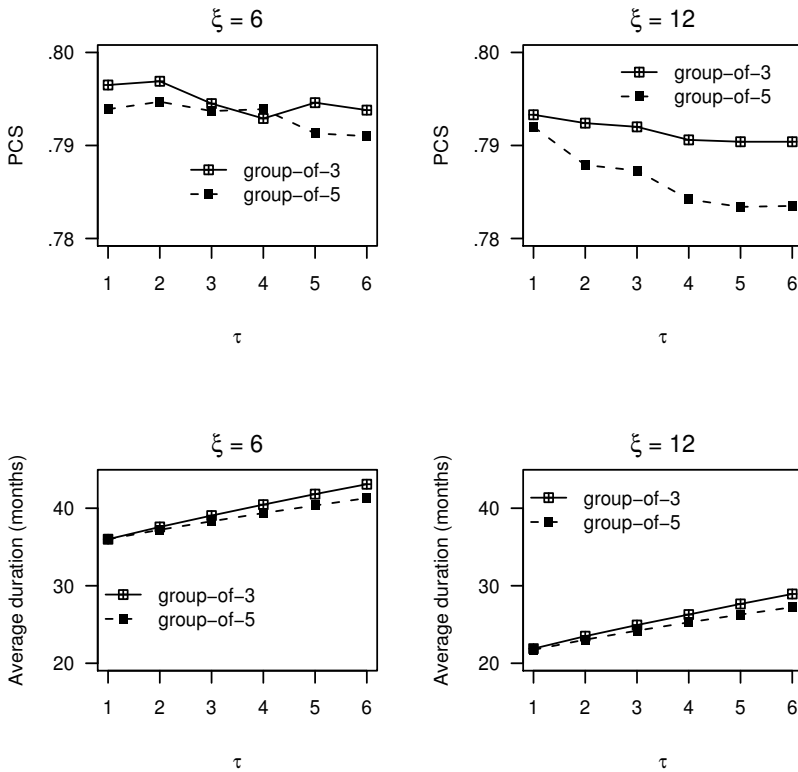


Figure 11.4 Operating characteristics of the two-stage TITE-CRM using model (11.15) and various initial design under dose-toxicity curve: $p_1 = p_2 = 0.10$ and $p_3 = 0.25$.

11.7 Bibliographic Notes

The problem of long trial duration of the CRM is noted by Goodman et al. [43] who propose the use of group accrual in the CRM. Subsequently, Thall et al. [103] study

a look-ahead strategy in conjunction with the CRM. The look-ahead strategy follows similar steps to Algorithm 11.1 by enumerating all possible eventual outcomes of the partially followed patients, and continues accrual only when all the projected recommended doses point to the same dose. Comparatively, duration reduction due to the look-ahead strategy is less substantial than the group accrual approach.

The time-to-event formulation of the MTD is first introduced by Cheung and Chappell [20], who also propose the TITE-CRM. The first documented TITE-CRM trial is published by Muler et al. [69]. Hüsing et al. [47] considers a time-to-event design analogous to the look-ahead extension of the 3+3 algorithm, whereby the times-to-toxicity are incorporated in dose decisions through an excess recruitment function.

There are several proposals on the weighting schemes directed to enhance the safety of the TITE-CRM under late-onset toxicities. Braun [12] obtains the weights through estimation under the cure model framework. Cheung [17] studies a dose-adjusted adaptive weight in the context of radiation therapy, which is expected to induce late toxicities. Cheung and Thall [24] propose a general class of adaptive weights similar to (11.7) and extend the use of the weighted likelihood to phase II clinical trial settings where the binary outcome is defined in terms of one or more time-to-event variables over a fixed observation window. A recent review of the TITE-CRM geared toward oncologists is given in Normolle and Lawrence [73].

11.8 Exercises and Further Results

Exercise 11.1. A TITE-CRM trial may be generated using the function `titesim` in the `'dfcrm'` package. For example, the following code generates the simulated trial in Figure 11.2:

```
> ### A simulated trial of two-stage TITE-CRM of 30 subjects
> PI <- c(0.05, 0.10, 0.25) # true toxicity probabilities
> p0 <- c(0.15, 0.25, 0.35) # initial guesses
> x0 <- c(rep(1,3),rep(2,3),rep(3,24)) # initial design
> N <- 30
> xi <- 6 # Patient accrual rate
> tau <- 4 # Waiting window
> foo <- titesim(PI,p0,0.25,N,x0,obswin=6,rate=xi,tgrp=tau)
> foo$MTD # MTD estimate
[1] 3
>
```

The argument `rate` specifies the patient accrual rate ξ per observation window. For a two-stage design, the waiting window τ in stage 1 is specified via `tgrp`, which should use the same time unit as the observation window `obswin`. Two options of patient accrual patterns are provided by the function via the specification of the argument `accrual`: setting `accrual="poisson"` generates random interpatient arrivals in accordance with the Poisson process, whereas `accrual="fixed"` (the default value) gives fixed patient enrollment. The other arguments of the function are the same as those in the function `titecrm`.

Use the function `titesim` to evaluate the operating characteristics of the Poly E design under various waiting window; cf., Section 11.6.2.

Exercise 11.2. Verify the toxicity risks of the patients given in Table 11.1.

Exercise 11.3 (sketch of proof of Theorem 11.4). Let $t_0 = \min\{t : Y_i(t) = 1\}$ be the time the first toxic outcome appears.

- a. Prove that $\lim_{h \rightarrow 0} P_{\mathcal{D}}\{U(t_0, h) > 0\} = 0$ implies Condition (11.12).
- b. Assume that $x_{i,0}$ is a nondecreasing sequence. For given $t_0, h > 0$, prove that $U(t_0, h)$ attains its maximum when $Y_{N(t_0)} = 1$ and $Y_i(t) = 0$ for $i < N(t_0)$, that is, the most recent patient experiences toxicity at time t_0 ,
- c. Suppose that $t_0 \in (e_n, e_{n+1}]$, that is, first toxicity occurs between the enrollments of patients n and $n + 1$. Prove that $U(t_0, h)$ attains its maximum when $t_0 = e_{n+1}$.
- d. Based on the above results, show that if an initial sequence is compatible with the regular CRM (i.e., TITE-CRM with $\xi = 1$), then it is compatible with the corresponding TITE-CRM regardless of waiting window.

CRM with Multiparameter Models

12.1 Introduction

We have seen in Chapter 5 that the one-parameter CRM is consistent under certain model misspecifications, but is not generally so (Theorem 5.3). In order to avoid bias due to model misspecification, one may be compelled to consider a more realistic dose–toxicity function in the CRM than a one-parameter function. In this chapter we discuss the use of multiparameter models in the CRM literature. Section 12.2 first reviews dose finding methods based on *nonparametric* estimation, the so-called curve-free methods; and Section 12.3 presents the potential problems associated with these methods and provides remedies. Section 12.4 explores the use of two-parameter CRM and illustrates its pitfalls through numerical examples. Section 12.5 provides a literature review on the use of these multiparameter models in the CRM.

12.2 Curve-Free Methods

12.2.1 The Basic Approach

The basic idea of curve-free methods is to treat the next group of patients at a dose with toxicity probability estimated to be closest to θ , that is,

$$x_{i+1} = \arg \min_{d_k} |\hat{\pi}_i(d_k) - \theta| \quad (12.1)$$

where $\hat{\pi}_i$ is a nonparametric isotonic estimate of the dose–toxicity curve π based on the first i groups of observations; that is, the form of $\hat{\pi}_i(x)$ is unspecified, except that it is monotone increasing. This modeling strategy in effect assumes a K -parameter dose–toxicity curve. Because there is not much borrowing strength from information across doses in the estimation of the dose–toxicity curve, a curve-free method usually requires treating a few patients at a dose before the next update, that is, a group accrual enrollment plan with $m > 1$. Therefore, we let x_i denote the dose assigned to the i th *group* of patients in this and the next section. In addition, in align with the tendency for being conservative, these designs usually start at the lowest dose d_1 . The main variations among the curve-free methods thus arise from the modeling and estimation approaches. The following subsection outlines some of the most common approaches.

12.2.2 Product-of-Beta Prior Distribution

Gasparini and Eisele [40] propose the use of Bayesian nonparametric estimates of the dose–toxicity curve. Specifically, the toxicity probabilities p_j s are reparameterized in terms of

$$\phi_1 = 1 - p_1 \text{ and } \phi_j = \frac{1 - p_j}{1 - p_{j-1}} \text{ for } j = 2, \dots, K \quad (12.2)$$

where each $\phi_j \sim \text{beta}(a_{1j}, a_{2j})$ and is independent of the others a priori. The prior mean and variance of ϕ_j are, respectively,

$$E(\phi_j) = \frac{a_{1j}}{a_{1j} + a_{2j}} \text{ and } \text{var}(\phi_j) = \frac{a_{1j}a_{2j}}{(a_{1j} + a_{2j})^2(a_{1j} + a_{2j} + 1)}.$$

Converting (12.2) gives

$$p_k = 1 - \prod_{j=1}^k \phi_j \quad (12.3)$$

so that the beta prior distribution of ϕ_j s imposes strict monotonicity, that is, $p_1 < p_2 < \dots < p_K$, without additional parametric assumptions. The prior on p_k s is said to follow the product-of-beta distribution [101].

Using (12.3), the binomial likelihood based on i groups of observations can be written as

$$\prod_{k=1}^K p_k^{Z_{ki}} (1 - p_k)^{i_k - Z_{ki}} = \prod_{k=1}^K \left(1 - \prod_{j=1}^k \phi_j \right)^{Z_{ki}} \left(\prod_{j=1}^k \phi_j \right)^{i_k - Z_{ki}} \quad (12.4)$$

where $i_k = i_k(i)$ is the number of patients treated at dose level k among the first i groups of patients, and Z_{ki} is the number of toxic outcomes in those patients. Thus, we have $\sum_{k=1}^K i_k = im$. Then the $(i + 1)$ st group of patients will be treated according to (12.1) with

$$\hat{\pi}_i(d_k) = E_i(p_k) = 1 - E_i \left(\prod_{j=1}^k \phi_j \right)$$

where $E_i(\cdot)$ denotes expectation taken with respect to the posterior distribution given the first i groups of observations.

By the binomial expansion of the likelihood (12.4), the posterior distribution of ϕ_j s can be expressed as a mixture of beta distributions. Thus, an exact computation formula can be derived [40]. However, exact computation will become numerically unstable as the sample size grows: the computation involves the addition of many positive and negative terms, which eventually amount to a value close to zero, and causes the cancelation of significant digits, a well-known numerical problem. Thus, one will need to resort to simulation integration methods such as Markov Chain Monte Carlo (MCMC) for general posterior computations [42].

Another practical issue is the elicitation of the hyperparameters (a_{1j}, a_{2j}) in the product-of-beta prior of ϕ_j . Gasparini and Eisele [40] suggest:

1. Choose the initial guesses p_{0j} s of toxicity probabilities as in the CRM.

2. Treat p_j as if it were a beta variable with hyperparameters (A_{1j}, A_{2j}) . Then determine A_{1j} and A_{2j} by matching the median of $\text{beta}(A_{1j}, A_{2j})$ to p_{0j} , such that its variance is maximized subject to unimodality. In effect, if $p_{0j} \leq 0.5$, then $A_{1j} = \log(0.5)/\log(p_{0j})$ and $A_{2j} = 1$; else, if $p_{0j} > 0.5$, then $A_{1j} = 1$ and $A_{2j} = \log(0.5)/\log(1 - p_{0j})$.
3. Determine the parameters a_{1j} and a_{2j} of the beta prior for ϕ_j by matching the first two moments via (12.2):

$$E_0(\phi_j) = \frac{a_{1j}}{a_{1j} + a_{2j}} \approx \frac{E_0(1 - p_j)}{E_0(1 - p_{j-1})}$$

$$E_0(\phi_j^2) = \frac{a_{1j}a_{2j}}{(a_{1j} + a_{2j})^2(a_{1j} + a_{2j} + 1)} + \left[\frac{a_{1j}}{a_{1j} + a_{2j}} \right]^2 \approx \frac{E_0(1 - p_j)^2}{E_0(1 - p_{j-1})^2}$$

for $j = 1, \dots, K$, where $E_0(1 - p_j) = A_{2j}/(A_{1j} + A_{2j})$ and

$$E_0(1 - p_j)^2 = \frac{A_{1j}A_{2j}}{(A_{1j} + A_{2j})^2(A_{1j} + A_{2j} + 1)} + \left[\frac{A_{2j}}{A_{1j} + A_{2j}} \right]^2$$

with $E_0(1 - p_0) = E_0(1 - p_0)^2 = 1$.

According to this algorithm, the prior median of p_j will be equal to the initial guess p_{0j} , and mean $E_0(p_j) = A_{1j}/(A_{1j} + A_{2j})$. Table 12.1 displays the specification of the product-of-beta prior in the context of the NeuSTART, with initial guesses $p_{01} = 0.02, p_{02} = 0.06, p_{03} = 0.10, p_{04} = 0.18$, and $p_{05} = 0.30$. Figure 12.1 shows that the prior distribution of p_j is skewed to the right for all j , thus implying that its mean will likely be greater than its median. This explains why $E_0(p_j) > p_{0j}$ in Table 12.1—and such right-skewedness is generally true if $p_{0j} < 0.5$. In this regard, the product-of-beta prior utilizes the initial guesses in a conservative manner.

Table 12.1 A product-of-beta prior for the NeuSTART

j	p_{0j}	(A_{1j}, A_{2j})	(a_{1j}, a_{2j})	$E_0(p_j)$	Corr. between p_j and			
					p_2	p_3	p_4	p_5
1	.02	(0.18, 1)	(1.00, 0.18)	.15	.86	.79	.70	.60
2	.06	(0.25, 1)	(1.18, 0.07)	.20		.92	.81	.70
3	.10	(0.30, 1)	(1.25, 0.06)	.23			.88	.76
4	.18	(0.40, 1)	(1.30, 0.10)	.29				.87
5	.30	(0.58, 1)	(1.40, 0.17)	.37				

While the prior of p_j is not a beta distribution in general, Figure 12.1 shows that the former is well approximated by the latter. This suggests that the posterior computations can be approximated and simplified by assuming $p_j \sim \text{beta}(A_{1j}, A_{2j})$ a priori, which is conjugate for the binomial likelihood of p_{0j} s. Precisely, the posterior mean can be approximated by

$$E_i(p_j) \approx \frac{A_{1j} + Z_{ji}}{A_{1j} + A_{2j} + i_j} \tag{12.5}$$

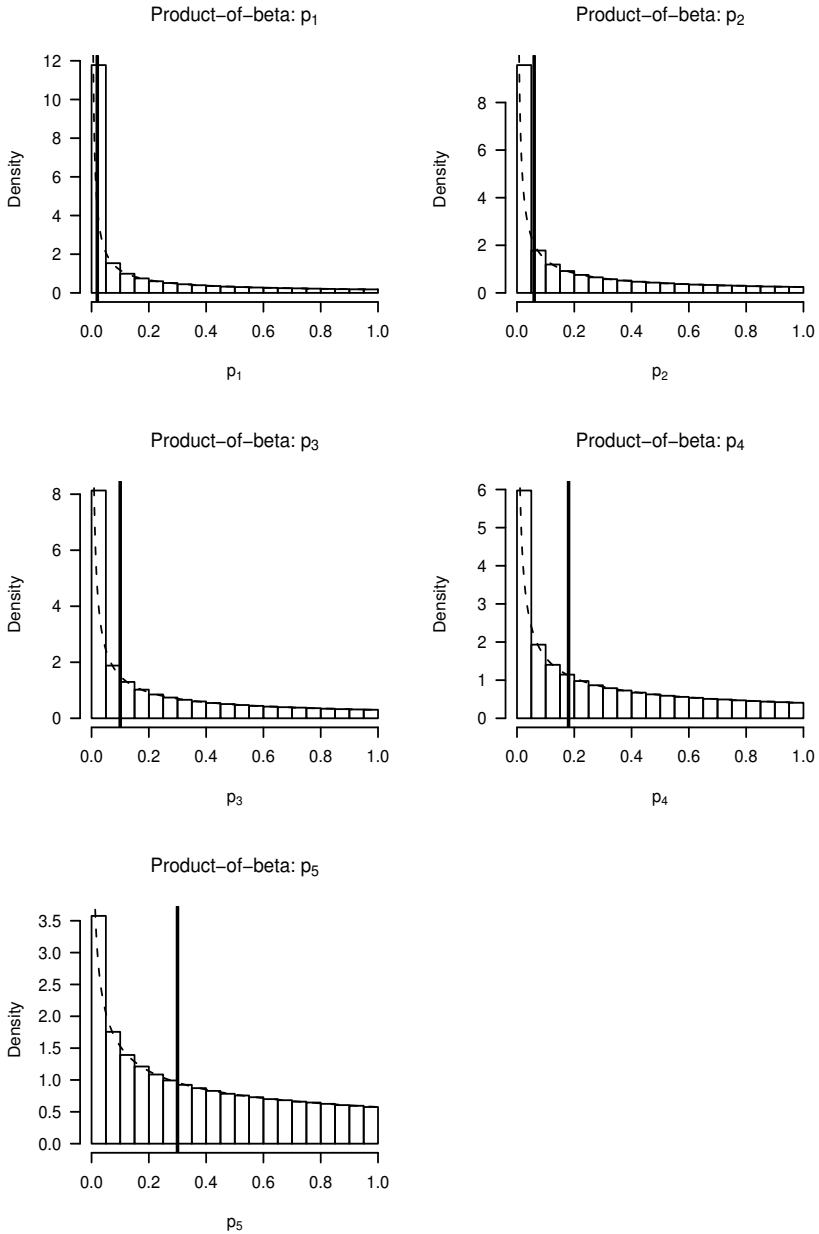


Figure 12.1 Marginal distributions of p_j s according to the product-of-beta prior in Table 12.1. In each figure, the dotted line indicates the density of $\text{beta}(A_{1j}, A_{2j})$, and the dark vertical line the initial guess p_{0j} of toxicity probability.

Example 12.1. Consider using the product-of-beta prior in Table 12.1 for the NeuSTART, in which the MTD is defined as a dose associated with a toxicity rate $\theta = 0.10$. Suppose that the trial starts at dose level 1 and enrolls $m = 3$ patients at a time, and that the first three patients have no indication of toxicity. Then the next group enters the trial at dose 2, and 1 out of 3 has a toxic outcome. That is, $i_1 = i_2 = 3$, $Z_{1i} = 0, Z_{2i} = 1$ with $i = 2$. Based on these observations, the exact posterior means of p_1 and p_2 are $E_2(p_1) = 0.11$ and $E_2(p_2) = 0.22$. This will bring the trial back to dose level 1. Suppose we use the approximation (12.5), we will get

$$E_2(p_1) \approx 0.04 \text{ and } E_2(p_2) \approx 0.29.$$

which also prescribe dose level 1 for the next group. Suppose, then, the next group of three patients is treated at dose 1 with one toxic outcome, so that $i_1 = 6$ and $Z_{1i} = 1$ with $i = 3$. Then the exact posterior means of p_1 and p_2 are 0.20 and 0.24, with respective approximation 0.16 and 0.29 according to the beta posterior (12.5).

Example 12.1 illustrates that the beta approximation (12.5) is reasonably accurate with $i = 3$ and will become increasingly accurate as i increases. This approximation thus provides a quick and easy posterior computation instead of using MCMC. See Exercise 12.2 for the exact computation of $E_i(p_1)$ and $E_i(p_2)$ in this example.

12.2.3 Dirichlet Prior Distribution

An alternative approach to enforce the monotonicity of p_k s is by modeling the non-negative increments of toxicity probabilities. Specifically, let

$$p_k = \sum_{j=1}^k \varphi_j$$

for $k = 1, \dots, K$, where $0 < \varphi_j < 1$ such that $\sum_{j=1}^K \varphi_j < 1$. A convenient choice of prior on φ_j s is the Dirichlet distribution, that is, having density proportional to

$$\left(\prod_{j=1}^K \varphi_j^{c_j-1} \right) \left(1 - \sum_{j=1}^K \varphi_j \right)^{c_{K+1}-1} \quad \text{for } c_j > 0 \text{ for all } j$$

The likelihood based on i groups of patients can be written as

$$\prod_{k=1}^K \left(\sum_{j=1}^k \varphi_j \right)^{Z_{ki}} \left(1 - \sum_{j=1}^k \varphi_j \right)^{i_k - Z_{ki}}.$$

Using the binomial expansion, we can show that the posterior distribution of φ_j s is a mixture of beta distributions, and derive an exact expression for the posterior mean of φ_j . However, it has similar numerical problem as in the product-of-beta prior, and MCMC should be used to obtain posterior quantities as sample size becomes moderate.

Under the Dirichlet prior formulation, the prior mean of p_k is $\sum_{j=1}^k c_j / C_+$, where $C_+ = \sum_{j=1}^{K+1} c_j$ can be viewed as the “prior sample size.” We may elicit c_j by matching the prior mean and the initial guess and solving

$$c_j = (p_{0j} - p_{0,j-1})C_+$$

for $j = 1, \dots, K$. To determine c_j s uniquely, however, one will need to specify the prior sample size. To maximize vagueness, one may iterate C_+ over the positive integers $\{1, 2, 3, \dots\}$ until the prior density of p_k is unimodal for all k . Table 12.2 gives the details of the Dirichlet prior using the initial guesses for the five doses in the NeuSTART, with $C_+ = \sum_{j=1}^6 c_j = 2$. Figure 12.2 shows that the prior marginals are unimodal, and are more informative and have shorter tails than the product-of-beta priors.

Table 12.2 A Dirichlet prior with $C_+ = 2$ for the NeuSTART

j	p_{0j}	c_j	Corr. between p_j and			
			p_2	p_3	p_4	p_5
1	.02	0.04	.56	.42	.30	.21
2	.06	0.08		.75	.53	.38
3	.10	0.08			.71	.51
4	.18	0.16				.71
5	.30	0.24				

12.2.4 Isotonic Design

Leung and Wang [63] take a frequentist nonparametric estimation approach, in which the toxicity probability of dose level k is estimated using isotonic regression [5]:

$$\tilde{p}_k = \min_{k \leq s \leq K} \max_{1 \leq r \leq k} \frac{\sum_{j=r}^s Z_{ji}}{\sum_{j=r}^s i_j}. \tag{12.6}$$

The expression (12.6) guarantees that $\tilde{p}_r \leq \tilde{p}_k \leq \tilde{p}_s$ for $r \leq k \leq s$ and enforces the monotonicity assumption. When the sample proportions of toxicity at the doses are monotone increasing, that is, $Z_{1i}/i_1 \leq Z_{2i}/i_2 \leq \dots \leq Z_{Ki}/i_K$, the isotonic estimates \tilde{p}_{ks} are equal to the sample proportions.

Similar to the other curve-free methods, the isotonic design treats the next group of patients according to (12.1) with $\hat{\pi}_i(d_k) = \tilde{p}_k$. Unlike the Bayesian nonparametric method, which enforces strict monotonicity through the prior structure, extra care is needed to handle the tied values among $|\tilde{p}_k - \theta|$ s because the nonparametric isotonic regression is highly unsmooth. Suppose the current group is treated at $x_n = d_j$. Then

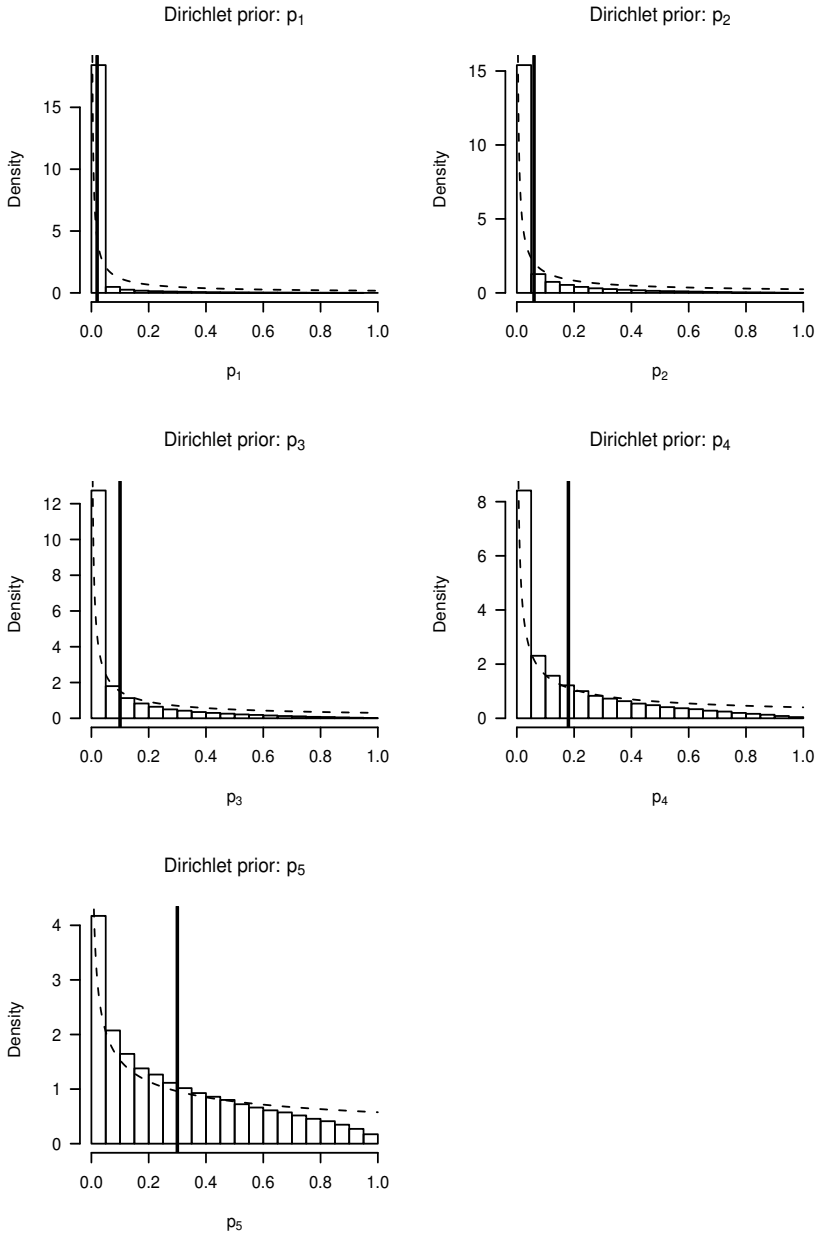


Figure 12.2 Marginal distributions of p_j s according to the Dirichlet prior in Table 12.2. In each figure, the dotted line indicates the density of $\text{beta}(A_{1j}, A_{2j})$, and the dark vertical line the initial guess p_{0j} of toxicity probability.

the dose for the next is determined as

$$x_{n+1} = \begin{cases} d_{j+1} & \text{if } \tilde{p}_j \leq \min(\theta, 2\theta - \tilde{p}_{j+1}) \\ d_{j-1} & \text{if } \tilde{p}_j \geq \theta \text{ and } \tilde{p}_j > 2\theta - \tilde{p}_{j-1} \\ d_j & \text{otherwise.} \end{cases} \quad (12.7)$$

For example, if $\theta = 0.10$, $\tilde{p}_j = 0.00$, and $\tilde{p}_{j+1} = 0.20$, then the next patient will receive dose d_{j+1} .

12.3 Rigidity

12.3.1 Illustrations of the Problem

The conventional intuition for using nonparametric estimation is to avoid bias and hence enhance the prospect of consistency. However, this conventional wisdom does not hold in dose finding trials where doses assignments are made sequentially. On the contrary, the curve-free methods may incur a rigidity problem that confines dose assignments to suboptimal levels *as a result of* the use of these “flexible” curve-fitting approaches:

Definition 12.1 (rigidity). A dose finding design \mathcal{D} is said to be rigid if for every $0 < p_L < \theta < p_U < 1$ and all $i \geq 1$,

$$P_{\mathcal{D}}\{x_i \in \mathcal{I}_{\pi}(p_L, p_U)\} < 1,$$

where $\mathcal{I}_{\pi}(p_L, p_U) = \{x : p_L \leq \pi(x) \leq p_U\}$.

It can be verified that δ -sensitivity excludes rigidity; see Section 5.3.1 for the definition of δ -sensitivity. Therefore, the one-parameter CRM, being δ -sensitive, is nonrigid.

In contrast, curve-free methods are rigid. To illustrate the problem, consider an outcome sequence specified by i_1, Z_{1i}, i_2 for given $Z_{2i} = 1$ and $i_k \equiv 0$ for $k \geq 3$. Using the product-of-beta prior, the exact posterior means of p_1 and p_2 are equal to

$$E_i(p_1) = 1 - Q_1 \left(1 + \frac{a_{22}}{Q_0} \right) \quad (12.8)$$

$$E_i(p_2) = 1 - Q_1 Q_2 \left(1 + \frac{a_{21} + a_{22} + Z_{1i}}{Q_0} \right) \quad (12.9)$$

where $Q_0 = (a_{12} + i_2 - 1)(a_{21} + Z_{1i}) + a_{22}(a_{11} + a_{21} + i_1 + i_2 - 1)$,

$$Q_1 = \frac{a_{11} + i_1 + i_2 - Z_{1i} - 1}{a_{11} + a_{21} + i_1 + i_2} \text{ and } Q_2 = \frac{a_{12} + i_2 - 1}{a_{12} + a_{22} + i_2}.$$

See Exercise 12.2. Applying the above expressions with $Z_{1i} = 0$ and $i_2 = 3$, the top panel of Figure 12.3 reveals that

$$|E_i(p_1) - 0.1| < |E_i(p_2) - 0.1|$$

for all i_1 . In other words, the curve-free method will select dose level 1 even if there is no toxic outcome (i.e., $Z_{1i} = 0$) in a long string of patients (i.e., as i_1 grows). And this happens primarily because there is one toxic outcome $Z_{2i} = 1$ among $i_2 = 3$ patients at dose level 2.

A similar problem occurs for the isotonic design: Consider the outcome sequence in Example 12.1. The isotonic estimates for p_1 and p_2 are $\tilde{p}_1 = 0$ and $\tilde{p}_2 = 0.33$ after the first six patients, and the next group of patients will be treated at level 1. From this point on, the estimate \tilde{p}_2 will be at least 0.33, that is, $|\tilde{p}_2 - 0.1| \geq 0.23$, because the data collected at dose level 1 will not affect estimation of p_2 (unless dose 1 is very toxic). As a result, if $\tilde{p}_1 \leq 0.1$, then $|\tilde{p}_1 - 0.1| < 0.23 \leq |\tilde{p}_2 - 0.1|$ and the trial will stay at dose level 1 indefinitely.

These two examples demonstrate that the rigidity problem is caused by the use of nonparametric estimation (which does not allow much borrowing strength from across doses) in conjunction with a sequential sampling plan. As a result, the design allows the possibility of an “extreme” overestimate of p_2 with a small sample at dose 2; in the above examples, this extremity is reached if there is at least one toxic outcome in a group of three patients. Now suppose that dose 2 is in truth safe with $p_2 = 0.10$. Then the probability that the trial is confined to the suboptimal dose 1 is least 0.27, which is nonnegligible and cannot be improved by increasing the sample size. If the true MTD is among the higher doses, the rigidity probability from the lower doses will accumulate and impede the method’s ability to escalate near the true MTD.

12.3.2 Remedy 1: Increase m

One way to alleviate the rigidity problem is by reducing the probability of extreme outcomes by using a larger group size m . For the isotonic design, with $\theta = 0.1$, rigidity will occur when there are two or more toxic outcomes in a group of $m = 5$ patients. Hence, the rigidity probability at a dose with $p_j = 0.1$ is $\Pr(Z_{ji} \geq 2 \mid i_j = 5) = 0.08$. This remedial measure reduces the rigidity probability substantially although it does not do away with the problem altogether.

Similarly, for the curve-free CRM, if we enroll patients in groups of size $m = 5$, with $Z_{1i} = 1$ and $i_2 = 5$, the trial will reescalate to dose level 2 with 10 consecutive nontoxic outcomes, that is, $i_1 = 10, Z_{1i} = 0$; see the middle panel of Figure 12.3.

The use of a large m , however, may reduce the number of interim decisions during the trial, and undermine the adaptiveness of a sequential design.

12.3.3 Remedy 2: Increase Prior Correlations

When using a Bayesian estimation approach, one may adopt prior distributions with high correlations among p_j s so that the estimation of p_j is influenced also by the observations from doses other than d_j .

For the product-of-beta prior, we may introduce strong prior correlations among p_j s by reducing the prior variance of ϕ_j for $j = 2, 3, \dots, K$. In other words, we may use an informative prior distribution. For example, if we replace the hyperparameters

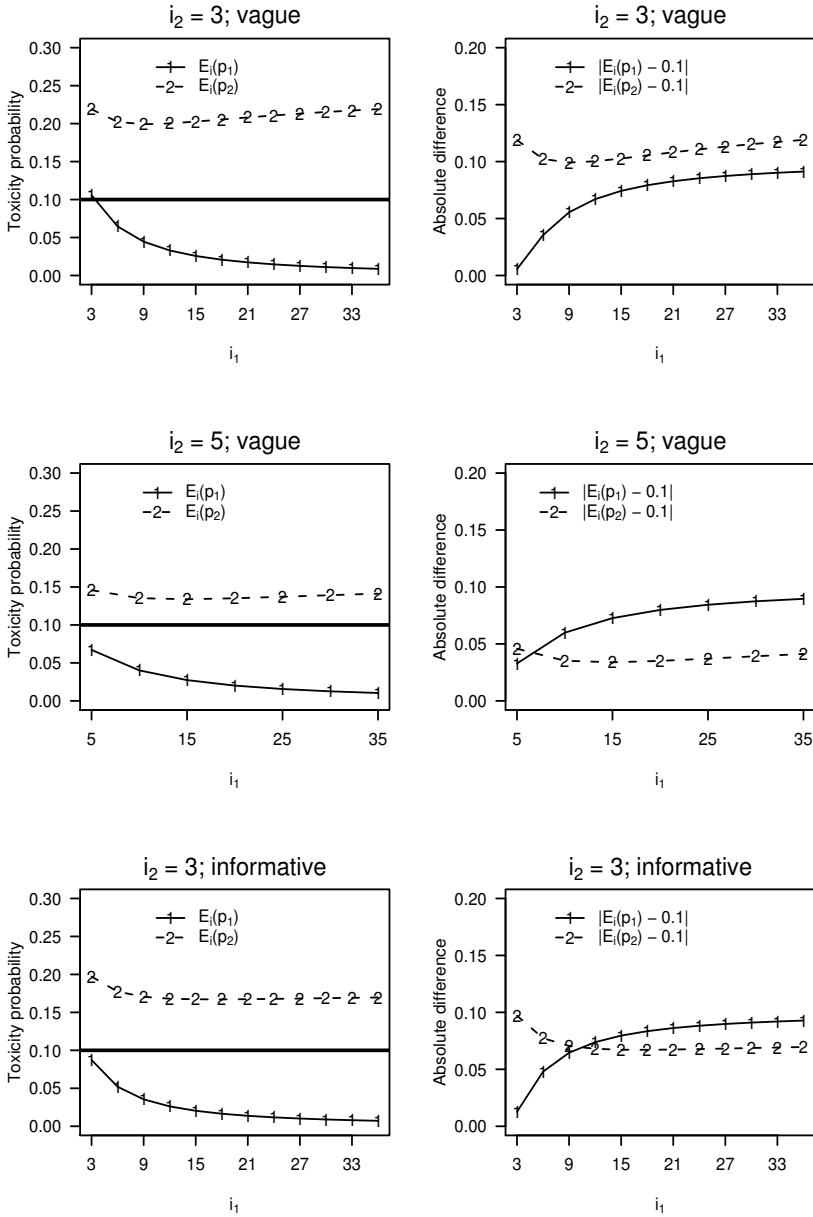


Figure 12.3 Left panel: The posterior means of toxicity probability associated with dose levels 1 and 2 against the number i_1 of patients treated at dose level 1 for $Z_{1i} = 0, Z_{2i} = 1$, and $i_2 = 3$ or 5. Right panel: the distance between the posterior means and the target $\theta = 0.1$. Top and middle panels: The prior specified in Table 12.1 is used in the calculation of the posterior means. Bottom panel: An informative prior in Section 12.3.3 is used.

$a_{12} = 1.00$ and $a_{22} = 0.18$ in Table 12.1 with $a_{12} = 3.53$ and $a_{22} = 0.21$, the prior correlation between p_1 and p_2 will increase to 0.93 from 0.86. The bottom panel of Figure 12.3 shows that, under this informative prior, the posterior mean $E_i(p_2)$ is pulled down to a greater extent by nontoxic observations at dose 1; as a result, even if we enroll in group size of $m = 3$, with $Z_{1i} = 1$ and $i_2 = 3$, the trial will reescalate to dose level 2 after three toxicity-free groups, that is, $i_1 = 9, Z_{1i} = 0$.

Simulation suggests that using an informative product-of-beta prior in the curve-free method reduces the rigidity probability and shifts the distribution of the selected dose away from the lower ends; and that these are achieved without increasing the frequency of selecting a toxic dose [15]. It leads to the investigation of the limiting case where $a_{1j}, a_{2j} \rightarrow \infty$ such that $\phi_j = a_{1j}(a_{1j} + a_{2j})^{-1} \equiv \phi_{0j}$ has a degenerated distribution for $j = 2, \dots, K$. In this case, the parameterization (12.2) is identical to a one-parameter model:

$$p_k \equiv F(d_k, \beta) = 1 - \beta/d_k \quad (12.10)$$

where the dose labels are $d_1 = 1$ and $d_k = (\phi_{02} \cdots \phi_{0k})^{-1}$ for $k = 2, \dots, K$, and $\beta = \phi_1$ is the only free parameter in the model. In addition, if the constants ϕ_{0j} s are chosen such that

$$\beta \prod_{j=2}^k \phi_{0j} \geq 1 - \theta \text{ for all } k, \quad (12.11)$$

then the limiting model (12.10) will satisfy the conditions for consistency stated in Section 4.3. In light of this, when using the product-of-beta prior, one may choose pointed priors for ϕ_j centering around ϕ_{0j} for $j \geq 2$, where ϕ_{0j} s satisfy (12.11).

The consideration of prior correlations also informs us as to the model choice. For example, since the prior correlations under the Dirichlet prior (Table 12.2) are smaller than those under the product-of-beta prior distribution (Table 12.1), the latter appears to be the preferred parameterization [40].

12.4 Two-parameter CRM

12.4.1 The Basic Approach

As the dose finding objective is formulated as a percentile estimation problem, it shares similar features with the well-studied problem in bioassay, for which there is an extensive set of analytical tools. Detailed expositions and discussions on this topic can be found in monographs such as Finney [37], McCullagh and Nelder [66], and Morgan [68]. In the literature, it is conventional to use two-parameter functions to capture the shape of a dose–toxicity curve. Some common examples are the logistic and the probit models. In particular, the logistic model postulates

$$F(x, \alpha, \beta) = \frac{\exp(\alpha + \beta x)}{1 + \exp(\alpha + \beta x)} \quad (12.12)$$

where α and β are free parameters to be estimated based on data. Specifically, the mle $(\tilde{\alpha}_i, \tilde{\beta}_i)$ is obtained by maximizing

$$L_i(\alpha, \beta) = \prod_{j=1}^i \left\{ \frac{\exp(\alpha + \beta x_j)}{1 + \exp(\alpha + \beta x_j)} \right\}^{Y_j} \left\{ 1 - \frac{\exp(\alpha + \beta x_j)}{1 + \exp(\alpha + \beta x_j)} \right\}^{1-Y_j} \quad (12.13)$$

where x_j and Y_j , respectively, denote the dose and binary outcome of patient j . It is therefore natural to consider using the two-parameter function (12.12) in conjunction with the CRM. Also, since a dose finding trial typically confines the test doses to a discrete set of levels, $\{d_1, \dots, d_K\}$, an estimate of the MTD based on the first i observations is given by

$$x_{i+1} = \arg \min_{d_k} \left| F(d_k, \tilde{\alpha}_i, \tilde{\beta}_i) - \theta \right|. \quad (12.14)$$

Other estimation methods are certainly possible; Chevret [25], for example, uses the posterior means of α and β instead of the mle. However, the choice of estimation is a relatively minor issue. Rather, the main difficulty with the two-parameter CRM arises from the fact that the design points x_i s are determined sequentially and thus are random. This is where the dose finding problem diverges from the classical bioassay problem, and rigidity may occur as a result.

12.4.2 A Rigid Two-Parameter CRM: Illustration

Suppose the outcome sequence in Example 12.1 is observed, that is, zero toxic outcome out of i_1 subjects at dose level 1, and 1 out of 3 at dose level 2. Consider the logistic model (12.12) with $x \in \{1, \dots, K\}$, that is, $d_k \equiv k$. Then the likelihood (12.13) can be simplified as

$$L_i(\alpha, \beta) = \frac{\exp(\alpha + 2\beta)}{\{1 + \exp(\alpha + \beta)\}^{i_1} \{1 + \exp(\alpha + 2\beta)\}^3}. \quad (12.15)$$

The mle $\tilde{\alpha}_i$ and $\tilde{\beta}_i$ can be shown to solve

$$\frac{\partial}{\partial \alpha} \log L_i(\alpha, \beta) = 1 - \frac{i_1 \exp(\alpha + \beta)}{1 + \exp(\alpha + \beta)} - \frac{3 \exp(\alpha + 2\beta)}{1 + \exp(\alpha + 2\beta)} = 0$$

and

$$\frac{\partial}{\partial \beta} \log L_i(\alpha, \beta) = 2 - \frac{i_1 \exp(\alpha + \beta)}{1 + \exp(\alpha + \beta)} - \frac{6 \exp(\alpha + 2\beta)}{1 + \exp(\alpha + 2\beta)} = 0$$

from which we can derive

$$F(d_2, \tilde{\alpha}_i, \tilde{\beta}_i) = \frac{\exp(\tilde{\alpha}_i + 2\tilde{\beta}_i)}{1 + \exp(\tilde{\alpha}_i + 2\tilde{\beta}_i)} = \frac{1}{3} \quad (12.16)$$

and

$$\frac{\partial \log L_i(\alpha, \beta)}{\partial \alpha} < 0 \quad (12.17)$$

at the mle. For technical purposes, suppose the parameters (α, β) are confined to a finite rectangle. If α and β are allowed to live in an open parameter space in this example, the mle does not exist because of (12.17). Then, we can further establish that

$$F(d_1, \tilde{\alpha}_i, \tilde{\beta}_i) = \frac{\exp(\tilde{\alpha}_i + \tilde{\beta}_i)}{1 + \exp(\tilde{\alpha}_i + \tilde{\beta}_i)} \approx 0 \quad (12.18)$$

Now, with a target $\theta = 0.10$, substituting (12.16) and (12.18) into (12.14) will lead to dose level 1 as the next recommended dose, that is, $x_{i+1} \equiv d_1$.

To complete the rigidity illustration, we claim that the two-parameter CRM will always recommend the lowest dose from this point on, regardless of what happens at dose level 1. Exercise 12.3 outlines the proof of this claim.

12.4.3 Three-Stage Design

Like the curve-free method, when the dose–toxicity curve of a two-parameter model is estimated based on very few data points, extreme observations early in the trial may confine (with nonnegligible probability) the trial to suboptimal dose. Silvapulle [96] shows that the mles $\tilde{\alpha}_i$ and $\tilde{\beta}_i$ exist and are unique if and only if any of the following conditions holds:

$$\begin{aligned} (i) \quad & (x_{\min}^+, x_{\max}^+) \cap (x_{\min}^-, x_{\max}^-) \text{ is nonempty} \\ (ii) \quad & x_{\min}^+ < x_{\min}^- = x_{\max}^- < x_{\max}^+ \\ (iii) \quad & x_{\min}^- < x_{\min}^+ = x_{\max}^+ < x_{\max}^- \end{aligned} \quad (12.19)$$

where $x_{\max}^+ = \max\{x_j : Y_j = 1\}$, $x_{\max}^- = \max\{x_j : Y_j = 0\}$, $x_{\min}^+ = \min\{x_j : Y_j = 1\}$ and $x_{\min}^- = \min\{x_j : Y_j = 0\}$. In contrast, if a one-parameter CRM is used, the mle will exist as soon as there is heterogeneity in the toxicity outcomes among patients, which is a weaker condition than (12.19). In other words, it will take more observations before the model-based MTD estimate (12.14) can come into effect than when its one-parameter counterpart (3.11) is used. Thus, it is natural to consider a three-stage CRM design if one would like to eventually adopt a two-parameter model:

1. Treat the initial patients according to a predetermined dose sequence $x_{1,0}, \dots, x_{N,0}$.
2. Switch to one-parameter CRM once a toxicity is observed among patients.
3. Switch to two-parameter CRM once Condition (12.19) is met.

The stage of using one-parameter CRM may also delay the use of a two-parameter model until sufficient data are available. Table 12.3a displays the outcomes of a simulated trial by a three-stage CRM for the NeuSTART with $\theta = 0.10$ in $N = 33$ subjects. The trial starts with a group-of-three initial design, switches to the one-parameter model used in the NeuSTART (Chapter 10, Section 10.5) upon the first observed toxic outcome, and uses the two-parameter estimate (12.14) once Condition (12.19) comes to hold.

The trial is simulated under scenario 3 in Chapter 10, Table 10.5, where dose level 3 is the true MTD. To summarize, the trial takes 22 patients before the two-parameter model comes into effect, and treats the remaining 11 patients at dose level

Table 12.3 Simulated trials of two-parameter versus one-parameter CRM, with $\theta = 0.10$ and $N = 33$ under the scenario $p_1 = 0.01$, $p_2 = 0.04$, $p_3 = 0.10$, $p_4 = 0.0.25$, and $p_5 = 0.30$

Stage	i	u_i	(a) Three-stage two-parameter CRM:							(b) Two-stage one-parameter CRM:						
			x_i	Y_i	Model-based estimates of					x_i	Y_i	Model-based estimates of				
					p_1	p_2	p_3	p_4	p_5			p_1	p_2	p_3	p_4	p_5
Initial	1	.613	1	0	—	—	—	—	—	Identical to the three-stage CRM						
	2	.636	1	0	—	—	—	—	—							
	3	.780	1	0	—	—	—	—	—							
	4	.058	2	0	—	—	—	—	—							
	5	.189	2	0	—	—	—	—	—							
	6	.018	2	1	.11	.20	.27	.38	.50							
CRM (1p)	7	.992	1	0	.09	.18	.25	.35	.48	Identical to the three-stage CRM						
	8	.891	1	0	.08	.16	.23	.33	.46							
	9	.821	1	0	.07	.15	.21	.32	.44							
	10	.500	1	0	.06	.14	.20	.30	.43							
	11	.165	1	0	.06	.13	.19	.29	.42							
	12	.368	2	0	.05	.12	.17	.27	.40							
	13	.142	2	0	.04	.11	.16	.26	.38							
	14	.645	2	0	.04	.10	.15	.24	.37							
	15	.979	2	0	.04	.09	.14	.23	.36							
	16	.412	2	0	.03	.08	.13	.22	.35							
	17	.905	2	0	.03	.08	.12	.21	.33							
	18	.249	2	0	.03	.07	.12	.20	.33							
	19	.670	3	0	.02	.07	.11	.19	.31							
	20	.499	3	0	.02	.06	.10	.18	.30							
	21	.718	3	0	.02	.06	.09	.17	.29							
CRM (2p)	22	.009	3	1	.02	.07	.28	.66	.90	3	1	.04	.10	.15	.24	.37
	23	.246	2	0	.01	.07	.28	.67	.91	2	0	.04	.09	.14	.23	.36
	24	.727	2	0	.01	.07	.28	.67	.92	2	0	.03	.07	.14	.23	.35
	25	.414	2	0	.01	.06	.27	.68	.92	2	0	.03	.06	.13	.22	.34
	26	.242	2	0	.01	.06	.27	.69	.93	2	0	.03	.06	.13	.21	.34
	27	.144	2	0	.01	.06	.27	.69	.93	2	0	.03	.06	.12	.21	.33
	28	.852	2	0	.01	.05	.27	.70	.94	3	0	.03	.05	.12	.20	.32
	29	.303	2	0	.01	.05	.27	.71	.94	3	0	.02	.05	.11	.20	.32
	30	.443	2	0	.01	.05	.26	.72	.95	3	0	.02	.05	.11	.19	.31
	31	.486	2	0	.01	.05	.26	.72	.95	3	0	.02	.05	.10	.18	.30
	32	.117	2	0	.01	.05	.26	.73	.95	3	0	.02	.05	.10	.18	.30
	33	.742	2	0	.01	.04	.26	.73	.96	3	0	.02	.04	.10	.17	.29

2, which is the selected MTD in this simulated trial. In fact, if we continue the trial with additional patients, the trial will remain at dose 2 or below regardless of what happens. This example thus illustrates that while rigidity may be delayed by a three-stage design strategy, the problem remains: as the two-parameter function is flexible enough to fit well at two dose levels, the estimate for dose 3, $F(d_3, \tilde{\alpha}_i, \tilde{\beta}_i)$, does not change much after patient 22 and stays close to the sample proportion $\hat{p}_3 = 1/4$.

A simulated trial by the NeuSTART design is also given in Table 12.3b for comparison. To match the two designs, the same set of uniform variates (u_i s) for the toxicity tolerance is used to generate the outcomes. The one-parameter CRM would reescalate to dose level 3 after an additional five nontoxic outcomes at dose 2 (patients 23 through 27). This example sheds favorable light on the one-parameter CRM when compared to the two-parameter CRM in terms of the rigidity property. In addition, the two-parameter model does not come into play until the trial is over halfway through in the example. Since in general, it takes many observations before Condition (12.19) is satisfied, the typical sample size used in early phase dose finding trial may not allow effective use of a two-parameter model.

12.4.4 Continuous Dosage

Piantadosi et al. [81] use a two-parameter CRM in a dose finding study in which a continuum of test doses is available. The use of continuous dosage is an exception where using a two-parameter CRM model is reasonable, at least from a theoretical viewpoint. Since we are allowed to test two doses that are arbitrarily close to each other, the rigidity problem is resolved. Ying and Wu [116] prove that a “CRM-like” sequential design, called maximum likelihood recursion, with two-parameter model consistently estimates the target percentile.

However, there are still two practical issues to consider when a two-parameter model is used. First, a nontrivial number of observations are usually needed before Condition (12.19) is satisfied. In the simulated trial in Table 12.3, for instance, it takes two-thirds of the final sample size before the mles for a two-parameter model exist. To get around with the problem, Piantadosi et al. [81] suggest including pseudo-data in model fitting so as to ensure the mle exists. While this approach gives a pragmatic solution from a computational perspective, the fundamental problem remains: the informational content in the middle of a trial is so low—few data points and binary outcomes—that a two-parameter curve may overfit. Thus, the pseudo-data approach should be used with caution.

A second practical issue is that continuous dosage may not be feasible or even physically plausible in many clinical situations. If a discrete set of doses is used in a trial, we will fall back to the rigidity problem described in this chapter. Chapter 14 will review the stochastic approximation and connect it with the CRM. As it turns out, the use of discrete test doses is the primary feature that distinguishes the CRM from the large stochastic approximation literature.

12.5 Bibliographic Notes

Simulation studies in the CRM literature compare the one-parameter CRM favorably with the two-parameter CRM in terms of operating characteristics. Chevret [25] finds that the two-parameter CRM does not improve bias and mean squared error when compared to the one-parameter CRM. O'Quigley et al. [77] find that two-parameter CRM could be less efficient than one-parameter CRM; this is true even when the two-parameter model used is a correct specification.

In line with the use of underparameterized models in the CRM, Shen and O'Quigley [95] study the general use of one-parameter models in estimating the root of a regression function. O'Quigley [74] describes the computation difficulties associated with fitting a two-parameter model in the CRM.

Gasparini and Eisele [40] and Leung and Wang [63] introduce nonparametric estimation of the dose–toxicity curves in conjunction with the CRM. Cheung [15] subsequently points out the problem of rigidity with these nonparametric methods.

12.6 Exercise and Further Results

Exercise 12.1. Verify the numerical values obtained in Table 12.1 for the skeleton $p_{01} = 0.02, p_{02} = 0.06, p_{03} = 0.10, p_{04} = 0.18,$ and $p_{05} = 0.30$.

Exercise 12.2. Verify that the exact posterior means in Example 12.1 are given by (12.8) and (12.9).

Exercise 12.3. Consider the two-parameter CRM used in Section 12.4.2, where we show that $x_{i+1} = d_1$ after the first six observations given in Example 12.1.

a. Verify that, after the first six observations given, the mles $\tilde{\alpha}_i$ and $\tilde{\beta}_i$ solve

$$Z_{1i} + 1 - \frac{i_1 \exp(\alpha + \beta)}{1 + \exp(\alpha + \beta)} - \frac{3 \exp(\alpha + 2\beta)}{1 + \exp(\alpha + 2\beta)} = 0$$

and

$$Z_{1i} + 2 - \frac{i_1 \exp(\alpha + \beta)}{1 + \exp(\alpha + \beta)} - \frac{6 \exp(\alpha + 2\beta)}{1 + \exp(\alpha + 2\beta)} = 0$$

so that (12.16) still holds.

b. Show that $F(d_1, \tilde{\alpha}_i, \tilde{\beta}_i) = Z_{1i}/i_1$ when $Z_{1i} \geq 1$, and hence

$$|F(d_1, \tilde{\alpha}_i, \tilde{\beta}_i) - 0.1| < |F(d_2, \tilde{\alpha}_i, \tilde{\beta}_i) - 0.1|$$

unless $Z_{1i}/i_1 > 1/3$.

The rigidity of the two-parameter CRM can then be completed by considering the case $Z_{1i} = 0$ using similar arguments as in Section 12.4.2. This exercise also presents the potential problem that $F(d_1, \tilde{\alpha}_i, \tilde{\beta}_i) > F(d_2, \tilde{\alpha}_i, \tilde{\beta}_i)$ without putting a constraint on $\tilde{\beta}_i$.

When the CRM Fails

13.1 Introduction

Most dose finding methods, including the CRM, are suitable for clinical settings that satisfy at least one of the following assumptions. First, there is a *strictly* increasing dose–toxicity relationship around the target dose. Second, the toxicity endpoint is an adequate surrogate for efficacy, which often takes much longer to observe. While these assumptions are reasonable for traditional cancer chemotherapy, they are debatable in other disease areas or treatment options. When neither assumption holds, the surrogacy perspective of the MTD is not an appropriate operational objective, and as a result the CRM is not a suitable design for dose finding. This chapter introduces alternative paradigms to the surrogacy perspective of the MTD. Section 13.2 presents the notion of maximum safe dose, and Section 13.3 outlines some basic ideas of bivariate dose finding in combined phase I/II trials.

13.2 Trade-Off Perspective of MTD

13.2.1 Motivation

The ASCENT trial introduced in Example 2.3 involves six physical therapy (PT) regimens with escalating intensity via the timing of the PT sessions or the duration per session; see Chapter 2, Table 2.1. The target adverse event rate $\theta = 0.25$ is approximately the same as the background adverse event rate in the untreated population. As such, we expect the lower end of the dose–toxicity curve to be quite flat at 25% as shown in Figure 2.1c. Consequently, the surrogacy definition of the MTD

$$v = \arg \min_k |p_k - 0.25|$$

may not be uniquely defined, and dose finding designs that adopt this MTD definition can lead to arbitrary dose selection. Table 13.1 shows the operating characteristics of a two-stage CRM under three plausible toxicity scenarios for the ASCENT trial. On the one hand, the CRM performs in the way it is intended, by selecting a dose associated with 25% toxicity rate. On the other hand, it is clear in this context that the desired dose is the largest dose with a 25% toxicity rate, because the toxicity endpoint is not a surrogate for efficacy. (In contrast, if therapeutic benefits of a drug work through the toxic side effects, as in cytotoxic agents, it may be acceptable to

choose any dose with toxicity rate θ ; this is essentially the concept of indifference intervals used in Chapter 8.) In other words, one should take the trade-off perspective to define the MTD:

$$\gamma = \max\{k : p_k \leq \theta\} \quad (13.1)$$

where $\theta = 0.25$ for the ASCENT trial.

Table 13.1 Dose selection and the average toxicity number (ATN) of a two-stage CRM and a stepwise SPRT procedure for the ASCENT trial

Model	Probability of selecting dose							ATN	ASN ^a
	0	1	2	3	4	5	6		
Scenario L1:	.00	.25	.25	.25	.45	.45	.45		
CRM	.00	.47	.20	.28	.05	.00	.00	13	50
SPRT	.12	.10	.13	.60	.03	.02	.01	16	49
Scenario L2:	.00	.25	.25	.25	.25	.45	.45		
CRM	.00	.45	.17	.15	.20	.03	.00	13	50
SPRT	.11	.10	.09	.12	.54	.03	.01	15	49
Scenario L3:	.00	.25	.25	.25	.25	.25	.45		
CRM	.00	.45	.16	.13	.12	.13	.02	13	50
SPRT	.12	.09	.09	.08	.11	.49	.03	14	47

Note: The CRM uses the logistic model (10.4) specified with the following design parameters: $\theta = 0.25$, $K = 6$, $v_0 = 3$, $\delta = 0.05$, $N = 50$; the initial design is extracted from Table 10.3. Numbers associated with the MTD γ are in bold.

^aASN = average sample number.

13.2.2 Maximum Safe Dose and Multiple Testing

A related concept to MTD (13.1) is the maximum safe dose (MAXSD). A dose k is said to be safe if $p_k < \vartheta$ for some $\vartheta > \theta$. Then the MAXSD is defined as

$$\text{MAXSD} = \max\{k : p_k < \vartheta\}. \quad (13.2)$$

In the ASCENT trial, we set $\vartheta = 0.45$. It is easy to verify that the MAXSD is always at least as large as the MTD. However, the MAXSD (13.2) is equal to the MTD (13.1) under the following class of toxicity scenarios

$$\{\pi : p_k \notin (\theta, \vartheta) \text{ for all } k\}. \quad (13.3)$$

This class of scenarios is characterized by the exclusion of the so-called indifference zone [7]. Our goal is to estimate the MTD (13.1). However, in light of the equivalence between the MTD (13.1) and MAXSD (13.2) under (13.3), we can work on an altered objective: design a study to have reasonable estimation property for the MAXSD under (13.3). Therefore, we set the first design principle for the MTD objective (13.1) as follows:

(P1) Maintaining a reasonably high probability of selecting the MAXSD under all toxicity scenarios belonging to class (13.3).

To avoid additional notation, we shall use γ to denote MAXSD in this section unless there is ambiguity. Also, with the notion of an unsafe dose defined with respect to ϑ , we have another obvious design principle:

(P2) Keeping the probability of selecting an unsafe dose low under *all* toxicity scenarios.

These two design principles for MTD estimation can be formally addressed via a multiple testing framework. Specifically, consider testing a family of hypotheses:

$$H_{0k} : p_k \geq \vartheta \text{ versus } H_{1k} : p_k < \vartheta \text{ for } k = 1, \dots, K. \quad (13.4)$$

As in conventional hypothesis testing, a type I error is committed when a true null H_{0k} is falsely rejected, or equivalently, when an unsafe dose is declared safe by the data. Then the familywise (type I) error rate of a test procedure $\hat{\gamma}$, is defined as

$$\begin{aligned} \text{FWER}(\hat{\gamma}) &\equiv \Pr(\text{Any true } H_{0k} \text{ is rejected}) \\ &= \max_{0 \leq m \leq K} \sup_{\mu \in \Theta_m} P_\pi(\hat{\gamma} > m) \end{aligned} \quad (13.5)$$

$$= \Pr(\hat{\gamma} > \text{MAXSD}) \quad (13.6)$$

where P_π in (13.5) denotes probability computed under the dose-toxicity curve π , and $\Theta_m = \{\pi : p_m < \vartheta, p_k \geq \vartheta; k > m\}$ is the parameter subspace wherein $\gamma = m$. Because of the equality between $\text{FWER}(\hat{\gamma})$ and (13.6), the probability of selecting an unsafe dose can be controlled by keeping

$$\text{FWER}(\hat{\gamma}) \leq \alpha_0, \quad (13.7)$$

an objective to which a statistical test is naturally suited. Since $\{\Theta_m : m = 0, 1, \dots, K\}$ in (13.5) jointly partition the entire parameter space $[0, 1]^K$, a dose finding design (or a statistical test) that satisfies (13.7) provide strong control of the familywise error without any additional assumptions on the parameter space (cf., principle [P2]).

In addition, the selection probability of the true MAXSD (and hence the MTD) under (13.3) can be controlled via the multiple testing framework as follows:

$$\text{PCS}(\hat{\gamma}) = \min_{0 \leq m \leq K} \inf_{\pi \in \Theta_m^*} P_\pi(\hat{\gamma} = m) \geq 1 - \alpha_1 \quad (13.8)$$

where $\Theta_m^* = \{\pi : \max(p_1, \dots, p_m) \leq \theta, p_k \geq \vartheta; k > m\}$ is a subset of Θ_m so that $\cup_{m=0}^K \Theta_m^*$ form the scenario class (13.3). Thus, constraint (13.8) in effect controls the probability of selecting the MTD γ under all scenarios of no indifference zone (cf., design principle [P1]).

13.2.3 A Sequential Stepwise Procedure

Once the dose finding problem is formulated as a multiple testing problem, a large number of statistical procedures can, in theory, be used to estimate the MAXSD.

The most well-known procedure is the Bonferroni adjustment. Among many others, stepwise procedures can be easily adapted to the dose finding setting to provide a strong control of the FWER [100, 46, 28]. Briefly, a stepwise procedure tests one hypothesis from the family (13.4) at a time, and proceeds to the next hypothesis only when certain decisions are reached regarding the current hypothesis. Exploiting the *stepwise* nature of stepwise procedures, Cheung [18] proposes a two-stage stepwise procedure that can be implemented in a *sequential* manner:

Stage 1 An escalation stage where the family (13.4) is tested in the order of

$$H_{01} \rightarrow H_{02} \rightarrow \cdots \rightarrow H_{0K} \quad (13.9)$$

so that patients are enrolled to dose i only if all doses below i have been tested and declared safe, that is, H_{0j} is rejected for all $j < i$. Once a null hypothesis H_{0i} is accepted for some i , the stepwise procedure will declare dose i and the doses above to be unsafe, and it will be unnecessary to enroll patients at higher doses. Stage 1 ends upon the first acceptance of a null hypothesis. The sequence (13.9) of testing and this stopping rule is called a step-down (SD) test. The term “step-down” is used as a mathematical term here, while the sequence (13.9) starts testing a low dose *up* to a high dose. Generally, a SD test starts with testing the most restrictive hypothesis—in this case, H_{01} —and continues to the next restrictive one until the data fail to reject a null hypothesis.

Stage 2 A deescalation stage where the family (13.4) is tested in the order of

$$H_{0,S_2} \rightarrow H_{0,S_2-1} \rightarrow \cdots \rightarrow H_{01} \quad (13.10)$$

where S_2 is the dose immediately below the stopping dose in stage 1. That is, patients are enrolled to dose i only if all doses above i have been tested and declared unsafe. Deescalation in stage 2 stops once a hypothesis H_{0i} is rejected using the *cumulative* data accrued in both stages, thus avoid unnecessarily treating patients at low doses. The stopping dose in Stage 2 is used to estimate the MAXSD. The sequence (13.10) of testing and this stopping rule is called a step-up (SU) test; like the SD test, the term “step-up” has a mathematical connotation and describes a procedure that starts testing at the least restrictive hypothesis. This two-stage design is thus a step-down-step-up (SDSU) test [82]. The MAXSD estimate resulted from the SDSU procedure is denoted as $\hat{\gamma}_{SDSU}$.

To make the SDSU procedure more precise, let \mathcal{R}_{i1} and \mathcal{A}_{i1} , respectively, denote the rejection and acceptance regions of H_{0i} based on the *cumulative* data accrued to dose i up to stage l for $l = 1, 2$. The starting dose in stage 2 is then given as

$$S_2 = \min\{i : \mathcal{A}_{i1} \text{ is observed}\} - 1$$

and the final MAXSD estimate is given as

$$\hat{\gamma}_{SDSU} = \max\{i : \mathcal{R}_{i1} \cap \mathcal{R}_{i2} \text{ is observed}\}. \quad (13.11)$$

The operating characteristics of $\hat{\gamma}_{SDSU}$ can be easily computed for any given test

regions \mathcal{R}_{i1} and \mathcal{A}_{i1} . Let $\xi_i = \Pr(\mathcal{R}_{i1} \mid p_i)$ and $\rho_i = \Pr(\mathcal{R}_{i1} \cap \mathcal{R}_{i2} \mid p_i)$. Then

$$P_\pi(\hat{\gamma}_{SDSU} = m) = \left(\prod_{j=1}^{m-1} \xi_j \right) \rho_m \Omega_{m+1,K} \tag{13.12}$$

for $m = 1, \dots, K$, where

$$\Omega_{m,k} = 1 - \xi_m + \sum_{i=m}^{K-1} \left\{ (1 - \xi_{i+1}) \prod_{j=m}^i (\xi_j - \rho_j) \right\} + \prod_{j=m}^K (\xi_j - \rho_j)$$

with $\Omega_{K+1,K} \equiv 1$. Also,

$$P_\pi(\hat{\gamma}_{SDSU} = 0) = 1 - \sum_{m=1}^K P_\pi(\hat{\gamma}_{SDSU} = m). \tag{13.13}$$

In addition, the key properties $\text{FWER}(\hat{\gamma}_{SDSU})$ and $\text{PCS}(\hat{\gamma}_{SDSU})$ can be calculated by using the expressions (13.12) and (13.13), as follows:

$$\text{FWER}(\hat{\gamma}_{SDSU}) = 1 - P_{\pi_0^L}(\hat{\gamma}_{SDSU} = 0) \text{ and } \text{PCS}(\hat{\gamma}_{SDSU}) = \min_{0 \leq m \leq K} P_{\pi_m^L}(\hat{\gamma}_{SDSU} = m)$$

where π_m^L is the so-called *least favorable configuration* under which $p_i = \theta$ for $i \leq m$ and $= \vartheta$ for $i > m$. The simulations in Table 13.1, for example, are done under three least favorable configurations, $\pi_3^L, \pi_4^L, \pi_5^L$. In other words, we can choose test regions \mathcal{R}_{i1} and \mathcal{A}_{i1} that respect the error constraints (13.7) and (13.8) by calibrating the individual test properties, namely, ξ_i and ρ_i . While there are many possible choices of \mathcal{R}_{i1} and \mathcal{A}_{i1} , Cheung [18] details the use of the likelihood ratio test (LRT), test region that mimics the standard 3+3 algorithm, and the sequential probability ratio test (SPRT). Furthermore, while (13.12) and (13.13) apply to the SDSU procedure that starts at the lowest dose level, the distribution of $\hat{\gamma}_{SDSU}$ for a general starting dose is also available [18].

13.2.4 Case Study: The ASCENT Trial

The SDSU procedure with the SPRT, denoted as SDSU-SPRT, was used to design the ASCENT trial. Let $Z_i(j)$ denote the number of toxic outcomes in the first j patients at dose i . The SPRT prescribes test regions

$$\mathcal{R}_{i1} = \{\lambda_{i,\tau_{il}} \geq c_l\} \text{ and } \mathcal{A}_{i1} = \{\lambda_{i,\tau_{il}} \leq c^*\} \tag{13.14}$$

for $l = 1, 2$, where $\lambda_{i,n}$ is the likelihood ratio for dose i :

$$\lambda_{i,n} = \frac{\theta^{Z_i(n)} (1 - \theta)^{n - Z_i(n)}}{\vartheta^{Z_i(n)} (1 - \vartheta)^{n - Z_i(n)}}$$

and $\tau_{il} = \inf\{n > 0 : \lambda_{i,n} \geq c_l \text{ or } \leq c^*\}$ is a stopping time.

The test regions (13.14) and, hence, SDSU-SPRT are completely specified by

the boundary parameters, $c^* < c_1 < c_2$, and their statistical properties ξ_i and ρ_i have well-known approximations in terms of these boundary parameters [108]. Thus, the SDSU procedure can be calibrated to satisfy the error constraints (13.7) and (13.8) so that

- The probability of erroneous escalation is controlled, that is, $\Pr(\mathcal{R}_{i1} | \vartheta) \leq \varepsilon^*$.
- The average number of patients treated at an overdose is minimized.

See Algorithm 2 in Cheung [18] for details. For the ASCENT trial, we calibrated SDSU-SPRT with the following design parameters and constraints: $K = 6$, $\theta = 0.25$, $\vartheta = 0.45$, $\alpha_0 = 0.1$, $1 - \alpha_1 = 0.4$, and $\varepsilon^* = 0.5$. and obtained $c^* = 0.24$, $c_1 = 1.76$, and $c_2 = 13.74$. For convenience during the trial, the decision boundaries of the SPRT can be charted in terms of the number of toxic outcomes. Table 13.2 displays the escalation/deescalation decisions according to this SPRT for the first 20 patients at a dose. While the SPRT manages to escalate quickly stage 1 with two consecutive nontoxic outcomes, it also is cautious to make a final safety declaration needing at least nine subjects at a dose (if none exhibits any adverse event).

Table 13.2 Decisions for dose i for the first 20 subjects by the SDSU-SPRT with $c^* = 0.24$, $c_1 = 1.76$, and $c_2 = 13.74$ in the ASCENT trial

j	Stage 1: intervals of $Z_i(j)$			Stage 2: intervals of $Z_i(j)$		
	Escalate	Stay	Deescalate ^a	Stop ^b	Stay	Deescalate
2	0	[1,2]	—	—	[0,2]	—
3	0	[1,2]	3	—	[0,2]	3
4	0	[1,2]	[3,4]	—	[0,2]	[3,4]
5	[0,1]	[2,3]	[4,5]	—	[0,3]	[4,5]
6	[0,1]	[2,3]	[4,6]	—	[0,3]	[4,6]
7	[0,1]	[2,3]	[4,7]	—	[0,3]	[4,7]
8	[0,2]	[3,4]	[5,8]	—	[0,4]	[5,8]
9	[0,2]	[3,4]	[5,9]	0	[1,4]	[5,9]
10	[0,2]	[3,5]	[6,10]	0	[1,5]	[6,10]
11	[0,3]	[4,5]	[6,11]	0	[1,5]	[6,11]
12	[0,3]	[4,5]	[6,12]	[0,1]	[2,5]	[6,12]
13	[0,3]	[4,6]	[7,13]	[0,1]	[2,6]	[7,13]
14	[0,4]	[5,6]	[7,14]	[0,1]	[2,6]	[7,14]
15	[0,4]	[5,6]	[7,15]	[0,2]	[3,6]	[7,15]
16	[0,4]	[5,7]	[8,16]	[0,2]	[3,7]	[8,16]
17	[0,5]	[6,7]	[8,17]	[0,2]	[3,7]	[8,17]
18	[0,5]	[6,7]	[8,18]	[0,3]	[4,7]	[8,18]
19	[0,5]	[6,8]	[9,19]	[0,3]	[4,8]	[9,19]
20	[0,6]	[7,8]	[9,20]	[0,3]	[4,8]	[9,20]

^aStage 1 ends once a deescalation is invoked.

^bStage 2 ends and the dose is declared safe once the stopping criterion is reached.

The operating characteristics of the ASCENT design under three least favorable

configurations are given in Table 13.1, along with the CRM. The results verify the two theoretical properties of SDSU-SPRT:

- The probability of selecting an overdose is less than $\alpha_0 = 0.10$
- The probability of selecting the MAXSD is at least $1 - \alpha_1 = 0.40$

In fact, the actual performance of SDSU-SPRT is much better than what the nominal values α_0 and $1 - \alpha_1$ suggest under these three particular toxicity scenarios. It shows that the constraints (13.7) and (13.8) can be quite conservative as they are intended for a wide class of dose–toxicity curves.

13.2.5 Practical Notes

While the SDSU-SPRT is tailored for the ASCENT trial where the true dose–toxicity curve is flat at the lower tail, it is interesting to see how the method will compare with the CRM for situations when the CRM is expected to work. Table 13.3 shows the operating characteristics of a two-stage CRM and the ASCENT design under toxicity configurations where the surrogacy MTD v is uniquely defined. Interestingly, the SDSU-SPRT is comparable (or even slightly superior) to the CRM in terms of the selection probability with comparable *average* sample size.

Table 13.3 The operating characteristics of a two-stage CRM and the SDSU-SPRT under three strictly increasing dose–toxicity curves

Model	Probability of selecting dose							ATN	ASN
	0	1	2	3	4	5	6		
Scenario S1	.00	.10	.20	.25	.40	.50	.60		
CRM	.00	.03	.32	.52	.12	.00	.00	12	50
SPRT	.01	.05	.16	.65	.12	.01	.01	17	53
Scenario S2	.00	.05	.05	.18	.25	.40	.50		
CRM	.00	.00	.01	.26	.59	.14	.00	12	50
SPRT	.00	.00	.04	.16	.67	.12	.01	16	54
Scenario S3	.00	.00	.01	.05	.15	.25	.45		
CRM	.00	.00	.00	.00	.19	.70	.11	11	50
SPRT	.00	.00	.00	.02	.18	.75	.04	13	47

Note: The details of the CRM are given in Table 13.1. Numbers associated with the MTD are in bold.

This advantage, however, should be viewed in light of the fact that SDSU-SPRT is an open-ended procedure. Based on Table 13.2, it is possible for the SPRT to enroll more than 20 subjects to a dose in stage 1 if $Z_i(j)$ stays in the continuation region (defined by the intervals under the Stay column in the table). While there is only a small probability for the SPRT to test such a large sample size without reaching a conclusion, the open-endedness of the SPRT renders the procedure undesirable and infeasible in practice. In contrast, the CRM with a fixed sample size $N = 50$ fits well with issues involving budget and scheduling.

There are two approaches to prevent an open-ended trial. First, one could use a statistical test that has with a maximum sample size. For example, the LRT prescribes $\mathcal{R}_{il} = \{Z_j(N_l) \leq c_l\}$ where the constants $c_1 \leq c_2$ and $N_1 \leq N_2$ are fixed in advance, thereby imposing a maximum total sample size KN_2 . Based on extensive simulation in Cheung [18], SDSU-LRT is comparable to SDSU-SPRT (and the CRM) in terms of accuracy but will require a slightly larger average sample size. Second, one could apply a truncated SPRT in the SDSU procedure [114]. This approach has theoretical advantage because the SPRT is optimal under the hypothesized values [109]. On the other hand, while truncating a single SPRT is straightforward, more work is needed to identify an optimal way to truncate multiple SPRTs with a given total maximum sample size.

In addition to the choice of the target rate θ , the MAXSD paradigm requires the specification of an unsafe rate ϑ . In contrast, no such choice is needed (explicitly) for the CRM. In the ASCENT trial, the choice $\theta = 0.25$ is quite convincing as no benefit from a PT regimen is expected to outweigh an increase in the adverse event, which can be quite severe. The choice of ϑ is not as clear. In theory, since there is no explicit control of selection probability under toxicity scenarios outside class (13.3), we are in a sense indifferent to whether a dose with $p_i \in (\theta, \vartheta)$ is considered safe. Thus, we may choose ϑ by asking the question whether selecting a dose with, say, 35% adverse event rate is still acceptable. However, in many situations, the choice of ϑ may be driven by feasibility; generally, the further θ and ϑ are apart, the smaller the sample size is needed to satisfy both (13.7) and (13.8). As a pragmatic approach, we may calibrate the SDSU procedure with respect to the error constraints for given ϑ and θ , but will also examine the performance of the procedure under toxicity scenarios that contain the indifference zone. For example, although the SDSU design for the ASCENT is calibrated with $\vartheta = 0.45$, the results in Table 13.3 show that the method selects a dose with $p_i = 0.40 \in (0.25, 0.45)$ with a reasonably low probability under Scenarios S1 and S2. From a theoretical viewpoint, Cheung [18] proves that the SDSU procedure will select an overdose, that is, $p_i > \theta$, less often as p_i increases. This property is called *unbiasedness*, which provides some assurance on the nonselection of an overdose even though there is no explicit control of the probability of selecting dose i when $p_i \in (\theta, \vartheta)$.

13.3 Bivariate Dose Finding

The bortezomib trial was designed as a phase I/II study, with a dose finding stage that solely used toxicity and a testing stage where all patients were treated at the estimated MTD. The ultimate objective was efficacy-driven, with the hope to demonstrate that the identified dose was promising in terms of 2-year progression-free survival [90]; but the toxicity-driven dose finding stage was an appropriate approach for two reasons. First, toxicity in this trial might be viewed as a surrogate for efficacy. Second, the use of a meaningful efficacy endpoint, 2-year progression-free survival in this case, would take too long to observe for sequential dose finding in an early phase trial, even with a TITE approach (cf. Chapter 11). This second point renders dose finding based on efficacy as infeasible in many situations.

In some situations, a short-term biomarker may be available and can be used as an efficacy endpoint. When this is the case, it is natural to perform dose finding using the efficacy endpoint as well as the toxicity endpoint, so as to avoid the concern whether toxicity is an adequate surrogate. For clarity in classification, a bivariate dose finding design is defined as a design that uses both toxicity and efficacy as bases of interim dose decisions. According to this definition, the bortezomib trial design is not a bivariate design.

While the CRM in its original prescription is not applicable to dealing with two endpoints, the notion of continual reassessment recurs in the increasing literature on bivariate dose finding; just to name a few, see O'Quigley et al. [76], Braun [11], Thall and Cook [102], and Yin et al. [115]. A major theme in most of the methods is the use of model-based and utility-based decisions. To illustrate, consider a trial with a binary efficacy response as well as a binary toxicity outcome. Let $F_E(x, \beta)$ and $F_T(x, \beta)$, respectively, denote the model-based probabilities of achieving efficacy and toxicity at dose x , indexed by a vector-valued parameter β . Then a model-based utility $\mathcal{U}(x, \beta)$ of dose x is defined as a function $F_E(x, \beta)$ and $F_T(x, \beta)$. Common choices of the utility function include

- Probability ratio:

$$\mathcal{U}(x, \beta) = \frac{F_E(x, \beta)}{F_T(x, \beta)};$$

- Odds ratio:

$$\mathcal{U}(x, \beta) = \frac{F_E(x, \beta)\{1 - F_T(x, \beta)\}}{\{1 - F_E(x, \beta)\}F_T(x, \beta)};$$

- Success probability (under independence of toxicity and efficacy):

$$\mathcal{U}(x, \beta) = F_E(x, \beta)\{1 - F_T(x, \beta)\}.$$

The basic idea is to select the dose with the highest utility value. An immediate practical problem of this utility-based approach is the choice of utility function. Figure 13.1 plots the probability ratio, the odds ratio, and the success probability for a given dose–toxicity curve and a given dose–efficacy curve. It is clear that different utility function \mathcal{U} can lead to different recommendations: low doses are preferred based on the probability ratio or odds ratio utilities, whereas the trend is reversed when success probability is used. This is hardly surprising because, by projecting (F_E, F_T) on a two-dimensional grid onto a one-dimensional utility \mathcal{U} , ambiguous orderings among some pairs of (F_E, F_T) are unavoidable. Therefore, clinician inputs are crucial in choosing the utility \mathcal{U} so as to make the utility-based designs clinically relevant.

This being the case, there are useful ideas that may help enhance the robustness of the trial objective and conduct:

Acceptable doses. It is often the case that a utility function does not make clinical sense on the entire dose range. For example, based on the probability ratio utility in Figure 13.1a, the zero dose is the optimal dose. But this is simply due to the mathematical artifact that a virtually zero toxicity probability at zero dose unduly inflates




Figure 13.1 Some examples of utility function (dark solid) for a given dose–toxicity curve (light solid) and a given dose–efficacy curve (dashed). For a trial with toxicity tolerance $\theta_T = 0.25$ and efficacy threshold $\theta_E = 0.50$, the range of acceptable doses is indicated by the dark region on the x-axis and is bracketed by two vertical dotted lines in each panel. The cross “+” in each panel indicates the acceptable dose with the largest utility.

its probability ratio. The odds ratio utility also has this problem, as indicated in Figure 13.1b. However, based on the success probability in Figure 13.1c, the optimal dose is about 0.80, which is associated with a 30% toxicity rate. This dose may be considered unacceptable, based on safety. It leads to the notion of acceptable doses: a dose x is *acceptable* if

$$F_E(x, \beta) \geq \theta_E \text{ and } F_T(x, \beta) \leq \theta_T \quad (13.15)$$

for some prespecified efficacy threshold θ_E and toxicity tolerance θ_T . Then one may view bivariate dose finding as a constrained utility maximization problem subject to (13.15). Bayesian versions of (13.15) have also been adopted in the literature [102].

Alternatively, one may set a bivariate dose finding objective as one of identifying *an* acceptable dose (instead of utility maximization). This objective is less ambitious than the utility-based approach, and is an easier estimation problem. In early-phase trials where sample sizes are limited, this objective may be the only feasible option. Besides feasibility, this objective also has some clinical merits. For one thing, the constraints (13.15) are easy to interpret. Furthermore, the constraint is effective at reducing the differences among the utilities. In Figure 13.1, the acceptable doses range roughly from 0.5 to 0.7: the probability ratio and odds ratio utilities prefer 0.5 on this dose range, and the success probability is maximized at 0.7. Thus, the discrepancy among different utility functions is much reduced on the dose scale. And importantly, utilities of the acceptable doses have a relatively narrow range. The implication is that if we can identify an acceptable dose, we are not far off from some optimal dose regardless of which utility function we use.

Coherence. The notion of coherence can be easily extended to bivariate dose finding where the toxicity and efficacy endpoints are binary. With bivariate binary outcomes, there are four possible outcome configurations for each patient:

- Response, toxic

- Nonresponse, toxic
- Response, nontoxic
- Nonresponse, nontoxic

For each of these four possible outcomes, Table 13.4 lists the subsequent moves that are not allowed due to coherence considerations. As discussed earlier in Chapter 5, escalation is incoherent after a toxic outcome, and deescalation is incoherent after a nontoxic outcome. Also, it seems ethically difficult to give the next patient a lower and, hence, less efficacious dose if the current dose does not yield an efficacy response in the current patient. Thus, if the current subject has a nonresponse and a toxic outcome, the coherence restrictions will allow neither escalation nor deescalation. It then requires an unambiguous decision that the next patient should receive the same dose as the current patient; except that when there is strong evidence that the current dose is unacceptable with respect to constraints (13.15), we should terminate the trial and declare that there is no acceptable dose. While bivariate designs in the literature typically involve complex decision algorithms and computations, this example shows how coherence can simplify clinical decisions that are ethically defensible.

Table 13.4 Incoherent moves under four possible outcomes

	Toxic outcome	Nontoxic outcome
Response	Incoherent escalation: toxicity	Incoherent deescalation: toxicity
Nonresponse	Incoherent escalation: toxicity Incoherent deescalation: efficacy	Incoherent deescalation: toxicity, efficacy

Our discussion on bivariate dose finding so far assumes we have a good estimate of the dose–toxicity F_T and dose–efficacy F_E curves. However, this is an assumption that cannot be taken for granted. As discussed in Chapter 12, for the toxicity-only CRM, an elaborate multiparameter model is not necessarily a good idea because it may generate a rigid sequential dose finding design. At the same time, a one-parameter model (i.e., β is a scalar in both F_E and F_T) may not be adequately flexible to approximate all possible true states of nature. Mimicking Condition 4.4, it may be postulated that the model F_E and F_T should be adequately flexible in the following sense:

Condition 13.1. *For any given $p_T, p_E \in (0, 1)$ and a given x , there exists β in the parameter space such that $F_E(x, \beta) = p_E$ and $F_T(x, \beta) = p_T$.*

While there is currently no unified modeling theory for bivariate dose finding, one possible choice is to extend the conditions given in Chapter 4, Section 4.3 for the bivariate setting and examine their implications. Model calibration in bivariate

dose finding (for model-based designs) will be an interesting and very practical area to pursue.

Stochastic Approximation

14.1 Introduction

Robbins and Monro (1951) introduce the first stochastic approximation method to address the problem of finding the root of a regression function $M(x)$. Precisely, let $Y = Y(x)$ denote a random outcome of interest at the stimulus level x with expectation $E(Y) = M(x)$. The objective is to sequentially approach the root x^* of the equation

$$M(x^*) = \theta \tag{14.1}$$

for a given θ . In the special case when Y is a binary indicator for toxicity and the stimulus level is dose, the regression function is equal to the toxicity probability at dose x , that is, $M(x) = \pi(x) = \Pr(Y = 1 \mid x)$, and the root x^* is the θ th percentile on the dose–toxicity curve π . Using dose finding terminology, solving (14.1) amounts to estimating the MTD under the surrogacy perspective (Definition 2.1). As such, the Robbins-Monro stochastic approximation is a natural method for dose finding in clinical trials. In this chapter we explore the clinical relevance of the stochastic approximation method via its connections to the CRM. Section 14.2 reviews the Robbins-Monro method and some of its refinements, and discusses their implications on the CRM. Section 14.3 presents some adaptations of stochastic approximation for dose finding trials. Section 14.4 discuss some future directions in dose finding methodology. Section 14.5 ends this chapter with some technical details regarding the regularity conditions for $M(x)$ and $Y(x)$.

14.2 The Past Literature

14.2.1 The Robbins-Monro Procedure

The Robbins-Monro procedure starts the experiment at a predetermined level x_1 and estimates x^* with successive approximation x_i for $i > 1$ recursively:

$$x_{i+1} = x_i - \frac{1}{ib} \{Y(x_i) - \theta\} \text{ for some } b > 0. \tag{14.2}$$

Under mild assumptions about the regression function $M(x)$ and the outcome $Y(x)$, recursion (14.2) yields a consistent sequence for x^* , that is, $x_i \rightarrow x^*$ with probability

one. If in addition, the constant b is chosen such that $b < 2M'(x^*)$, then $i^{1/2}(x_i - x^*)$ is asymptotically normally distributed with variance

$$\frac{\sigma_0^2}{b(2\beta - b)} \quad (14.3)$$

where $\beta = M'(x^*)$ and $\sigma_0^2 = \text{var}\{Y(x^*)\}$. The Robbins-Monro procedure is a non-parametric method in that these convergence results depend only very weakly on the underlying distribution of $Y(x)$ and the regression function $M(x)$. The conditions for the convergence results are given in Section 14.5.

In practice, the choice of b in the Robbins-Monro method (14.2) is crucial. First, for asymptotic normality to hold, the constant b is assumed to be less than the upper bound 2β which is unknown. Second, in view of efficiency, the asymptotic variance (14.3) attains its minimum σ_0^2/β^2 when b is set to β . Since β is unknown in most situations, these theoretical results are not directly applicable. But suppose that we can estimate β from the data already observed. Then it is natural to use that estimate in (14.2) instead of the constant b . To be exact, an *adaptive* stochastic approximation is defined as

$$x_{i+1} = x_i - \frac{1}{ib_i} \{Y(x_i) - \theta\} \quad (14.4)$$

where $\{b_i\}$ is a sequence of positive random variables such that $b_i \rightarrow \beta$. In particular, Lai and Robbins [58] study the adaptive procedure (14.4) with b_i set to a truncated least squares estimate $\hat{\beta}_i$ of β as in a linear model with

$$M(x) = \theta + \beta(x - x^*). \quad (14.5)$$

Write $Y_j = Y(x_j)$, and let \bar{x}_i and \bar{Y}_i denote the sample means of x_j s and Y_j s based on the first i observations. Then we may set $b_i = \max\{\underline{b}, \min(\hat{\beta}_i, \bar{b})\}$ for some prespecified \underline{b}, \bar{b} , with

$$\hat{\beta}_i = \frac{\sum_{j=1}^i (x_j - \bar{x}_i)(Y_j - \bar{Y}_i)}{\sum_{j=1}^i (x_j - \bar{x}_i)^2}. \quad (14.6)$$

It can be proved that $b_i \rightarrow \beta$ and that $\{x_i\}$ is asymptotically optimal, that is, having asymptotic variance σ_0^2/β^2 . In the same line of work, Lai and Robbins [57] describe the conditions for the general choice of $\{b_i\}$ that will guarantee asymptotic normality of the adaptive recursion (14.4).

While the adaptive recursion is theoretically ideal, its implementation in practice can be quite limited because estimating the slope β requires a substantial amount of data. In dose finding trials with small samples, the additional variation due to the least squares estimator $\hat{\beta}_i$ may outweigh its benefits.

14.2.2 Maximum Likelihood Recursion

Suppose that we consider a continuous outcome Y with mean specified by the simple linear model (14.5), and suppose further that we believe the slope $\beta = b$. Then it is natural to estimate x^* with the least squares estimate

$$x_i^* = \bar{x}_i - b^{-1}(\bar{Y}_i - \theta) \quad (14.7)$$

repeatedly and set the next design point $x_{i+1} = x_i^*$ in a sequential manner. Lai and Robbins [57] show that the design $\{x_i\}$ generated by repeated least squares recursion is identical to the nonadaptive Robbins-Monro method if the same value of b is used in (14.2) and (14.7). This identity is quite interesting, as it establishes an unlikely link between the nonparametric (14.2) and the least squares estimate (14.7) under a parametric model (14.5) with a strong assumption on the slope β . An implication is that one may use a parametric model $\tilde{M}(x)$ as a *working model* to generate a design sequence, which will be consistent for x^* even though $\tilde{M}(x)$ is not a correct specification of the true $M(x)$. For linear model (14.5), the repeated least squares (14.7) is consistent for x^* if b is not equal to the true β —in fact, consistency holds even if the true $M(x)$ is not linear.

This identity motivates the study of *maximum likelihood recursion* for data arising from the exponential family [113]. First, by postulating that $\tilde{M}(x)$ depends on x via $x - x^*$, we may rewrite the regression function as $\tilde{M}_c(x - x^*)$ such that $\tilde{M}_c(0) = \theta$; see the linear model (14.5), for example. Then, the design sequence $\{x_i\}$ is defined recursively based on the maximum likelihood estimate of x^* , that is, obtain x_i^* by solving

$$\sum_{j=1}^i w(x_j) \{Y_j - \tilde{M}_c(x_j - x_i^*)\} = 0 \tag{14.8}$$

for some appropriately chosen and prespecified weight function $w(x)$, and then

$$x_{i+1} = \max\{x_{\min}, \min(x_i^*, x_{\max})\} \tag{14.9}$$

where x_{\min}, x_{\max} are, respectively, the lower and upper limits of the dose range. Ying and Wu [116] show that maximum likelihood recursion (14.9) can be expressed as an adaptive stochastic approximation with large samples, so that they can apply the standard results for stochastic approximation to establish the asymptotic properties of the maximum likelihood recursion. Briefly, under the regularity conditions specified in Section 14.5, the recursion formed by (14.8) and (14.9) is consistent for x^* and is asymptotically normal. These results hold without assuming the working model \tilde{M}_c is a correct specification of the true function M .

14.2.3 Implications on the CRM

For binary outcome Y , Wu [112] proposes the logit-MLE that is a special case of the maximum likelihood recursion, using the logistic function as a working model

$$\tilde{M}(x) = \tilde{M}_c(x - x^*) = \frac{\theta \exp\{b(x - x^*)\}}{1 - \theta + \theta \exp\{b(x - x^*)\}} \tag{14.10}$$

for some $b > 0$, with $w(x) \equiv 1$ in the estimating equation (14.8). Using the results in Ying and Wu [116], we can show that the logit-MLE is asymptotically normal, that is,

$$i^{1/2}(x_i - x^*) \rightarrow_{\mathcal{L}} N\left(0, \frac{1}{b\{2M'(x^*) - b\theta(1 - \theta)\}}\right). \tag{14.11}$$

if $\tilde{M}'_c(0) = b\theta(1 - \theta) < 2M'(x^*)$.

If we disregard the truncation (14.9), the logit-MLE effectively places the next subject at the maximum likelihood estimate x_i^* . This will yield the same design as the likelihood CRM (3.11) if there always exists a test dose d_k with $\tilde{M}(d_k - x_i^*) = \theta$. Thus, the CRM may be viewed as an analog of the logit-MLE for situations that only allow a discrete set of test doses. This analogy between these two methods gives some practical insights.

First, the one-parameter CRM is now justified by its (asymptotic) connection to the nonparametric stochastic approximation via the logit-MLE, which uses a one-parameter location model (14.10). While the theoretical properties of the CRM have been studied in the recent dose finding literature (cf., Chapter 5), it is reassuring that similar model recommendations can be independently reached by a comparison to the rich literature on stochastic approximation. As a case in point, for the logit-MLE to converge at a \sqrt{n} -rate as prescribed in (14.11), the working model \tilde{M}_c needs to be relatively shallow, that is, $b < 2M'(x^*)/\{\theta(1 - \theta)\}$. To translate its implication for the CRM, this suggests that a shallow skeleton $\{p_{0k}\}$ should be used. In Chapter 8, we see that it is asymptotically desirable to choose a working model so that the CRM has a narrow indifference interval, which indeed corresponds to a shallow set of p_{0k} s; see Exercise 8.2. This recommendation thus concurs with that based on the large sample properties of the logit-MLE.

Second, while the justification of the model-based logit-MLE relies on its asymptotic equivalence to the nonparametric Robbins-Monro procedure, the former has been demonstrated to be superior to the latter in finite sample settings [50, 112]. Thus, a model-based approach seems to retain information that is otherwise lost when using nonparametric procedures. This speculation is made without assuming much confidence about the working model. In a typical dose finding trial with binary data and small sample sizes, this efficiency comparison lends credence to the use of model-based methods such as the CRM over the nonparametric and algorithm-based designs.

14.3 The Present Relevance

14.3.1 Practical Considerations

While the work of Robbins and Monro has motivated numerous refinements and generated a large number of applications in the statistical and engineering literature, its clinical applications have been rare—for a number of practical considerations. We elucidate the three main issues here.

First, as mentioned in the previous section, with a binary outcome and small samples, the Robbins-Monro method is generally less efficient than the model-based methods, and hence may not be suitable for clinical dose finding where the study endpoint is classified as toxic versus nontoxic.

Second, the method entails the availability of a continuum of doses. As discussed in Chapter 3 in the context of the bortezomib trial and the ASCENT trial, this is unrealistic. In the NeuSTART, where ordering of the test regimens is determined on the sole basis of the dose of a drug, dose availability is limited by the dosage of a

tablet and convenience of drug preparation. To accommodate the clinical situations that allow only a discrete set of test doses, one may naturally use a discretized version of the stochastic approximation. Suppose that the test doses are labeled $\{1, \dots, K\}$. A discretized stochastic approximation yields the next dose via

$$x_{i+1} = C \left\{ x_i - \frac{1}{ib} (Y_i - \theta) \right\} \quad (14.12)$$

where $C(x)$ is the rounded value of x for $x \in [0.5, K + 0.5)$, and is equal to 1 and K , respectively, if $x < 0.5$ or $\geq K + 0.5$. This discretized stochastic approximation is rigid (Definition 12.1) and may be stuck at a wrong dose indefinitely. To illustrate, consider applying (14.12) with $b = 0.5$, $\theta = 0.1$, and a starting dose $x_1 = 1$. It can be shown that $x_2 = C(1.2) = 1$ if $Y_1 = 0$ and $x_2 = C(-0.8) = 1$ if $Y_1 = 1$, that is, the second patient will be treated at dose level 1 regardless of the first outcome Y_1 . In the same manner, we can further verify that the trial will in fact never leave dose 1. Generally, using simple algebra, we can show that the update x_{i+1} in accordance with (14.12) will stay the same as x_i if

$$\left| \frac{1}{ib} (Y_i - \theta) \right| < 0.5. \quad (14.13)$$

For binary outcomes, Condition (14.13) is implied by $i > 2 \max(\theta, 1 - \theta)/b$. For outcomes of a general variable type, the likelihood that (14.13) occurs is

$$\Pr \left(\left| \frac{1}{ib} (Y_i - \theta) \right| < 0.5 \right) \geq 1 - \frac{4E(Y_i - \theta)^2}{i^2 b^2}$$

by Chebyshev's inequality. That is, for any given b and any outcome type (with a finite variance), the discretized recursion (14.12) will quickly be stuck at a dose level, which may or may not be the correct dose. This rigidity problem is a result of rounding the stochastic approximation output to a discrete set of levels, so that the term $(ib)^{-1}(Y_i - \theta)$ does not contribute to future updates once i gets sufficiently large. Hence, stochastic approximation cannot be directly applied to dose finding trials with naive discretization.

Third, at small sample sizes, especially in the early enrollment stage, the least squares estimator (14.6) will fluctuate with substantial variation, and can take on a negative value or is undefined (when all patients are treated at the same dose). The low informational content of a binary outcome aggravates this numerical instability. As a result, the prespecified truncated values \underline{b} and \bar{b} will be in effect for most part of the trial. This is in essence running a nonadaptive procedure. Thus, the choice of b in the recursion (14.2) or \underline{b} and \bar{b} for (14.4) remains a practical issue.

14.3.2 Dichotomized Data

In many situations, the binary toxicity outcome Y is defined by dichotomizing an observed quantitative biomarker Q , that is, $Y = I\{Q \geq t_0\}$ for some prespecified threshold t_0 . The quantitative measurement Q apparently contains more information

than Y , and may be used to estimate x^* in (14.1) with greater efficiency. Specifically, consider the regression model

$$Q = R(x) + \sigma_Q(x)e \quad (14.14)$$

where e is an unobserved noise from a known distribution η that has mean 0 and unit variance, so that $E(Q) = R(x)$ and $\text{var}(Q) = \sigma_Q^2(x)$. The mean function $R(x)$ is assumed to satisfy the regularities for the stochastic approximation specified in Section 14.5. Note that model (14.14) allows heteroscedasticity in the outcome at different doses. Since biological lab values tend to exhibit larger variations at higher doses, an assumption of homoscedasticity, that is, $\sigma_Q(x) \equiv \sigma_Q$, is probably too strong for practical purposes. Figure 14.1a shows that variability of the liver enzymes in the NeuSTART increases with dose as the mean increases. To make the modeling effort parsimonious, we further assume that the variability $\sigma_Q(x)$ increases in x via the mean function $R(x)$. Apart from these, the forms of $R(x)$ and $\sigma_Q(x)$ are unspecified. Finally, to identify the distribution η , the normal QQ-plot in Figure 14.1b shows that the noise e fits well to a standard normal distribution.

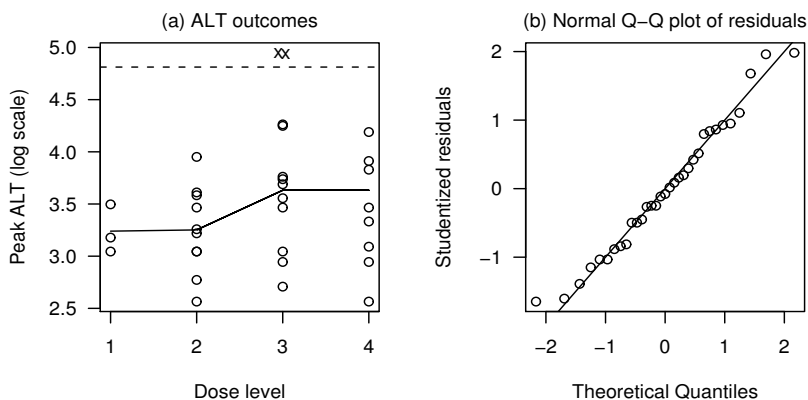


Figure 14.1 (a) The liver enzyme data in the NeuSTART. (b) A diagnostic plot for the normality assumption of the noise.

Now, under (14.14), the toxicity probability at dose x , that is, also the mean of Y at dose x , can be expressed as

$$M(x) = 1 - \eta \left\{ \frac{t_0 - R(x)}{\sigma_Q(x)} \right\}.$$

We can then show that solving the objective (14.1) is equivalent to solving

$$f(x) = t_0 \quad (14.15)$$

where $f(x) = R(x) + c_\theta \sigma_Q(x)$ and c_θ is the upper θ th percentile of η .

Suppose for the moment that a continuum of dose x is available, and that patients are enrolled in small groups of size m . Let x_i denote the dose given to the i th group, and Q_{ij} the measurement of the j th patient in the group. Then define

$$O_i = \bar{Q}_i + \left[E \left\{ \frac{S_i}{\sigma_Q(x_i)} \right\} \right]^{-1} c_\theta S_i \tag{14.16}$$

where

$$\bar{Q}_i = \frac{1}{m} \sum_{j=1}^m Q_{ij} \text{ and } S_i^2 = \frac{1}{m-1} \sum_{j=1}^m (Q_{ij} - \bar{Q}_i)^2$$

are, respectively, the sample mean and variance of all measurements in group i . There are two important facts about the variable O_i :

- O_i is observable.
- $E(O_i) = f(x_i)$.

See Exercises 14.1 and 14.2 for details. As a result of these two facts, we can use the Robbins-Monro method based on O_i s to solve the objective (14.15):

$$x_{i+1} = x_i - \frac{1}{ib} (O_i - t_0). \tag{14.17}$$

Following standard convergence results, the recursion (14.17) can be shown to be consistent for x^* , and if in addition $b < 2f'(x^*)$, the sequence $\{x_i\}$ is asymptotically normal with variance

$$v_O = \text{var}\{O_i(x^*)\} [b\{2f'(x^*) - b\}]^{-1}.$$

In particular, when e is standard normal, for $m \geq 2$,

$$v_O = \frac{\sigma_Q^2(x^*) \{1 + mc_\theta^2(\lambda_m - 1)\}}{mb\{2f'(x^*) - b\}}$$

where

$$\lambda_m = \frac{(m-1)\Gamma^2\{(m-1)/2\}}{2\Gamma^2(m/2)}. \tag{14.18}$$

and $\Gamma(\cdot)$ is the gamma function.

The efficiency gain due to the use of the quantitative measurement Q can be measured by comparing v_O to the asymptotic variance of the logit-MLE. We can easily modify the result (14.11) and obtain the asymptotic variance of the logit-MLE with a group size of $m \geq 2$:

$$v_Y = \frac{1}{mb\{2M'(x^*) - b\theta(1-\theta)\}}.$$

The asymptotic variance v_O is minimized by setting $b = f'(x^*)$ in (14.17), and v_Y by setting $b = M'(x^*)/\{\theta(1-\theta)\}$ in (14.10). Thus, the optimal choice of b depends

on the unknown $f'(x^*)$ and $M'(x^*)$ in both procedures. For the purpose of comparing efficiencies, however, suppose we know these optimal values. Then, for normal noise, the variance ratio is equal to

$$\frac{v_Y}{v_O} = \frac{\theta(1-\theta)}{\eta'(c_\theta)\sigma_Q^2(x^*)\{1+mc_\theta^2(\lambda_m-1)\}} \quad (14.19)$$

where $\eta'(x)$ is the standard normal density. For $m=3$, the minimum value of the ratio (14.19) is 1.238 when the target $\theta=0.12$ or 0.88. The variance ratio, also the asymptotic relative efficiency of the recursion (14.17), increases as an extreme θ is targeted [19].

While the efficiency gain due to the use of a quantitative biomarker Q is expected, the use of (14.17) poses two practical issues additional to those raised previously in Section 14.3.1. First, the gain is achieved by making an additional assumption on the noise distribution η , whereas the logit-MLE makes no such assumption as the design uses only the dichotomized data Y . Therefore, it would be useful to get some sense of η by looking at pilot data before the trial, and perform diagnostics afterward (cf., Figure 14.1b). Some preliminary investigation of the impact of a misspecified η is given in [22], and further study is warranted.

A second issue is that, in order to be able to estimate $\sigma_Q(x_i)$, it is necessary to enroll patients in groups with size $m \geq 2$. The logit-MLE, in contrast, can be applied in a fully sequential manner. This issue does not cause much practical difficulty, as the clinicians are used to the group accrual designs such as the 3+3 algorithm.

14.3.3 Virtual Observations

To address the issue with a discrete design space, Cheung and Elkind [22] introduce the notion of virtual observations. Consider using model (14.14) and using the same notation developed in the previous subsection, we may define the virtual observation of the i th group of patients as

$$V_i = O_i + b(x_i^* - x_i) \quad (14.20)$$

where x_i^* denotes the *assigned* dose of the group. In the situations when the dose x_i *actually given* to the group can take on any real value (i.e., a continuum of dose is available for testing), the given dose x_i always equals the assigned dose x_i^* . For the far more common situations where $x_i \in \{1, \dots, K\}$, we may set $x_i = C(x_i^*)$. Now, since the assigned dose x_i^* can take values on a continuous scale, we may generate the assigned sequence $\{x_i^*\}$ by the Robbins-Monro procedure based on the virtual observations V_i s:

$$x_{i+1}^* = x_i^* - \frac{1}{ib}(V_i - t_0) \quad (14.21)$$

To initiate the recursion, we set $x_1^* = x_1 \in \{1, \dots, K\}$ according to the clinical choice of starting dose. Note that the given dose x_i here is obtained by rounding as in the discretized stochastic approximation (14.12). However, the update (14.21) is done on

the assigned dose x^* without rounding. As a result, the term $(ib)^{-1}(V_i - t_0)$, which is of the order $O(i^{-1})$, can be carried over to the future updates.

Since it is possible that $M(k) \neq \theta$ for all k , the dose finding objective is to identify

$$v = \arg \min_k |M(k) - \theta|. \tag{14.22}$$

Under mild assumptions on $R(x)$ and $\sigma_Q(x)$, the assigned sequence $\{x_i^*\}$ converges to v_b for some $v_b \in v \pm 0.5$, and the rounded sequence $\{x_i\}$ to the true v . Briefly, under a dose-toxicity curve where $M(k) \notin (\theta_L, \theta_U)$ for all $k \neq v$, for some $\theta_L < \theta < \theta_U$, consistency requires

$$b < 2\sigma_Q(x^*) \min(c_\theta - c_\theta^2/c_L, c_\theta - c_U) \tag{14.23}$$

where $c_L = \eta^{-1}(1 - \theta_L)$ and $c_U = \eta^{-1}(1 - \theta_U)$. Details of the required assumptions are given in [22].

The idea of virtual observation is to create an objective function

$$h(x) = E(V_i | x_i^* = x) = f\{C(x)\} + b\{x - C(x)\}$$

that is defined on the real line, and has a local slope at $\{1, \dots, K\}$, such that the solution of (14.15) can be approximated by the solution v_b of $h(x) = t_0$. Since now the objective function h has a known slope b around v_b , we can use the same b in the recursion (14.21) as in the definition of virtual observations (14.20). This feature allows us to achieve asymptotic optimality without resorting to adaptive estimation of the slope (cf., Section 14.2). In other words, by using the virtual observation recursion, we also address the practical choice of b . We can show that the asymptotic variance of the assigned sequence $\{x_i^*\}$ is inversely proportional to b^2 . Therefore, in view of efficiency, we should choose b to be as large as the consistency condition (14.23) permits it. Under the monotonicity assumption of $\sigma_Q(x)$, we may estimate the upper bound in (14.23) by replacing $\sigma_Q(x^*)$ with an estimate $\hat{\sigma}_Q^*$ for $\sigma_Q(0)$ using pilot data in untreated patients. Then, we may set b as $2\hat{\sigma}_Q^* \min(c_\theta - c_\theta^2/c_L, c_\theta - c_U)$.

14.3.4 Quasi-Likelihood Recursion

The concept of virtual observations applies to other outcome types. Suppose, for instance, the dose finding objective (14.22) is achieved using the binary outcome Y . Then a virtual observation may be defined as

$$V_i = Y_i + b(x_i^* - x_i) \tag{14.24}$$

so that $E(V_i | x_i^* = x) = M\{C(x)\} + b\{x - C(x)\}$ and the assigned dose sequence may be generated by stochastic approximation:

$$x_{i+1}^* = x_i^* - (ib)^{-1}(V_i - \theta). \tag{14.25}$$

However, since the Robbins-Monro method is not efficient with binary outcomes, it is likely to also be the case for the virtual observation recursion (14.25). We could

therefore apply the virtual observation (14.24) with a model-based approach such as the maximum likelihood recursion (14.9). That is, the assigned sequence $\{x_j^*\}$ is generated by recursively solving

$$\sum_{j=1}^i w(x_j^*) \{V_j - \tilde{M}_c(x_j^* - x_{i+1}^*)\} = 0 \quad (14.26)$$

where the model \tilde{M}_c is the logistic function (14.10). The actual given dose is then obtained by rounding, that is, $x_{i+1} = C(x_{i+1}^*)$. Solving (14.26) corresponds to quasi-likelihood estimation [66]. Since the asymptotic properties of maximum likelihood recursion rely only on the first two moments of the outcome (cf., Section 14.5), we expect that the quasi-likelihood recursion (14.26) has similar asymptotic behaviors under mild assumptions on \tilde{M}_c , although the details will need to be worked out in a rigorous manner.

14.4 The Future Challenge

While simulation is an indispensable tool for design evaluation (this book has relied significantly on it), pathological behaviors such as rigidity and incoherence may not be detected in aggregate via simulations. These pathologies are *pointwise* properties that can only be found by careful analytical study. In this book, the CRM is presented at a rigorous level, so that the calibration techniques are built upon solid theoretical ground. The attempt here, however, primarily addresses the “standard” phase I trial settings where the objective is to determine the MTD defined in terms of binary toxicities over a short-term period. In some chapters, there are discussions on dose finding in specialized situations such as with late toxicities (Chapter 11), bivariate outcomes (Section 13.3), and dichotomized outcomes (Section 14.3.2). While the literature has seen an increasing number of designs for these nonstandard situations, a theoretical framework is yet to be worked out before we can systematically study the calibration of these designs for actual practice.

Obviously, as we deal with different data types other than the binary outcomes, the theoretical criteria presented in this book have to be reexamined. Take coherence, for example. When toxicity is defined by dichotomizing a continuous biomarker, the escalation restrictions should not only take the dichotomized outcomes into account, but also the mean and variability of the continuous observations. As Figure 14.1a suggests, a drug may have a more pronounced impact on the variance of a biomarker than on its mean level. Therefore, it is sensible to use caution in escalation when increased variability is observed. The formal definition of coherence in this setting yet needs further research.

Despite these complications, it is worthwhile to study the theoretical properties of dose finding methods in general, especially for the nonstandard situations. Consider coherence again: its extension to the bivariate case in Section 13.3 provides clear and clinically sensible guidance on an otherwise complex model-based dose-decision blackbox. In addition, as the number of parameters used in a bivariate CRM model increases, the concern for rigidity certainly warrants more scrutiny.

How do we proceed from here then? This section illustrates a promising direction through a connection to the rich stochastic approximation literature. This chapter deals with the root-finding objective (14.1) in accordance with the surrogacy MTD definition in the standard dose finding setting. However, there remains a vast number of untapped ideas in the stochastic approximation literature that can be borrowed to dose finding. For example, one may formulate the bivariate dose finding into a problem of finding the maximum of a regression function, for which the work by Kiefer and Wolfowitz [52] and Fabian [35] may prove relevant. Among topics that are not covered in this book, individualized dosing is an increasingly popular concept; the basic idea is to account for patient heterogeneity by using each patient's baseline covariate [14, 104]. This problem can in fact be reformulated into a multivariate contour-finding problem for which Blum [10] proposes a stochastic approximation. While it can be an arduous task to study the theoretical properties of the extensions of the CRM to these specialized situations, a connection to the well-studied stochastic approximation procedures may allow us to do so with relative ease and elegance, as in the case for the virtual observation recursions (14.21) and (14.26).

14.5 Assumptions on $M(x)$ and $Y(x)$

This section states the conditions on $M(x)$ and $Y(x)$, so that the Robbins-Monro recursion (14.2) will yield a consistent estimate for the root x^* . These conditions are extracted from Sacks [92] with some slight modifications.

Condition 14.1. $M(x^*) = \theta$ and $(x - x^*)\{M(x) - \theta\} > 0$ for all $x \neq x^*$.

Condition 14.2. $\sup_x \text{var}\{Y(x)\} < \infty$ and $\lim_{x \rightarrow x^*} \text{var}\{Y(x)\}$ exists.

Condition 14.3. For some $\varepsilon_1, \varepsilon_2 > 0$,

$$\sup_{|x - x^*| < \varepsilon_1} E |Y(x)|^{2 + \varepsilon_2} < \infty.$$

Condition 14.1 ensures that the equation (14.1) has a unique solution, so that the problem is well defined. Conditions 14.2 and 14.3 impose $Y(x)$ has finite higher order moments; these two conditions are satisfied when $Y(x)$ is a binary variable. Also, we can often modify Condition 14.2 with a slightly stronger condition that says $\text{var}\{Y(x^*)\} = \sigma_0^2$ for some $\sigma_0 > 0$. Under these three conditions, we can prove strong consistency of the Robbins-Monro procedure (14.2), that is, $x_i \rightarrow x^*$ with probability one, using the results in Robbins and Siegmund [87].

Condition 14.4. For some positive constant C , and for all x , $|M(x) - \theta| \leq C|x - x^*|$.

Condition 14.5. $M(x) = \alpha_0 + \alpha_1(x - x^*) + \tau(x, x^*)$ such that $\tau(x, \theta) = o(|x - x^*|)$ as $x \rightarrow \theta$.

Conditions 14.4 and 14.5 impose mild restrictions on the form of the regression function $M(x)$. These two conditions, together with Conditions 14.1–14.3 and the assumption that $b < 2M'(x^*)$, are sufficient for the asymptotic normality of $\{x_i\}$, that is,

$$i^{1/2}(x_i - x^*) \rightarrow_{\mathcal{L}} N\left(0, \frac{\sigma_0^2}{b\{2M'(x^*) - b\}}\right)$$

where $\sigma_0^2 = \lim_{x \rightarrow x^*} \text{var}\{Y(x)\}$. It can be easily verified that the asymptotic variance is minimized when the constant b in the recursion (14.2) is set to $M'(x^*)$.

14.6 Exercises and Further Results

Exercise 14.1. Consider regression model (14.14) with a standard normal noise e .

a. Show that

$$E\left\{\frac{S_i}{\sigma_Q(x_i)}\right\} = \frac{1}{\sqrt{\lambda_m}}$$

where λ_m is defined in (14.18).

b. For the observed outcome O_i defined in (14.16), verify that $E(O_i) = f(x_i)$ and show that

$$\text{var}(O_i) = \frac{\sigma_Q^2(x_i)}{m} \{1 + mc_\theta^2(\lambda_m - 1)\}.$$

Exercise 14.2. Repeat Exercise 14.1 for any random noise with $E(e) = 0$ and $\text{var}(e) = 1$. You may consider the following steps:

a. Describe how you will calculate $E\{S_i/\sigma_Q(x_i)\}$ by simulation.

b. Define $\lambda_m = [E\{S_i/\sigma_Q(x_i)\}]^{-2}$. Evaluate λ_m for a logistic noise e and $m = 2, 3, 4$.

c. Verify that $E(O_i) = f(x_i)$ and show that

$$\text{var}(O_i) = \frac{\sigma_Q^2(x_i)}{m} \left\{1 + mc_\theta^2(\lambda_m - 1) + 2mc_\theta \sqrt{\lambda_m} E(\bar{e}_i S_i / \sigma_Q(x_i))\right\}.$$

Exercise 14.3. Under model (14.14), show that $M'(x^*) = \eta'(c_\theta)f'(x^*)/\sigma_Q(x^*)$, and hence verify the variance ratio (14.19) for normal noise.

References

- [1] C. Ahn. An evaluation of phase I cancer clinical trial designs. *Statistics in Medicine*, 17:1537–1549, 1998.
- [2] D. Anbar. Stochastic approximation methods and their use in bioassay and phase I clinical trials. *Communications in Statistics*, 13:2451–2467, 1984.
- [3] M. Ayer, H. D. Brunk, G. M. Ewing, W. T. Reid, and E. Silverman. An empirical distribution function for sampling with incomplete information. *Annals of Mathematical Statistics*, 26:641–647, 1995.
- [4] J. Babb, A. Rogatko, and S. Zacks. Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statistics in Medicine*, 17:1103–1120, 1998.
- [5] D. Bartholomew. Isotonic inference. In *Encyclopedia of Statistical Science: Volume 4*, pages 260–265. New York: Wiley, 1983.
- [6] J. Bartroff and T. L. Lai. Approximate dynamic programming and its applications to the design of phase I cancer trials. *Statistical Science*, 2010 in press.
- [7] R. E. Bechhofer, J. Kiefer, and M. Sobel. *Sequential Identification and Ranking Procedures (with Special Reference to Koopman-Darmois Populations)*. Chicago: University of Chicago Press, 1968.
- [8] B. N. Bekele, Y. Ji, Y. Shen, and P. F. Thall. Monitoring late-onset toxicities in phase I trials using predicted risks. *Biostatistics*, 9:442–457, 2008.
- [9] J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. New York: Springer-Verlag, 1980.
- [10] J. R. Blum. Multidimensional stochastic approximation methods. *Annals of Mathematical Statistics*, 25:737–744, 1954.
- [11] T. M. Braun. The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials*, 23:240–256, 2002.
- [12] T. M. Braun. Generalizing the TITE-CRM to adapt for early- and late-onset toxicities. *Statistics in Medicine*, 25:2071–2083, 2006.
- [13] P. P. Carbone, M. J. Krant, S. P. Miller, T. C. Hall, B. I. Shnider, Colsky J., J. Horton, H. Hosley, J. M. Miller, E. Frie, and M. Schneiderman. The feasibility of using randomization schemes early in the clinical trials of new chemotherapeutic agents: Hydroxyurea. *Clinical Pharmacology and Therapeutics*, 6:17–24, 1965.

- [14] J. D. Cheng, J. S. Babb, C. Langer, S. Aamdal, F. Robert, L. R. Engelhardt, O. Fernberg, J. Schiller, G. Forsberg, R. K. Alpaugh, L. M. Weiner, and A. Rogatko. Individualized patient dosing in phase I clinical trials: the role of Escalation With Overdose Control in PNU-214936. *Journal of Clinical Oncology*, 22:602–609, 2004.
- [15] Y. K. Cheung. On the use of nonparametric curves in phase I trials with low toxicity tolerance. *Biometrics*, 58:237–240, 2002.
- [16] Y. K. Cheung. Coherence principles in dose finding studies. *Biometrika*, 92:863–873, 2005.
- [17] Y. K. Cheung. Dose finding with delayed binary outcomes in cancer trials. In S. Chevret, editor, *Statistical Methods for Dose Finding Experiments*, pages 225–242. New York: Wiley, 2006.
- [18] Y. K. Cheung. Sequential implementation of stepwise procedures for identifying the maximum tolerated dose. *Journal of the American Statistical Association*, 102:1448–1461, 2007.
- [19] Y. K. Cheung. Stochastic approximation and modern model-based designs for dose finding clinical trials. *Statist. Sci.*, 25(2):191–201, 2010.
- [20] Y. K. Cheung and R. Chappell. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics*, 56:1177–1182, 2000.
- [21] Y. K. Cheung and R. Chappell. A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics*, 58:671–674, 2002.
- [22] Y. K. Cheung and M. S. V. Elkind. Stochastic approximation with virtual observations for dose finding on discrete levels. *Biometrika*, 97:109–121, 2010.
- [23] Y. K. Cheung, P. H. Gordon, and B. Levin. Selecting promising ALS therapies in clinical trials. *Neurology*, 67:1748–1751, 2006.
- [24] Y. K. Cheung and P. F. Thall. Monitoring the rates of composite events with censored data in phase II clinical trials. *Biometrics*, 58:89–97, 2002.
- [25] S. Chevret. The continual reassessment method in cancer phase I clinical trials: a simulation study. *Statistics in Medicine*, 12:1093–1108, 1993.
- [26] S. Chevret, editor. *Statistical Methods for Dose Finding Experiments*. New York: Wiley, 2006.
- [27] S. Chevret and S. Zohar. The continual reassessment method. In S. Chevret, editor, *Statistical Methods for Dose Finding Experiments*, pages 131–148. New York: Wiley, 2006.
- [28] C. W. Dunnett and A. C. Tamhane. A step-up multiple test procedure. *Journal of the American Statistical Association*, 87:162–170, 1992.
- [29] S. D. Durham and N. Flournoy. Random walks for quantile estimation. In S.S. Gupta and J.O. Berger, editors, *Statistical Decision Theory and Related Topics V*, pages 467–476. New York: Springer-Verlag, 1994.
- [30] S. D. Durham and N. Flournoy. Up-and-down designs I: Stationery treatment distributions. In N. Flournoy and W. F. Rosenberger, editors, *Adaptive De-*

- signs*, pages 139–157. Hayward, California: Institute of Mathematical Statistics, 1995.
- [31] S. D. Durham, N. Flournoy, and A. A. Montazer-Haghighi. Up-and-down designs II: exact treatment moments. In N. Flournoy and W. F. Rosenberger, editors, *Adaptive Designs*, pages 158–178. Hayward, California: Institute of Mathematical Statistics, 1995.
- [32] S. D. Durham, N. Flournoy, and W. F. Rosenberger. A random walk rule for phase I clinical trials. *Biometrics*, 53:745–760, 1997.
- [33] B. H. Eichhorn and S. Zacks. Sequential search of an optimal dosage, I. *Journal of the American Statistical Association*, 68:594–598, 1973.
- [34] M. S. Elkind, R. L. Sacco, R. B. Macarthur, D. J. Fink, E. Peerschke, H. Andrews, G. Neils, J. Stillman, J. Chong, S. Connolly, T. Corporan, D. Leifer, and K. Cheung. The neuroprotection with statin therapy for acute recovery trial (NeuSTART): an adaptive design phase I dose-escalation study of high-dose lovastatin in acute ischemic stroke. *International Journal of Stroke*, 3:210–218, 2008.
- [35] V. Fabian. Stochastic approximation of minima with improved asymptotic speed. *Annals of Mathematical Statistics*, 38:191–200, 1967.
- [36] D. Faries. Practical modifications of the continual reassessment method for phase I cancer trials. *Journal of Biopharmaceutical Statistics*, 4:147–164, 1994.
- [37] D. J. Finney. *Statistical Methods in Biological Assay*. Griffin, London, 1978.
- [38] M. Fisher, K. Cheung, G. Howard, and S. Warach. New pathways for evaluating potential acute stroke therapies. *International Journal of Stroke*, 1:52–58, 2006.
- [39] E. Garrett-Mayer. The continual reassessment method for dose finding studies: a tutorial. *Clinical Trials*, 3:57–71, 2006.
- [40] M. Gasparini and J. Eisele. A curve-free method for phase I clinical trials. *Biometrics*, 56:609–615, 2000.
- [41] Constantine Gatsonis and Joel B. Greenhouse. Bayesian methods for phase I clinical trials. *Statistics in Medicine*, 11:1377–1389, 1992.
- [42] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors. *Markov Chain Monte Carlo in Practice*. London: Chapman & Hall, 1996.
- [43] S. N. Goodman, M. L. Zahurak, and S. Piantadosi. Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine*, 14:1149–1161, 1995.
- [44] L. M. Haines, I. Perevozskaya, and W. F. Rosenberger. Bayesian optimal design for phase I clinical trials. *Biometrics*, 59:591–600, 2003.
- [45] J. M. Heyd and B. P. Carlin. Adaptive design improvements in the continual reassessment method for phase I studies. *Statistics in Medicine*, 18:1307–1321, 1998.

- [46] J. C. Hsu and R. L. Berger. Stepwise confidence intervals without multiplicity adjustment for dose-response and toxicity studies. *Journal of the American Statistical Association*, 94:468–482, 1999.
- [47] J. Hüsing, W. Sauerwein, K. Hideghéty, and K. H. Jöckel. A scheme for a dose-escalation study when the event is lagged. *Statistics in Medicine*, 20:3323–3334, 2001.
- [48] A. Iasonos, A. S. Wilton, E. R. Riedel, V. E. Seshan, and D. R. Spriggs. A comprehensive comparison of the continual reassessment method to the standard 3+3 dose escalation scheme in phase I dose finding studies. *Clinical Trials*, 5:465–477, 2008.
- [49] Y. Ji, Y. Li, and B. N. Bekele. Dose finding in phase I clinical trials based on toxicity probability intervals. *Clinical Trials*, 4:235–244, 2007.
- [50] V. R. Joseph. Efficient Robbins-Monro procedure for binary data. *Biometrika*, 91:461–470, 2004.
- [51] H. L. Kaufman, S. Cohen, K. Cheung, G. DeRaffele, J. Mithcam, D. Moroziewicz, J. Schlom, and C. Hesdorffer. Local delivery of vaccinia virus expressing multiple costimulatory molecules for the treatment of established tumors. *Human Gene Therapy*, 17:239–244, 2006.
- [52] J. Kiefer and J. Wolfowitz. Stochastic estimation of the maximum of a regression function. *Annals of Mathematical Statistics*, 23:462–466, 1952.
- [53] E. L. Korn. Nontoxicity endpoints in phase I trial designs for targeted, non-cytotoxic agents. *Journal of National Cancer Institute*, 96:977–978, 2004.
- [54] E. L. Korn, D. Midthune, T. T. Chen, L. V. Rubinstein, M. C. Christian, and R. M. Simon. A comparison of two phase I trial designs. *Statistics in Medicine*, 13:1799–1806, 1994.
- [55] R. Kurzrock and R. S. Benjamin. Risks and benefits of phase I oncology trials, revisited. *New England Journal of Medicine*, 352:930–932, 2005.
- [56] T. L. Lai. Stochastic approximation. *Annals of Statistics*, 31:391–406, 2003.
- [57] T. L. Lai and H. Robbins. Adaptive design and stochastic approximation. *Annals of Statistics*, 7:1196–1221, 1979.
- [58] T. L. Lai and H. Robbins. Consistency and asymptotic efficiency of slope estimates in stochastic approximation schemes. *Zeit Schrift Fur Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 56:329–360, 1981.
- [59] C. Le Tourneau, J. J. Lee, and L. L. Siu. Dose escalation methods in phase I cancer clinical trials. *Journal of the National Cancer Institute*, 101:708–720, 2009.
- [60] S. M. Lee and Y. K. Cheung. Model calibration in the continual reassessment method. *Clinical Trials*, 6:227–238, 2009.
- [61] E. L. Lehmann. *Theory of Point Estimation*. New York: Wiley, 1983.
- [62] J. P. Leonard, R. R. Furman, Y. K. Cheung, E. J. Feldman, H. J. Cho, J. M. Vose, G. Nichols, P. W. Glynn, M. A. Joyce, J. Ketas, J. Ruan, J. Carew,

- R. Niesvizky, A. LaCasce, A. Chadburn, E. Cesarman, and M. Coleman. Phase I/II trial of bortezomib plus CHOP-Rituximab in diffuse large B cell (DLBCL) and mantle cell lymphoma (MCL): Phase I results. *Blood*, 106:147A–147A, 2005.
- [63] D. H. Y. Leung and Y. G. Wang. Isotonic designs for phase I trials. *Controlled Clinical Trials*, 22:126–138, 2001.
- [64] Y. Lin and W. J. Shih. Statistical properties of the traditional algorithm-based designs for phase I cancer trials. *Biostatistics*, 2:203–215, 2001.
- [65] S. M. Lippman, S. E. Benner, and W. K. Hong. Cancer chemoprevention. *Journal of Clinical Oncology*, 12:851–873, 1994.
- [66] P. McCullagh and J. A. Nelder. *Generalized Linear Models*. New York: Chapman & Hall/CRC, 1989.
- [67] S. Møller. An extension of the continual reassessment methods using a preliminary up-and-down design in dose finding study in cancer patients, in order to investigate a greater range of doses. *Statistics in Medicine*, 14:911–922, 1995.
- [68] B. J. T. Morgan. *Analysis of Quantal Response Data*. New York: Chapman & Hall, 1992.
- [69] J. H. Muler, C. J. McGinn, D. Normolle, T. Lawrence, D. Brown, G. Hejna, and M. M. Zalupski. Phase I trial using a time-to-event continual reassessment strategy for dose escalation of cisplatin combined with gemcitabine and radiation therapy in pancreatic cancer. *Journal of Clinical Oncology*, 22:238–243, 2004.
- [70] J. R. Murphy and D. L. Hall. A logistic dose-ranging method for phase I clinical investigations trials. *Journal of Biopharmaceutical Statistics*, 7(4):635–647, 1997.
- [71] National Cancer Institute. Common terminology criteria for adverse events. Version 3.0. <http://ctep.cancer.gov/forms/CTCAEv3.pdf>.
- [72] J. Naylor and A. Smith. Applications of a method for the efficient computation of posterior distributions. *Applied Statistics*, 31:214–225, 1982.
- [73] D. Normolle and T. Lawrence. Designing dose-escalation trials with late-onset toxicities using the time-to-event continual reassessment method. *Journal of Clinical Oncology*, 24:4426–4433, 2006.
- [74] J. O’Quigley. Theoretical study of the continual reassessment method. *Journal of Statistical Planning and Inference*, 136:1765–1780, 2006.
- [75] J. O’Quigley and S. Chevret. Methods for dose finding studies in cancer clinical trials: a review and results of a Monte Carlo study. *Statistics in Medicine*, 19:1647–1664, 1991.
- [76] J. O’Quigley, M. D. Hughes, and T. Fenton. Dose finding designs for HIV studies. *Biometrics*, 57:1018–1029, 2001.
- [77] J. O’Quigley, X. Paoletti, and J. MacCario. Non-parametric optimal design in

- dose finding studies. *Biostatistics*, 3:51–6, 2002.
- [78] J. O’Quigley, M. Pepe, and L. Fisher. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics*, 46:33–48, 1990.
- [79] J. O’Quigley and E. Reiner. A stopping rule for the continual reassessment method. *Biometrika*, 85:741–748, 1998.
- [80] J. O’Quigley and L. Z. Shen. Continual reassessment method: a likelihood approach. *Biometrics*, 52:673–684, 1996.
- [81] S. Piantadosi, J. D. Fisher, and S. A. Grossman. Practical implementation of the continual reassessment method for dose finding trials. *Cancer Chemotherapy and Pharmacology*, 41:429–436, 1998.
- [82] M.-Y. Polley and Y. K. Cheung. Two-stage designs for dose finding trials with a biologic endpoint using stepwise tests. *Biometrics*, 64:232–241, 2008.
- [83] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.r-project.org>, 2008.
- [84] M. J. Ratain, D. Collyar, B. Kamen, E. Eisenhauer, T. S. Lawrence, C. Runowicz, S. Turner, and J. L. Wade. Critical role of phase I clinical trials in cancer treatment. *Journal of Clinical Oncology*, 15:853–859, 1997.
- [85] M. J. Ratain, R. Mick, R. L. Schilsky, and M. Siegler. Statistical and ethical issues in the design and conduct of phase I and II clinical trials of new anticancer agents. *Journal of the National Cancer Institute*, 85:1637–1643, 1993.
- [86] H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22:400–407, 1951.
- [87] H. Robbins and D. Siegmund. A convergence theorem for non-negative almost supermartingales and some applications. In Jagdish S. Rustagi, editor, *Optimizing Methods in Statistics*, pages 233–257. Academic Press, New York, 1971.
- [88] A. Rogatko, D. Schoeneck, W. Jonas, M. Tighiouart, F. R. Khuri, and A. Porter. Translation of innovative designs into phase I trials. *Journal of Clinical Oncology*, 25:4982–4986, 2007.
- [89] D. D. Rosa, J. Harris, and G. C. Jayson. The best guess approach to phase I trial design. *Journal of Clinical Oncology*, 24:206–208, 2006.
- [90] J. Ruan, P. Martin, R. R. Furman, S. M. Lee, K. Cheung, J. M. Vose, A. Lacasce, J. Morrison, R. Elstrom, S. Ely, A. Chadburn, E. Cesarman, M. Coleman, and J. P. Leonard. Bortezomib plus CHOP-Rituximab for previously untreated diffuse large B cell lymphoma (DLBCL) and mantle cell lymphoma (MCL). *Journal of Clinical Oncology*, 2010 under review.
- [91] D. B. Rubin. Inference and missing data. *Biometrika*, 63:581–592, 1976.
- [92] J. Sacks. Asymptotic distribution of stochastic approximation procedures. *Annals of Mathematical Statistics*, 29:373–405, 1958.

- [93] M. A. Schneiderman. How can we find an optimal dose? *Toxicology and Applied Pharmacology*, 7:44–53, 1965.
- [94] L. Z. Shen and J. O’Quigley. Consistency of continual reassessment method under model misspecification. *Biometrika*, 83:395–405, 1996.
- [95] L. Z. Shen and J. O’Quigley. Using a one-parameter model to sequentially estimate the root of a regression function. *Computational Statistics and Data Analysis*, 34:357–369, 2000.
- [96] M. J. Silvapulle. On the existence of maximum likelihood estimators for the binomial response model. *Journal of the Royal Statistical Society Series B*, 43:310–313, 1981.
- [97] B. Storer and D. DeMets. Current phase I/II designs: Are they adequate? *Journal of Clinical Research and Drug Development*, 1:121–130, 1987.
- [98] B. E. Storer. Design and analysis of phase I clinical trials. *Biometrics*, 45:925–937, 1989.
- [99] M. Stylianou and N. Flournoy. Dose finding using the biased coin up-and-down design and isotonic regression. *Biometrics*, 58:171–177, 2002.
- [100] A. C. Tamhane, Y. Hochberg, and C. W. Dunnett. Multiple test procedures for dose finding. *Biometrics*, 52:21–37, 1996.
- [101] J. Tang and A. K. Gupta. On the distribution of the produce of independent beta variables. *Statistics and Probability Letters*, 2:165–168, 1984.
- [102] P. F. Thall and J. D. Cook. Dose finding based on efficacy–toxicity trade-offs. *Biometrics*, 60:684–693, 2004.
- [103] P. F. Thall, J. J. Lee, C. H. Tseng, and E. H. Estey. Accrual strategies for phase I trials with delayed patient outcome. *Statistics in Medicine*, 18:1155–1169, 1999.
- [104] P. F. Thall, H. Nguyen, and E. H. Estey. Patient-specific dose finding based on bivariate outcomes and covariates. *Biometrics*, 64:1126–1136, 2008.
- [105] N. Ting, editor. *Dose Finding in Drug Development (Statistics for Biology and Health)*. Berlin: Springer, 2006.
- [106] J. M. Treluyer, S. Zohar, E. Rey, P. Hubert, F. Iserin, M. Jugie, R. Lenclen, S. Chevret, and G. Pons. A strategy for dose finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics*, 54:251–264, 1998.
- [107] J. H. Venter. An extension of the Robbins-Monro procedure. *Annals of Mathematical Statistics*, 38:181–190, 1967.
- [108] A. Wald. Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics*, 16:117–186, 1945.
- [109] A. Wald and J. Wolfowitz. Optimum character of the sequential probability ratio test. *Annals of Mathematical Statistics*, 19:326–339, 1948.
- [110] G. B. Wetherill. Sequential estimation of quantal response curves. *Journal of the Royal Statistical Society, Series B*, 25:1–48, 1963.

- [111] J. Whitehead and H. Brunier. Bayesian decision procedures for dose determining experiments. *Statistics in Medicine*, 14:885–893, 1995.
- [112] C. F. J. Wu. Efficient sequential designs with binary data. *Journal of the American Statistical Association*, 80:974–84, 1985.
- [113] C. F. J. Wu. Maximum likelihood recursion and stochastic approximation in sequential designs. In Van Ryzin, editor, *Adaptive Statistical Procedures and Related Topics*, volume 8, pages 298–313. Institute of Mathematical Statistics, 1986.
- [114] X. Wu. Stepwise procedures for dose finding in an adaptive clinical trial of early rehabilitation after acute stroke. In *DrPH dissertation*. Columbia University, 2010.
- [115] G. Yin, Y. Li, and Y. Ji. Bayesian dose finding in phase I/II clinical trials using toxicity and efficacy odds ratio. *Biometrics*, 62:777–787, 2006.
- [116] Z. Ying and C. F. J. Wu. An asymptotic theory of sequential designs based on maximum likelihood recursion. *Statistica Sinica*, 7:75–91, 1997.
- [117] S. Zacks, A. Rogatko, and J. Babb. Optimal bayesian-feasible dose escalation for cancer phase I trials. *Statistics and Probability Letters*, 38:215–220, 1998.