

Springer Series in Statistics

# Sequential Experimentation in Clinical Trials

Design and Analysis

 Springer

# Springer Series in Statistics

*Series Editors:*

P. Bickel

P.J. Diggle

S.E. Fienberg

U. Gather

I. Olkin

S. Zeger

For further volumes:

<http://www.springer.com/series/692>



Jay Bartroff • Tze Leung Lai • Mei-Chiung Shih

# Sequential Experimentation in Clinical Trials

Design and Analysis

 Springer

Jay Bartroff  
Department of Mathematics  
University of Southern California  
Los Angeles, California, USA

Tze Leung Lai  
Department of Statistics  
and Cancer Institute  
Stanford University  
Stanford, California, USA

Mei-Chiung Shih  
Department of Health Research and Policy  
Stanford University  
VA Cooperative Studies Program  
Stanford, California, USA

ISSN 0172-7397

ISBN 978-1-4614-6113-5

ISBN 978-1-4614-6114-2 (eBook)

DOI 10.1007/978-1-4614-6114-2

Springer New York Heidelberg Dordrecht London

Library of Congress Control Number: 2012951627

© Springer Science+Business Media New York 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media ([www.springer.com](http://www.springer.com))

*To our parents*

*Jack and Barbara Bartroff*

*Chi Y. and Wai C. Lai*

*Ming-Tsan Shih and Min-Hui Hsu*



# Preface

The idea of writing a book on sequential experimentation in clinical trials arose 25 years ago when Lai was at Columbia University, collaborating with Dan Anbar of Abbott Laboratories on a university-industry cooperative research project, “Sequential Statistical Methods in Biopharmaceutical Research,” funded by the National Science Foundation. Anbar and Lai, together with Gordon K.K. Lan and Anastasio Tsiatis (at that time, at the NIH and Harvard, respectively), formed a focused research group that held a week-long meeting every 2 months and organized an annual workshop, with invited speakers from academia, industry, and the FDA and NIH, and dedicated to the development and discussion of sequential methods in the design and analysis of clinical trials. Although substantial progress was made by the group to advance this new area that attracted considerable attention from the pharmaceutical industry after the early termination of the Beta-Blocker Heart Attack Trial in 1981, the book project could not materialize when Lai moved to the West Coast in 1987, joining Stanford University, while the other collaborators remained on the East Coast but were busy with their own moves to new positions. On the other hand, the group members continued their separate research efforts in this area. These efforts and those by other researchers led to major advances and eventual widespread use of group sequential designs and interim analysis methods by the pharmaceutical industry and their acceptance by the FDA.

At the turn of the new century, the monograph by [Jennison and Turnbull \(2000\)](#) appeared, giving a comprehensive overview of group sequential methods developed up to that time. Besides continual developments in interim analysis and group sequential methods, the past decade has also witnessed new developments and growing interest in adaptive designs of clinical trials. The books by [Proschan et al. \(2006\)](#), [Chow and Chang \(2006\)](#), [Chang \(2007\)](#), and [Berry et al. \(2011\)](#) describe some of these developments and their applications. However, as pointed out in [Chap. 8](#), there is substantial disagreement in the literature concerning the appropriateness of these adaptive designs, which either use inefficient test statistics that are not supported by mainstream statistical principles to adjust for the adaptation in maintaining the Type I error of the test or use Bayesian posterior probabilities that do not guarantee the prescribed Type I error. [Chapter 8](#) describes our recent work that provides a new



class of adaptive designs which are both flexible and efficient, thereby resolving the dilemma between efficiency and flexibility in the adaptive design literature. Prior to this work, we have also developed a comprehensive methodology of flexible and efficient group sequential designs, to which Chap. 4 is devoted. In fact, the new adaptive designs in Chap. 8 are modifications of the corresponding group sequential designs in Chap. 4, and a unified approach is provided for the methodology and implementation of group sequential and adaptive designs.

Besides giving an up-to-date account of these flexible designs, we also present in Chap. 7 a comprehensive overview, including the most recent developments of inference after the termination of these clinical trials. Chapter 6 describes the Beta-Blocker Heart Attack Trial as an example for the design and analysis of clinical trials with failure-time endpoints and interim analyses. The material in Chaps. 4, 6, 7, and 8 can be used for short courses on group sequential and adaptive designs. We have given short courses based on this material in the First Joint Biostatistics Symposium in Beijing, July 2010, the Applied Statistics Symposium of the International Chinese Statistical Association in New York, June 2011, and the Workshop on the Design and Analysis of Clinical Trials at National University of Singapore, October 2011. We were greatly encouraged by the enthusiastic response and stimulating comments of the participants.

This book has also benefited from the Third International Workshop in Sequential Methodologies at Stanford University, June 2011. The workshop was very well attended and was truly international in nature. There it was pointed out that despite a resurgence of interest in sequential analysis, the subject was not in the graduate curriculum of most statistics departments. One reason that was mentioned was the lack of textbooks that could present the material in an appealing way to today's graduate students. In fact, only a handful of such books had been written and they were published more than 20 years ago. Although there are more recent books which we have mentioned in the second paragraph of this preface, they all deal with the specialized topics of group sequential and adaptive designs rather than general methods and principles in sequential analysis. Another reason that came up during workshop discussions was that sequential methods and adaptive designs seemed to involve special techniques and ideas that are detached from mainstream topics taught in the modern graduate statistics curriculum, e.g., likelihood inference, regression analysis, resampling, semiparametric theory, to name a few. Spurred by these comments, we have made particular efforts to change this perception in the selection and presentation of the materials. To make it suitable for an introductory course on sequential analysis, the book covers the much broader subject of sequential experimentation that includes group sequential and adaptive designs of Phase II and III clinical trials, which have attracted much attention in the past three decades. In particular, the broad scope of design and analysis problems in sequential experimentation clearly requires a wide range of statistical methods and models from nonlinear regression analysis, experimental design, dynamic programming, survival analysis, resampling, and likelihood and Bayesian inference. The background material in these building blocks is summarized in Chaps. 2 and 3 and certain sections in Chaps. 6 and 7. Besides group sequential tests and adaptive

designs, we also introduce sequential change-point detection methods in Chap. 5 in connection with pharmacovigilance and public health surveillance. Together with dynamic programming and approximate dynamic programming in Chap. 3, the book therefore covers all basic topics for a graduate course in sequential analysis.

Different parts of the book can be used for short courses on clinical trials, translational medical research, and sequential experimentation. Lai has used an early draft of the book to teach a course on innovative clinical trial designs and statistical methods for second-year Ph.D. students in the Department of Statistics at Stanford University. The course has led to supplements and exercises for various chapters and also to the web site for the book, <http://meichiun.web.stanford.edu/clinicaltrials/>, to which different parts of the book refer for links to software.

We thank the teaching assistant Alex Deng and the students in the aforementioned course for their stimulating interest and helpful comments on a first draft of the book that was used as lecture notes. In particular, we want to mention Olivia Liao who subsequently wrote her Ph.D. thesis on the subject and contributed to some recent publications that we have referenced. We also thank our collaborators Philip Lavori, Zheng Su, Ray Zhu, Joe Heyse, Jie Chen, Sam Wong, and Hock Peng Chan for working with us on related projects and for their valuable suggestions. We are particularly grateful to Balasubramanian Narasimhan who has been helping us develop open-source software for the innovative designs and analyses that will be posted at the book's web site and to Cindy Kirby who has put together the separate chapters in an efficient and timely fashion. We also acknowledge grant support for the research projects related to this book by the National Science Foundation (DMS-0907241 and 0403105 for Bartroff and 1106535 and 0805879 for Lai), the National Institutes of Health (GMS-068968 for Bartroff and 1 P30 CA124435 for Lai), the U.S. Department of Veterans Affairs Cooperative Studies Program (for Shih), and the Stanford Center for Clinical and Translational Education and Research (for Lai and Shih).

Los Angeles, California  
Stanford, California  
Stanford, California

Jay Bartroff  
Tze Leung Lai  
Mei-Chiung Shih



# Contents

<b>1</b>	<b>Introduction</b>	1
1.1	Sequential Experimentation in Translational Medical Research	1
1.2	Sequential Analysis: From Weapons Testing to Confirmatory Clinical Trials	2
1.3	Adaptation and Sequential Optimization	5
1.4	Two Time Scales and Time-Sequential Survival Analysis	6
1.5	Bayesian and Frequentist Approaches and Associated Software	7
<b>2</b>	<b>Nonlinear Regression, Experimental Design, and Phase I Clinical Trials</b>	11
2.1	Nonlinear Regression Models	12
2.1.1	Nonlinear Least Squares	12
2.1.2	Nonlinear Mixed Effects Models	13
2.2	Pharmacokinetics and Pharmacodynamics	16
2.3	Theory of Optimal Design	18
2.3.1	Optimal Design Theory in Linear Regression Models	19
2.3.2	Elfving's Method for $c$ -Optimal Design	20
2.3.3	Extension to Nonlinear Models	22
2.4	Phase I Clinical Trials for Relatively Benign Drugs	23
2.5	Early Phase Clinical Trials for Cytotoxic Cancer Treatments	24
2.5.1	Up-and-Down and Related Designs	24
2.5.2	Model-Based Designs	28
2.6	Supplements and Problems	29
<b>3</b>	<b>Sequential Testing Theory and Stochastic Optimization Over Time</b>	37
3.1	Likelihood Ratio Statistics and Likelihood Ratio Identity	37
3.2	Wald's SPRT and Its Error Probabilities	38
3.3	Wald's Equation and Lower Bounds of Wald and Hoeffding	39
3.4	Lorden's 2-SPRT and Sequential GLR Tests	41
3.5	Modifications for Group Sequential Testing	43

- 3.6 Optimality of SPRT and Dynamic Programming..... 45
  - 3.6.1 Dynamic Programming: Finite-Horizon Case ..... 45
  - 3.6.2 Infinite-Horizon Dynamic Programming ..... 46
  - 3.6.3 Bayes Sequential Tests of Simple Null Versus Simple Alternative ..... 47
  - 3.6.4 Auxiliary Bayes Problem and Optimality of the SPRT ..... 50
- 3.7 Approximations to Bayes Sequential Tests and Operating Characteristics ..... 51
  - 3.7.1 Derivation of Lorden’s 2-SPRT as an Approximate Bayes Rule ..... 51
  - 3.7.2 Approximation of Bayes Sequential Tests of One-Sided Hypotheses in an Exponential Family by Sequential GLR Tests ..... 52
  - 3.7.3 Laplace’s Asymptotic Formula and Approximations to Operating Characteristics ..... 56
- 3.8 ADP and Applications to Phase I Cancer Trial Designs ..... 60
  - 3.8.1 A Stochastic Optimization Problem Related to the Treatment Versus Experimentation Dilemma ..... 60
  - 3.8.2 Some ADP Methods ..... 63
  - 3.8.3 Hybrid Designs Derived by ADP ..... 66
  - 3.8.4 A Two-Stage Modification ..... 70
- 3.9 Supplements and Problems ..... 73
- 4 Group Sequential Design of Phase II and III Trials..... 77**
  - 4.1 Group Sequential Tests for a Normal Mean ..... 78
    - 4.1.1 The Pocock, O’Brien–Fleming, and Wang–Tsiatis Group Sequential Boundaries ..... 78
    - 4.1.2 One-Sided Group Sequential Tests: Power Family and Triangular Tests ..... 78
    - 4.1.3 The Lan–DeMets Error Spending Approach ..... 80
  - 4.2 Group Sequential Generalized Likelihood Ratio Tests with Modified Haybittle–Peto Boundaries ..... 80
    - 4.2.1 Maximum Sample Size and the Alternative Implied by It ..... 81
    - 4.2.2 Group Sequential GLR Tests with Nearly Optimal Power and Expected Sample Size ..... 81
    - 4.2.3 Two-Sided Tests with or Without Futility Boundaries ..... 84
    - 4.2.4 Extensions to Multiparameter and Multiarmed Problems .... 85
  - 4.3 Implementation and Computational Methods ..... 86
    - 4.3.1 Asymptotic Normality and Recursive Numerical Integration ..... 86
    - 4.3.2 Monte Carlo and the Bootstrap for Sequential GLR Statistics ..... 87
    - 4.3.3 Design, Interim Analysis, and Nonparametric Extensions.... 87

- 4.4 Comparison of Group Sequential Tests ..... 90
  - 4.4.1 One-Sided Tests ..... 90
  - 4.4.2 Two-Sided Tests Without Futility Boundaries ..... 91
  - 4.4.3 Unequal Group Sizes That Are Not Prespecified ..... 93
- 4.5 Supplements and Problems ..... 95
- 5 Sequential Methods for Vaccine Safety Evaluation and Surveillance in Public Health** ..... 101
  - 5.1 Vaccine Safety Evaluation ..... 101
    - 5.1.1 Design Considerations for Clinical Trials to Test Vaccine Safety ..... 102
    - 5.1.2 Application of SPRT and Other Sequential Tests ..... 103
  - 5.2 Sequential GLR Tests in Prelicensure Randomized Clinical Trials ..... 106
    - 5.2.1 An Exponential Family Representation of Sequential GLR Tests ..... 106
    - 5.2.2 Implementation and Example ..... 107
    - 5.2.3 Discussion ..... 109
  - 5.3 The Rotavirus Efficacy and Safety Trial (REST) ..... 109
  - 5.4 Post-marketing Surveillance of Drug Safety ..... 111
  - 5.5 Sequential Change-Point Detection Methods for Pharmacovigilance and Public Health Surveillance ..... 113
    - 5.5.1 Sequential Change-Point Detection Methodology: An Overview ..... 113
    - 5.5.2 Sequential Detection in Public Health Surveillance ..... 119
    - 5.5.3 A Hybrid of Sequential Testing and Detection for Pharmacovigilance ..... 119
  - 5.6 Supplements and Problems ..... 120
- 6 Time-Sequential Design of Clinical Trials with Failure-Time Endpoints** ..... 123
  - 6.1 An Overview of Survival Analysis ..... 123
    - 6.1.1 Nelson–Aalen and Kaplan–Meier Estimators ..... 124
    - 6.1.2 Regression Models for Hazard Functions with Covariates ... 126
    - 6.1.3 Rank Tests Based on Censored Survival Data ..... 128
  - 6.2 The Beta-Blocker Heart Attack Trial (BHAT) ..... 128
    - 6.2.1 Trial Design ..... 128
    - 6.2.2 Trial Execution and Interim Analysis by DSMB ..... 129
    - 6.2.3 Stochastic Curtailment ..... 130
  - 6.3 Terminal Analyses of BHAT Data ..... 131
  - 6.4 Developments in Group Sequential Methods Motivated by BHAT ..... 133
    - 6.4.1 The Lan–DeMets Error-Spending Approach ..... 133
    - 6.4.2 Two Time Scales and Haybittle-Type Boundaries ..... 133
    - 6.4.3 Conditional and Predictive Power ..... 135

6.5	Randomized Clinical Trials with Failure-Time Endpoints and Interim Analyses .....	136
6.5.1	Time-Sequential Censored Rank Statistics and Their Asymptotic Distributions .....	136
6.5.2	Modified Haybittle–Peto Tests, Power, and Expected Savings .....	139
6.6	Appendix: Martingale Theory and Applications to Sequential/Survival Analysis .....	142
6.6.1	Optional Stopping Theorem .....	142
6.6.2	Predictable Variation and Stochastic Integrals .....	144
6.6.3	Rebolledo’s CLT .....	145
6.6.4	Counting Processes and Applications to Survival Analysis ..	146
6.6.5	Application to Time-Sequential Censored Rank Statistics ...	147
6.7	Supplements and Problems .....	148
<b>7</b>	<b>Confidence Intervals and <math>p</math>-Values</b> .....	<b>153</b>
7.1	Inference Following a Sequential Trial: Some Developments .....	154
7.1.1	Naive Confidence Intervals, Anscombe’s Theorem, and Bayesian Methods .....	154
7.1.2	Exact Confidence Intervals in the Normal Case .....	155
7.1.3	Siegmund’s Ordering and Group Sequential Trials with Random Group Sizes .....	156
7.1.4	Bias Correction for A Modified Pivot .....	157
7.1.5	Bivariate Normal Outcomes .....	157
7.1.6	Extensions Beyond the Normal Cases .....	159
7.2	A Hybrid Resampling Approach .....	159
7.2.1	A General Formulation of Exact, Bootstrap, and Hybrid Resampling Methods .....	159
7.2.2	Hybrid Resampling Confidence Intervals for Population Means Following Group Sequential Tests Based on Sample Means .....	161
7.2.3	A Comparative Study .....	162
7.3	A General Ordering Scheme and $p$ -Values .....	163
7.3.1	Total Ordering of Sample Space for $p$ -Values .....	163
7.3.2	Bootstrap Methods for $p$ -Values and Hybrid Resampling ...	165
7.4	Hybrid Resampling Approach for Secondary Endpoints .....	166
7.4.1	Bivariate Mean Vectors .....	166
7.4.2	Bivariate Nonparametric Functionals .....	167
7.5	Applications to Time-Sequential Survival Trials .....	169
7.5.1	Implementation of the Hybrid Resampling Approach .....	170
7.5.2	Tests on Secondary Parameters in the Cox Model .....	171
7.6	Simulation Studies .....	174
7.6.1	Confidence Intervals for Primary Endpoint of Hazard Rate .....	174
7.6.2	Test for Secondary Endpoint in Cox Model .....	177
7.7	Supplements and Problems .....	178

**8 Adaptive Design of Confirmatory Trials** ..... 187

8.1 Internal Pilot, Midcourse Sample Size Re-estimation  
and Trial Extensions ..... 189

8.1.1 Stein’s Two-Stage Procedure, with the First Stage  
to Estimate the Variance as an Internal Pilot ..... 189

8.1.2 Midcourse Sample Size Re-estimation ..... 190

8.1.3 Midcourse Modification of the Maximum Sample  
Size in a Group Sequential Trial ..... 192

8.1.4 Optimal Adaptive Group Sequential Designs  
via Dynamic Programming ..... 193

8.2 Efficient Adaptive Design and GLR Tests ..... 193

8.2.1 An Adaptive 3-Stage GLR Test ..... 193

8.2.2 Midcourse Modification of Maximum Sample Size ..... 195

8.2.3 Asymptotic Theory ..... 196

8.2.4 Implementation via Normal Approximation  
or Monte Carlo ..... 198

8.3 Comparison of Adaptive Designs ..... 200

8.3.1 Normal Mean with Known Variance ..... 200

8.3.2 Difference of Means with Unknown Variances ..... 205

8.3.3 Comparison of Tests Allowing Midcourse  
Modification of Maximum Sample Size ..... 207

8.3.4 Coronary Intervention Study ..... 210

8.4 Adaptive Choice Between Superiority and Non-inferiority  
Objectives via a Flexible Group Sequential Design ..... 213

8.4.1 Adaptation and Group Sequential Tests of  $H_0$   
or  $H'_0$  When  $M' \leq M$  ..... 214

8.4.2 Binary Responses and a Comparative Study ..... 216

8.4.3 Adaptive Choice Between One Superiority  
and Two Non-inferiority Trials ..... 219

8.4.4 Discussion ..... 222

8.5 Supplements and Problems ..... 222

**References** ..... 225

**Index** ..... 235





# Chapter 1

## Introduction

This chapter gives an overview of (a) the prevalence of sequential experimentation in translational medical research and (b) developments of statistical methods to design and analyze these sequential experiments in evidence-based medical research. In this connection it also gives an outline of the topics covered in the subsequent chapters and discusses the complementary roles of Bayesian and frequentist approaches to sequential design and analysis.

### 1.1 Sequential Experimentation in Translational Medical Research

“From bench to bedside,” a maxim of translational medical research, reflects the sequential nature of the experiments involved. “Bench” refers to laboratory experiments to study new biochemical principles and discover novel treatments. The experiments with promising results are followed by preclinical animal studies. After understanding the effect of the treatment (say, a new drug) on animals (e.g., rodents), the next stage of drug development consists of clinical trials that involve human subjects, starting with Phase I studies to determine a safe dose or dosage regimen and to collect information on the pharmacokinetics (PK) and pharmacodynamics (PD) of the drug. PK is concerned with the concentration versus time curve that is associated with the kinetics of drug absorption, distribution, and elimination. PD is concerned with the steady-state relationship of drug concentration at an effector site to the effect/response produced. The information collected and the dosage regimen determined from Phase I studies are used to design Phase II clinical trials to evaluate the efficacy of the drug for particular indications (endpoints) in patients with the disease. Phase II trials are precursors of Phase III trials whose goal is to demonstrate effectiveness of the drug for its approval by the regulatory agency (the Food and Drug Administration in the United States) and to provide adequate evidence for its labeling. Besides testing efficacy, Phase III trials also collect safety information

from the relatively large samples of patients accrued to the trial. The safety of the drug is evaluated from the data obtained from all three phases of clinical trials prior to marketing approval of the drug and continues to be evaluated through post-marketing Phase IV trials.

Despite the sequential nature of Phase I–III trials, the trials are often planned separately, treating each trial as an independent study whose design depends on results from studies in previous phases. An advantage of this is that the reproducibility of the results of the trial can be evaluated on the basis of the prescribed design, without worrying about the statistical variability of the results of earlier-phase trials that determine the prescribed design. A disadvantage lies in the fact that the sample sizes of the trials are often inadequate because of the separate planning. A different strategy is to expand a trial seamlessly from one phase into the next phase; the Phase II–III cancer trial design in Sect. 6.7 is an example. Although Phase II–III design, which is an active area of current research undergoing new advances, is beyond the scope of this book, we give a brief introduction in Sect. 6.7 to show the power of an overarching sequential experimentation approach to translational medicine.

This book focuses on sequential methods for the design and analysis of Phase I, II, and III clinical trials, thereby providing the background for understanding and developing the new advances. Although these methods are developed in the context of clinical trials, they are also applicable to other fields that involve sequential experimentation. We therefore give an introduction to the statistical methods and the underlying principles and also relate them to basic topics taught in typical graduate statistics programs that assume the data to be generated by nonsequential designs. For example, while Chap. 2 considers Phase I clinical trials, it starts with nonlinear regression and experimental design before relating them to basic pharmacologic principles and models underlying dose determination.

## 1.2 Sequential Analysis: From Weapons Testing to Confirmatory Clinical Trials

The subject named *sequential analysis*, which also includes sequential design of experiments, was born in response to demands for more efficient testing of antiaircraft gunnery during World War II, which led to Wald’s development of the sequential probability ratio test (SPRT) in 1943 (Wallis 1980). Let  $X_1, X_2, \dots$  be i.i.d. random variables with common density function  $f$ . To test  $H_0 : f = f_0$  versus  $H_1 : f = f_1$ , the SPRT stops sampling at stage

$$N = \inf \left\{ n \geq 1 : \prod_{i=1}^n (f_1(X_i)/f_0(X_i)) \notin (A, B) \right\}, \quad (1.1)$$

where  $0 < A < 1 < B$  are the stopping boundaries. When stopping occurs,  $H_0$  or  $H_1$  is rejected according to whether the likelihood ratio  $\prod_{i=1}^N (f_1(X_i)/f_0(X_i))$  crosses

the upper boundary  $B$  or the lower boundary  $A$ . In Chap. 3 we give a summary of the theory of sequential tests of hypotheses, beginning with the SPRT on testing a simple null versus a simple alternative hypothesis and describing important subsequent developments that led to a relatively complete theory for composite hypotheses.

Within a few years after Wald's introduction of the SPRT, it was recognized that sequential hypothesis testing might provide a useful tool in clinical trials to test the efficacy of new medical treatments. A number of papers appeared during the 1950s on modifications of the SPRT for the design of clinical trials, and an overview of these developments was given in Armitage (1960). In 1969, Armitage et al. proposed a new alternative to the SPRT and its variants, called the *repeated significance test* (RST). The underlying motivation for the RST is that, since the strength of evidence in favor of a treatment from a clinical trial is conveniently indicated by the results of a conventional significance test, it is appealing to apply the significance test, with nominal significance level  $\alpha$ , repeatedly during the trial. Noting that the overall significance level  $\alpha^*$ , which is the probability that the nominal significance level is attained at some stage, is larger than  $\alpha$ , they developed a recursive numerical algorithm to compute  $\alpha^*$  in the case of testing a normal mean  $\theta$  with known variance  $\sigma^2$ , for which the RST of  $H_0 : \theta = 0$  is of the form

$$T = \inf \{n \leq M : |S_n| \geq a\sigma\sqrt{n}\}, \quad (1.2)$$

rejecting  $H_0$  if  $T < M$  or if  $T = M$  and  $|S_M| \geq a\sigma\sqrt{M}$ , where  $S_n = X_1 + \cdots + X_n$ . Haybittle (1971) proposed the following modification of the RST to increase its power. The stopping rule has the same form as (1.2) but the rejection region is modified to  $T < M$  or  $|S_M| \geq c\sigma\sqrt{M}$ , where  $a(\geq c)$  is so chosen that the overall significance level is equal to some prescribed number. In particular,  $a = \infty$  gives the fixed sample size test while  $a = c$  gives the RST.

In double-blind multicenter clinical trials, it is not feasible to arrange for continuous examination of the data as they accumulate to perform the RST. This led Pocock (1977) to introduce a "group sequential" version of (1.2), in which the  $X_n$  represents an approximately normally distributed statistic of the data in the  $n$ th group (instead of the  $n$ th observation) and  $M$  represents the maximum number of groups. Instead of the square-root boundary  $a\sigma\sqrt{n}$ , O'Brien and Fleming (1979) proposed to use a constant stopping boundary in

$$T = \inf \{n \leq M : |S_n| \geq b\}, \quad (1.3)$$

which corresponds to the group sequential version of an SPRT.

While sequential analysis had an immediate impact on weapons testing when it was introduced during World War II to reduce the sample sizes of such tests (Wallis 1980), its refinements for testing new drugs and treatments received little attention from the biomedical community until the Beta-Blocker Heart Attack Trial (BHAT) that was terminated in October 1981, prior to its prescheduled end in June 1982.

The main reason for this lack of interest is that the fixed sample size (i.e., the number of patients accrued) for a typical trial is too small to allow further reduction while still maintaining reasonable power at the alternatives of interest. On the other hand, BHAT, which was a multicenter, double-blind, randomized placebo-controlled trial to test the efficacy of long-term therapy with propranolol given to survivors of an acute myocardial infarction, drew immediate attention to the benefits of sequential methods not because it reduced the number of patients but because it shortened a 4-year study by 8 months, with positive results for a long-awaited treatment for MI patients.

The “success story” of BHAT paved the way for major advances in the development of group sequential methods in clinical trials and for the steadily increasing adoption of group sequential design. Chapter 4 gives a review of these advances and describes the current methodology that has moved far beyond the Pocock and O’Brien–Fleming boundaries (1.2) and (1.3). Chapter 6 presents the design details of BHAT and the interim analysis results considered by its Data and Safety Monitoring Board. Inspired by the statistical issues raised by BHAT, a number of important and difficult problems concerning the design and analysis of clinical trials with failure-time endpoints and interim analyses have been resolved in the past 2 decades, and Chap. 6 also describes the “time-sequential” methodology developed in this connection. Chapter 5, however, shows that the fully sequential methodology summarized in Chap. 3 has recently emerged as a standard for prelicensure (Phase III) vaccine safety trials and post-marketing (Phase IV) safety studies.

Analysis of the data at the conclusion of a clinical trial typically involves tests and confidence intervals not only for the primary endpoint but also for different secondary endpoints. The use of a stopping rule whose distribution depends on these parameters introduces substantial difficulties for such inference. Siegmund (1978) developed a method, based on ordering the sample space in a certain way, to construct confidence intervals for its mean of a normal population with known variance following a RST. Alternative orderings of the sample space were subsequently introduced for group sequential tests by Rosner and Tsiatis (1988) and Emerson and Fleming (1990). By making use of resampling methods, Chuang and Lai (1998, 2000) developed a general resampling approach to constructing accurate confidence intervals following sequential tests. Subsequently, Lai and Li (2006) introduced a general ordering scheme that can be used in conjunction with resampling to completely solve the long-standing problem of constructing valid confidence intervals for the primary endpoint of a group sequential trial. Chapter 7 summarizes these developments and describes the methods. Analysis of secondary endpoints following a group sequential trial is also considered in Chap. 7, which reviews the bias-correction approach of Whitehead (1986), Liu et al. (2000), Whitehead et al. (2000), and Hall and Yakir (2003) and describes the hybrid resampling methods of Lai et al. (2009).

### 1.3 Adaptation and Sequential Optimization

After sequential analysis was introduced in response to more efficient testing of weapons during World War II, it was soon realized that sequential methods could be used to address statistical problems for which there are no solutions with fixed sample sizes. While [Dantzig \(1940\)](#) had shown that no fixed sample size test exists for the problem of testing the null hypothesis  $H_0 : \mu = \mu_0$ , with prescribed error probabilities  $\alpha$  and  $\beta$  at  $\mu_0$  and  $\mu_0 + \delta$ , for the mean  $\mu$  of a normal distribution whose variance  $\sigma^2$  is unknown, [Stein \(1945\)](#) showed that a two-stage procedure that uses the first stage to estimate  $\sigma^2$  and thereby to determine an appropriate second-stage sample size can have power independent of  $\sigma$ . Stein's two-stage design is the first example to show that one can use data during the course of an experiment to learn about the unknown parameters and thereby adapt the experimental design (which is the sample size in Stein's example) as the experiment progresses. It also paved the way for the next generation of adaptive designs in clinical trials in the 1990s that are described in [Chap. 8](#). These adaptive designs, however, are inefficient because they do not incorporate the uncertainties of the parameter estimates at the end of the first stage. [Chapter 8](#) also describes a new class of adaptive designs, introduced by [Bartroff and Lai \(2008a,b\)](#), which use an additional stage to accommodate the uncertainties in the first-stage estimates.

Adaptation via sequential learning of unknown parameters is also a central idea in the theory of nonlinear optimal experimental design. As shown in [Sect. 2.3](#), the optimal design measure involves the unknown parameter vector  $\theta$  in a nonlinear regression model. To circumvent these difficulties, [Fedorov \(1972\)](#) and others proposed that designs be constructed sequentially, using observations made to date to estimate  $\theta$  and choosing the next design point by replacing the unknown  $\theta$  in the optimal design by the current estimate. [Lai et al. \(2012c\)](#) have recently shown the advantages of adaptation in an integrated plan for developing a new drug, which is mentioned earlier in the second paragraph of [Sect. 1.1](#). In the development of a new drug, an important component of the effort and costs involves clinical trials to provide clinical data to support a beneficial claim of the drug and, in case such claim is not valid, to support the termination of the development. The clinical trials progress in steps and are labeled Phase I, II, and III trials, as we have already noted in [Sect. 1.1](#). A project team steers their operations in which intensity, cost, and duration increase with the phase; in particular, Phase III often involves over 3000 professionals, several years to reach completion, and over \$100 million in cost. In addition, there is a core team that makes decisions guided by a clinical development plan (CDP). The CDP maps out the clinical development pathway, beginning with first-in-man studies and ending with submission to the regulatory agency or termination of development. It defines the number and type of clinical studies and their objectives, determines the time sequence of the studies, some of which may be carried out in parallel, identifies key risk areas, and sets key decision points and go/no-go criteria. [Julious and Swank \(2005\)](#) have noted that statistical

methods for clinical trial design have focused primarily on “optimizing individual clinical trials” but are lacking “at a more global level in the optimization of clinical development plans.” In practice, however, it is often difficult to specify in advance the cost of each clinical trial in the sequence and the prior probabilities of a go or no-go decision to perform the optimization of CDPs “at a more global level.” [Lai et al. \(2012c\)](#) use ideas from adaptive design of clinical trials, in particular, seamless Phase II–III designs, to adapt a CDP to information acquired during the course of its execution.

Optimization is an important technique in formulating and computing statistical procedures, which can be regarded as statistical decision rules. When the decision rule consists of a sequence of actions, determination of the optimal rule involves dynamic programming. In [Chap. 3](#) we give an introduction to dynamic programming and use it to prove the optimality of the SPRT for simple hypotheses and to derive approximately optimal tests based on generalized likelihood ratio statistics for composite hypotheses. We also give an introduction to recent advances in approximate dynamic programming and apply it to address the treatment versus experimentation dilemma in Phase I cancer trial designs.

## 1.4 Two Time Scales and Time-Sequential Survival Analysis

As pointed out in [Sect. 1.2](#), the early termination of BHAT paved the way for major advances in the development of group sequential designs. These advances are summarized in [Chap. 4](#), but BHAT and other trials with failure-time endpoints require more subtle methods than those described in [Chap. 4](#). In [Chap. 6](#) we describe these methods that address two time scales in time-sequential survival analysis. “Time-sequential” means that interim analyses are conducted over calendar times, rather than on the time scale measured by the number of subjects at each interim analysis as in group sequential methods in [Chap. 4](#), for which the number of subjects is proportional to the variance (under the null hypothesis) of the test statistic. We begin [Chap. 6](#) with a review of traditional (nonsequential) survival analysis, focusing on how the variances of the commonly used test statistics can be derived with relative ease, despite the complexities due to right censoring, by making use of martingale theory. In the time-sequential setting, calendar time is one time scale, and the other time scale is “information time,” which is measured by the null variance of the test statistic at the time of interim analysis. There is no simple connection between the two time scales and it has been a long-standing problem concerning how to address the difficulties caused by the two time scales in the design and analysis of time-sequential clinical trials with failure-time endpoints.

In [Sect. 6.5](#) we discuss these difficulties and describe the methods that have been developed to address them. These include a comprehensive asymptotic distribution theory for time-sequential censored rank statistics, relatively simple and yet efficient modified Haybittle–Peto tests, and interim Bayesian estimation of the maximum information at the scheduled end of the trial for futility stopping. [Section 6.7](#)

describes some recent advances, including the Phase II–III cancer trial designs that we have mentioned in Sect. 1.2 and a method that allows multiple test statistics to increase power if the trial should proceed to its scheduled end. Sections 7.3 and 7.5 describe an innovative hybrid resampling approach to statistical inference from survival data following a time-sequential trial.

## 1.5 Bayesian and Frequentist Approaches and Associated Software

Berry et al. (2011, p. 1) say that a primary purpose of their book is to describe the Bayesian approach as an alternative to the traditional frequentist approach, which is “the standard statistical approach to designing and analyzing clinical trials and other medical experiments.” They find the “flexibility in both design and analysis” and the “decision-oriented” underpinning of the Bayesian approach particularly suited to sequential analysis and adaptive design of clinical trials. On the other hand, they acknowledge that for Phase III confirmatory trials, which are “typically overseen and judged by a regulatory agency,” the statistical hurdle for regulatory approval of the new treatment is “to get a statistically significant result at a specified type I error,” and the type I error of an adaptive Bayesian design is “extremely difficult, if not impossible, to calculate” and has to be computed by Monte Carlo simulations. Their approach is to adjust the rejection threshold of the Bayesian adaptive/sequential test by using the Monte Carlo simulations carried out under some chosen parameter configuration(s) belonging to the null hypothesis. However, for a composite null hypothesis, there is no guarantee that the worst parameter configuration in the null hypothesis has been chosen for these simulations. An example is given by Lai et al. (2012a, Sect. 4.4) in their comparison of the Phase II–III design described in Chap. 6 with the Bayesian counterpart developed by Huang et al. (2009). Their numerical study shows that because the Bayesian design uses simulations under certain assumed survival rates to control the type I error, the type I error can be substantially inflated under other survival rates belonging to the highly composite null hypothesis. In contrast, the frequentist semiparametric approach used in their design and analysis is shown to maintain the prescribed type I error.

The argument of Berry et al. (2011, Chap. 1) that the Bayesian approach can handle adaptation and sequential learning much more efficiently than the frequentist approach is fair for the prevailing frequentist methods cited in their references, but it overlooks the possibility that *suitably chosen* frequentist methods can work as well, if not better. In fact, there is already a versatile arsenal of statistical methods and theories, including likelihood inference, semiparametric models for censored survival data, bootstrap, and other resampling methods, for nonsequential settings. We shall show in the subsequent chapters how these time-tested methods can be extended to sequential experiments and adaptive designs. In fact, in Chap. 3, we show that these extensions can also be derived as approximations to Bayes



rules. Our viewpoint, therefore, is that Bayesian and frequentist approaches should complement each other. One may start with a Bayesian formulation and end up with a frequentist implementation that may be more convenient and appropriate for the problem at hand, for example, confirmatory testing for drug approval. This idea is illustrated in Sect. 3.7 that starts with Bayes sequential tests of one-sided hypotheses and ends up with sequential generalized likelihood ratio tests which have approximately optimal Bayesian and frequentist properties and are also convenient for implementation and description. Another example, which is beyond the scope of this book, is the classical multiarmed bandit problem; see the survey in [Lai \(2001, pp. 337–339\)](#) which shows that while the Bayesian formulation of the infinite-horizon version of the problem has a solution in terms of the “Gittins index” for each arm, a closed-form approximation of the Gittins index yields an upper confidence bound for the arm’s mean parameter. Not only does this frequentist approximation provide an intuitive interpretation of the Bayes solution but it also leads to approximately optimal solutions of finite-horizon bandit problems with a frequentist formulation. Conversely, one may start with a frequentist problem and ends up with a Bayes solution. A classic example is the optimality theorem of Wald’s SPRT in Sect. 3.6. As explained in Sect. 3.3, Wald conjectured this result on the basis of certain lower bounds for the expected sample sizes under the simple null and alternative hypotheses. Section 3.6 shows that the proof of the conjecture requires solving an auxiliary Bayes problem to which dynamic programming can be applied.

Other than Bayesian designs for Phase I trials, which usually have small sample sizes, considered in Chap. 2 and Sect. 3.8, and the interplay between Bayesian and frequentist approaches to sequential hypothesis testing discussed in Chap. 3, we focus in the subsequent chapters on the frequentist approach and refer readers to the comprehensive treatment of Bayesian methods for clinical trials in [Berry et al. \(2011\)](#). On the other hand, we want to discuss here an irreconcilable difference, which is widely recognized and somewhat controversial, between frequentist and Bayesian inference at the conclusion of a clinical trial with a group sequential or adaptive design. Bayesian inference (e.g., credible sets for parameters) is based on the posterior distribution given the randomly stopped sample, and no adjustment is needed for early stopping or adaptive randomization. In contrast, frequentist inference such as confidence sets has to make adjustments to ensure the correctness of the prescribed coverage probability. In nonsequential designs, the difference between credible and confidence intervals is small for large sample sizes because of the central limit theorem and higher-order expansions of the posterior distribution ([Johnson 1970](#)) and the sampling distribution ([Gross and Lai 1996](#)) of the approximate pivot used to construct credible or confidence intervals. However, for group sequential designs, this large-sample theory no longer holds, and in fact, an approximate pivot in the fixed sample size case is no longer approximately pivotal in the group sequential setting, as will be explained in Chap. 7 which also describes how valid confidence intervals can be constructed by a resampling procedure, similar to [Efron’s \(1987\)](#) bootstrap method to construct confidence intervals based on samples of fixed size.