

Bayesian Adaptive Methods for Clinical Trials

Scott M. Berry, Bradley P. Carlin, J. Jack Lee, and Peter Müller

iv

Draft date: July 14, 2010

To
OUR FAMILIES

Contents

Foreword	xi
Preface	xiii
1 Statistical approaches for clinical trials	1
1.1 Introduction	1
1.2 Comparisons between Bayesian and frequentist approaches	4
1.3 Adaptivity in clinical trials	6
1.4 Features and use of the Bayesian adaptive approach	8
1.4.1 The fully Bayesian approach	8
1.4.2 Bayes as a frequentist tool	10
1.4.3 Examples of the Bayesian approach to drug and medical device development	12
2 Basics of Bayesian inference	19
2.1 Introduction to Bayes' Theorem	19
2.2 Bayesian inference	26
2.2.1 Point estimation	26
2.2.2 Interval estimation	27
2.2.3 Hypothesis testing and model choice	29
2.2.4 Prediction	34
2.2.5 Effect of the prior: sensitivity analysis	37
2.2.6 Role of randomization	38
2.2.7 Handling multiplicities	40
2.3 Bayesian computation	42
2.3.1 The Gibbs sampler	44
2.3.2 The Metropolis-Hastings algorithm	45
2.3.3 Convergence diagnosis	48
2.3.4 Variance estimation	49
2.4 Hierarchical modeling and metaanalysis	51
2.5 Principles of Bayesian clinical trial design	63
2.5.1 Bayesian predictive probability methods	64

2.5.2	Bayesian indifference zone methods	66
2.5.3	Prior determination	68
2.5.4	Operating characteristics	70
2.5.5	Incorporating costs	78
2.5.6	Delayed response	81
2.5.7	Noncompliance and causal modeling	82
2.6	Appendix: R Macros	86
3	Phase I studies	87
3.1	Rule-based designs for determining the MTD	88
3.1.1	Traditional 3+3 design	88
3.1.2	Pharmacologically guided dose escalation	91
3.1.3	Accelerated titration designs	92
3.1.4	Other rule-based designs	92
3.1.5	Summary of rule-based designs	92
3.2	Model-based designs for determining the MTD	93
3.2.1	Continual reassessment method (CRM)	94
3.2.2	Escalation with overdose control (EWOC)	102
3.2.3	Time-to-event (TITE) monitoring	105
3.2.4	Toxicity intervals	109
3.2.5	Ordinal toxicity intervals	113
3.3	Efficacy versus toxicity	116
3.3.1	Trial parameters	117
3.3.2	Joint probability model for efficacy and toxicity	117
3.3.3	Defining the acceptable dose levels	118
3.3.4	Efficacy-toxicity trade-off contours	118
3.4	Combination therapy	121
3.4.1	Basic Gumbel model	122
3.4.2	Bivariate CRM	126
3.4.3	Combination therapy with bivariate response	127
3.4.4	Dose escalation with two agents	129
3.5	Appendix: R Macros	134
4	Phase II studies	137
4.1	Standard designs	137
4.1.1	Phase IIA designs	138
4.1.2	Phase IIB designs	140
4.1.3	Limitations of traditional frequentist designs	142
4.2	Predictive probability	142
4.2.1	Definition and basic calculations for binary data	143
4.2.2	Derivation of the predictive process design	146
4.3	Sequential stopping	150
4.3.1	Binary stopping for futility and efficacy	150
4.3.2	Binary stopping for futility, efficacy, and toxicity	151

4.3.3	Monitoring event times	154
4.4	Adaptive randomization and dose allocation	155
4.4.1	Principles of adaptive randomization	155
4.4.2	Dose ranging and optimal biologic dosing	163
4.4.3	Adaptive randomization in dose finding	167
4.4.4	Outcome adaptive randomization with delayed survival response	168
4.5	Hierarchical models for phase II designs	173
4.6	Decision theoretic designs	176
4.6.1	Utility functions and their specification	176
4.6.2	Screening designs for drug development	179
4.7	Case studies in phase II adaptive design	183
4.7.1	The BATTLE trial	183
4.7.2	The I-SPY 2 trial	189
4.8	Appendix: R Macros	191
5	Phase III studies	193
5.1	Introduction to confirmatory studies	193
5.2	Bayesian adaptive confirmatory trials	195
5.2.1	Adaptive sample size using posterior probabilities	196
5.2.2	Futility analyses using predictive probabilities	200
5.2.3	Handling delayed outcomes	204
5.3	Arm dropping	208
5.4	Modeling and prediction	211
5.5	Prior distributions and the paradigm clash	218
5.6	Phase III cancer trials	221
5.7	Phase II/III seamless trials	228
5.7.1	Example phase II/III trial	230
5.7.2	Adaptive design	231
5.7.3	Statistical modeling	232
5.7.4	Calculation	233
5.7.5	Simulations	235
5.8	Case study: Ablation device to treat atrial fibrillation	241
5.9	Appendix: R Macros	247
6	Special topics	249
6.1	Incorporating historical data	249
6.1.1	Standard hierarchical models	250
6.1.2	Hierarchical power prior models	252
6.2	Equivalence studies	260
6.2.1	Statistical issues in bioequivalence	261
6.2.2	Binomial response design	263
6.2.3	2×2 crossover design	265
6.3	Multiplicity	268

6.3.1	Assessing drug safety	269
6.3.2	Multiplicities and false discovery rate (FDR)	275
6.4	Subgroup analysis	276
6.4.1	Bayesian approach	276
6.4.2	Bayesian decision theoretic approach	277
6.5	Appendix: R Macros	280
References		281
Author index		297
Index		303

Foreword

It's traditional to get a foreword written by an *éminence grise*, generally an aging researcher who has seen better days. I can provide plenty of *grise* although I am possibly a bit short on *éminence*. Perhaps I best qualify through sheer long-service in trying to promote Bayesian clinical trials, having started my small contribution to this epic effort nearly 30 years ago with Laurence Freedman, eliciting prior opinions from oncologists about the plausible benefits of new cancer therapies.

This fine book represents the most recent and exciting developments in this area, and gives ample justification for the power and elegance of Bayesian trial design and analysis. But it is still a struggle to get these ideas accepted. Why is this? I can think of four main reasons: ideological, bureaucratic, practical and pragmatic.

By *ideological*, I mean the challenge facing the “new” idea of using probability theory to express our uncertainty about a parameter or existing state of the world – our epistemic uncertainty. Of course “new” is ironic, given it is nearly 250 years since Bayes formalized the idea, but the idea is still unfamiliar and disturbing to those brought up on classical ideas of probability as long-run frequency. One can only sympathize with all that effort to master the correct definition of a p -value and a confidence interval, only to be told that the intuitive meanings can be right after all.

I really enjoy introducing students to this beautiful idea, but tend to leave Bayes' theorem to subsequent lectures. In fact I sometimes feel the role of Bayes' theorem in Bayesian analysis is overemphasized: the crucial element is being willing to put a distribution over a parameter, and it is not always necessary even to mention the “B-word.” Natural examples include models for informative dropout in clinical trials, and the size of possible biases in historical studies: in these situations there may be no information in the data about the parameter, and so Bayes' theorem is not used.

But of course there are *bureaucratic* obstacles: as the authors of this book make clear, regulatory agencies perform a gate-keeping role where the Neyman-Pearson framework of decision-making without a loss function still has merits. Although the posterior distribution tells us what it is reasonable to believe given the evidence in a specific study, the regulators do need to

consider a continuous sequence of drug approval decisions. So quantifying Type I and Type II error can still be a valuable element of trial design, and one that is excellently covered in this book.

Then there are *practical* problems: can we actually do the analysis, or is the mathematics too tricky and there's no software to help us along? The authors have done a great job in discussing computation and providing software, but I am sure would still admit that there's some way to go before all these wonderful techniques are easily available to the average trial designer. But it will happen.

Finally, the crucial *pragmatic* test. Do these techniques help us do things we could not do before? This has been the factor that has led to increasingly widespread penetration of Bayesian methods into subject domains over the last 20 years or so: people can fit models and make inferences that were previously impossible or very cumbersome. And this is where this book wins hands down, since adaptive trials are so natural, ethical and efficient, that everyone wants to do them.

This book, based on the many years of cumulative experience of the authors, manages to deal with all these difficulties. Adaptive studies are a perfect application for a Bayesian approach, and I am confident that this book will be a major contribution to the science and practice of clinical trials.

DAVID J. SPIEGELHALTER
MRC Biostatistics Unit and University of Cambridge
April 2010

Preface

As has been well discussed, the explosion of interest in Bayesian methods over the last 10 to 20 years has been the result of the convergence of modern computing power and efficient Markov chain Monte Carlo (MCMC) algorithms for sampling from and summarizing posterior distributions. Practitioners trained in traditional, frequentist statistical methods appear to have been drawn to Bayesian approaches for three reasons. One is that Bayesian approaches implemented with the majority of their informative content coming from the current data, and not any external prior information, typically have good frequentist properties (e.g., low mean squared error in repeated use). Second, these methods as now readily implemented in WinBUGS and other MCMC-driven software packages now offer the simplest approach to hierarchical (random effects) modeling, as routinely needed in longitudinal, frailty, spatial, time series, and a wide variety of other settings featuring interdependent data. Third, practitioners are attracted by the greater flexibility and adaptivity of the Bayesian approach, which permits stopping for efficacy, toxicity, and futility, as well as facilitates a straightforward solution to a great many other specialized problems such as dose-finding, adaptive randomization, equivalence testing, and others we shall describe.

This book presents the Bayesian adaptive approach to the design and analysis of clinical trials. The ethics and efficiency of such trials can benefit from Bayesian thinking; indeed the Food and Drug Administration (FDA) Center for Devices and Radiological Health (CDRH) has been encouraging this through its document *Guidance for the Use of Bayesian Statistics*; see www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm071072.htm. The FDA Center for Drug Evaluation and Research (CDER) and Center for Biologics Evaluation and Research (CBER) has issued its own *Guidance for Industry: Adaptive Design Clinical Trials for Drugs and Biologics*; www.fda.gov/downloads/Drugs/GuidanceComplianceRegulatoryInformation/Guidances/UCM201790.pdf. This document also mentions Bayes, albeit far less prominently. The recent series of winter Bayesian biostatistics conferences at the University

of Texas M.D. Anderson Cancer Center in Houston are also testament to the growing role Bayesian thinking plays in this field.

The outline of the book is as follows. In Chapter 1 we summarize the current state of clinical trial design and analysis, present the main ideas behind the Bayesian alternative, and describe the potential benefits of such an alternative. We also describe what we mean by the word “adaptive” in the book’s title. Chapter 2 then gives an overview of the basic Bayesian methodological and computational tools one needs to get started as a Bayesian clinical trialist. While this whirlwind tour is not a substitute for a full course in Bayesian methods (as from Gelman et al., 2004, or Carlin and Louis, 2009), it should enable those with a basic understanding of classical statistics to get “up and running” on the material. This chapter also includes overviews of hierarchical modeling (with special emphasis on its role in Bayesian metaanalysis) and the basics of Bayesian clinical trial design and analysis. The idea here is to establish the basic principles that will be expanded and made phase- and endpoint-specific in subsequent chapters.

The next two chapters of the book (Chapters 3–4) follow standard clinical trials practice by giving Bayesian tools useful in “early” and “middle” phase clinical trials, roughly corresponding to phases I and II of the U.S. drug regulatory process, respectively. While our own professional affiliations have led us to focus primarily on oncology trials, the techniques we describe are readily adapted to other disease areas. We also place primary emphasis on “partially Bayesian” designs that concentrate on probability calculations utilizing prior information and Bayesian updating while still maintaining good frequentist properties (power and Type I error). An exception to this general rule is Section 4.6, where we discuss “fully Bayesian” designs that incorporate a utility function (and often more informative priors) within a more formal decision-theoretic framework. Chapter 4 also contains brief reviews of two recent trials utilizing Bayesian adaptive designs, BATTLE and I-SPY 2.

Chapter 5 deals with late (phase III) studies, an important area and the one of potentially greatest interest to statisticians seeking final regulatory approval for their compounds. Here we emphasize modern adaptive methods, seamless phase II–III trials for maximizing information usage and minimizing trial duration, and describe in detail a case study of a recently approved medical device. Finally, Chapter 6 deals with several important special topics that fit into various phases of the process, including the use of historical data, equivalence studies, multiplicity and multiple comparisons, and the related problem of subgroup analysis. The historical data material is particularly relevant for trials of medical devices, where large historical databases often exist, and where the product being evaluated (say, a cardiac pacemaker) is evolving slowly enough over time that worries about the exchangeability of the historical and current data are relatively low.

Since this is not a “textbook” per se, we do not include homework prob-

lems at the end of every chapter. Rather, we view this book as a handbook enabling those engaged in clinical trials research to update and expand their toolkit of available techniques, so that Bayesian methods may be used when appropriate. See <http://www.biostat.umn.edu/~brad/data3.html> and <http://biostatistics.mdanderson.org/SoftwareDownload/> on the web for many of our datasets, software programs, and other supporting information. The final sections of Chapters 2–6 link to these software sites and provide programming notes on the R and WinBUGS code we recommend.

We owe a debt of gratitude to those who helped in our writing process. In particular, the second author is very grateful to Prof. Donald Berry and the Division of Quantitative Sciences at the University of Texas M.D. Anderson Cancer Center for allowing him to spend his fall 2008 sabbatic time in the same U.S. state as the other three authors. Key staff members worthy of special mention are Martha Belmares and the incomparable Lydia Davis. Sections 1.1, 1.2, 1.4, and 2.4 are based on Prof. Berry’s previous work in their respective areas. Indeed, many sections of the book owe much to the hard work of our research colleagues, including Lee Ann Chastain, Nan Chen, Jason Connor, Laura Hatfield, Brian Hobbs, Haijun Ma, Ashish Sanil, and Amy Xia. We also thank the 2010 spring semester “Topics in Clinical Trials” class at Rice University and the University of Texas Graduate School of Biomedical Sciences,” taught by the third author, for commenting on the text and testing the supporting software. Rob Calver and David Grubbs at Chapman and Hall/CRC/Taylor & Francis Group were pillars of strength and patience, as usual. Finally, we thank our families, whose ongoing love and support made all of this possible.

SCOTT M. BERRY
BRADLEY P. CARLIN
J. JACK LEE
PETER MÜLLER

College Station, Texas
Minneapolis, Minnesota
Houston, Texas
Houston, Texas

March 2010

Index

- 2x2 crossover design, 265

- adaptive dose allocation, 167
- adaptive MCMC, *see* Markov chain Monte Carlo (MCMC) methods
- adaptive randomization, 7, 155–173
- adherence, *see* noncompliance
- Akaike Information Criterion (AIC), 32
- assurance
 - Bayesian, 80
- autocorrelation time, 50
- auxiliary variables, 212

- backward induction, 8, 9
- batching, 50
- Bayes factor, 31
 - large sample approximations to, 32
- Bayes' Theorem, 9, 21
- Bayesian Central Limit Theorem, 28
- Bayesian inference
 - hierarchical modeling, 20, 51–63
 - hypothesis testing, 29–34
 - interval estimation, 27
 - multiplicities, 30, 40–42, 268–276
 - point estimation, 26
 - prediction, 34–37
 - prior sensitivity, 37–38
 - role of randomization, 38–40
 - using historical data, 249–260
- Bayesian Information Criterion (BIC), 31
- beta-binomial distribution, 143, 201
- bioequivalence, 260
 - average, 261
 - individual, 261
 - population, 261
- bivariate CRM, 126
 - with bivariate response, 127
- BRugs language, 59–61, 64
- BUGS language, 58–59
- burn-in, 45

- causal inference, 82–86
 - g*-estimation, 83
 - exclusion restriction, 84
 - frequentist approaches, 83
 - monotonicity assumption, 84
 - Rubin causal model, 83
- combination therapy, 7, 121–129
- complete conditional distribution, *see* full conditional distribution
- compliance, *see* noncompliance
- confidence interval
 - Bayesian, 27
- confirmatory trials, 193
- continual reassessment method (CRM), 94–102, 115
 - bivariate, 126
 - with bivariate response, 127
- copula model, 123
- cost-effectiveness, 78
- credible set, 27
 - equal tail, 27
 - highest posterior density (HPD), 28
- curse of dimensionality, 42

- decision theoretic designs, 8, 176–182
- decision theory, 9, 79
- design of experiments, 63–86
 - for cost effectiveness, 79

- Deviance Information Criterion (DIC), 30, 32–34
- Dirichlet-multinomial model, 77
- dose allocation
 - adaptive, 155
- dose-finding, 87
- dose-limiting toxicity (DLT), 87
- drug safety data, 269
- dynamic linear model (DLM), 232

- effective sample size, 50
- empirical Bayes (EB), 20
- equivalence testing, 29, 260–268
- escalation with overdose control (EWOC), 102–105
- exceedance probability, 271
- experimental design, *see* design of experiments, 70

- Fibonacci sequence, 88
- Fisher information matrix
 - observed, 28
 - generalized, 28
- Fisher’s exact test, 271
- forward filtering backward sampling (FFBS) algorithm, 164
- full conditional distribution, 25, 44
- futility, 7, 13, 64, 77, 139, 150–154, 164, 165, 167, 200–202

- Gaussian/Gaussian model, 21
- Gibbs sampling, *see* Markov chain Monte Carlo (MCMC) methods

- hyperparameter, 19
- hypothesis testing
 - Bayesian, 29–34
 - frequentist, 2, 29–30
 - null hypothesis, 29

- incremental cost effectiveness ratio (ICER), 81
- indifference zone, 66, 262
- intention to treat (ITT), 82
- interim analysis, 77

- isotonic regression, 109–111

- Laplace’s Rule of Succession, 36
- Likelihood Principle, 4, 30
- longitudinal model
 - for adaptive design, 232
- loss function, 9

- marginal distribution
 - improper, 32
- Markov chain Monte Carlo (MCMC)
 - methods, 20, 43–51
 - adaptive, 44
 - convergence
 - burn-in period, 45
 - convergence diagnosis, 48–49
 - Gibbs sampling, 20, 44–45
 - Metropolis-Hastings algorithm, 20, 45–48
 - slice sampling, 43
 - thinning, 49
 - variance estimation, 49–51
- maximum likelihood estimate (MLE), 27
- maximum tolerated dose (MTD), 87, 114, 121
- metaanalysis, 4, 51–63
 - for a single success proportion, 52
- Metropolis-Hastings algorithm, *see* Markov chain Monte Carlo (MCMC) methods
- minimax design, 139
- model selection
 - via Bayes factors, 31
 - via the DIC criterion, 32–34
- multiple endpoints, 77

- noncompliance, 82–86
- noninferiority testing, 263
- normal/normal model, 21
 - with unknown sampling variance, 24

- operating characteristics, 10, 70

- p-value, 29

- penalized likelihood criteria, 32
- pharmacodynamics, 87
- pharmacokinetics, 87
- pharmacologically guided dose
 - escalation (PGDE) method, 91
- Phase I studies
 - 3+3 design, 88
 - accelerated titration designs, 92
 - model-based methods, 88, 93–115
 - pharmacologically guided dose
 - escalation, 91
 - rule-based methods, 88–93
- pick-the-winner design, 141
- piecewise exponential model, 242
- posterior distribution, 4, 20
- power, 70–78
 - Bayesian, 72, 76
 - frequentist, 72, 77
- precision, 21
- predictive probability, 64–66, 200–203
- preposterior, 63
- preposterior analysis, 72
- prior distribution, 4, 9, 19
 - analysis, 69, 72, 76, 80, 207
 - clinical, 69
 - community of priors, 69
 - conjugate, 20, 23, 35, 45
 - design, 69, 72, 76, 80, 207
 - enthusiastic, 69
 - improper, 24
 - noninformative, 23, 69
 - Jeffreys, 25
 - regulatory, 207
 - skeptical, 69
- prospectively adaptive trial, 195

- quadrature methods, 42
- quality-adjusted life years (QALYs),
 - 8, 176

- range of equivalence, *see* indifference zone
- Rayleigh distribution, 73

- sample size estimation, 70
- Schwarz criterion, *see* Bayesian Information Criterion

- seamless trials, 7, 12, 194, 211, 228–240
 - inferentially seamless, 229
 - operationally seamless, 229
- Simon’s two-stage design, 139, 146, 147, 159
 - software for, 140
- slice sampling, *see* Markov chain Monte Carlo (MCMC) methods
- state space model, 163
- stochastic curtailment, 64
- stopping for futility, *see* futility

- target toxicity level (TTL), 87
- TITE-CRM method, 106
- two one-sided tests (TOST), 262
- Type I error, *see* power
- Type II error, *see* power

- utility function, 9, 176

- WinBUGS, *see* BUGS language