BAYESIAN DOSE-FINDING PROCEDURE BASED ON INFORMATION CRITERION AND EFFICACY-TOXICITY TRADE-OFFS

by

Lei Gao

A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
In Partial fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Statistical Science

Committee:

_____	Dr. William F. Rosenberger Dissertation Director
_____	Dr. Daniel B. Carr Committee Member
_____	Dr. Jacqueline M. Hughes-Oliver Committee Member
_____	Dr. Guoqing Diao Committee Member
_____	Dr. Zorayr Manukyan Committee Member
_____	Dr. William F. Rosenberger Department Chair
_____	Dr. Kenneth S. Ball Dean, The Volgenau School of Information Technology and Engineering

Date: _____ Spring 2014
George Mason University
Fairfax, VA

Bayesian Dose-Finding Procedure Based on Information Criterion and Efficacy-Toxicity Trade-offs

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

By

Lei Gao
Master of Science
University of New Hampshire, 2010
Bachelor of Science
Peking University, 2005

Director: Dr. William F. Rosenberger, Professor
Department of Department of Statistics

Spring 2014
George Mason University
Fairfax, VA

Dedication

I dedicate this dissertation to my parents for their unconditional love, encouragement and support.

Acknowledgments

The three year journey within the Department of Statistics at George Mason University has completely shaped my vision of statistics. Not only do I become increasingly passionate and confident about the subject the more deeply I have explore it, but for the first time I have found the field that I can continuously devote my passion and talent to – adaptive designs in clinical trials – for the rest of my life. Therefore I am deeply indebt to committee chair Dr. William F. Rosenberger, who led me to the fascinating field. He is a great supervisor and collaborator with insightful vision and patience. He is an extraordinary mentor by placing great confidence on me, which stimulates me to move from good to great. He provided me with various resources to achieve my career and research goals as well as professional growth. He pointed out my weakness of communication during our first meeting, placed me in the consulting center, and patiently reviewed my writing, so that I can continuously improve my communication skills and benefit from doing so in a long term. He is a great friend with good taste of humor, food, wine and coffee.

Special thanks go to committee member Dr. Zorayr Manukyan, whose creative mind and passion add extra and valuable dimensions to my research. He generously provided me with internships to refine my statistical expertise and presentation skills under his supervision. He sets a good career example and provided wise career guidance.

I want to thank other committee members for their constant support throughout my research. In particular, Dr. Daniel B. Carr led me to the interesting world of data visualization; Dr. Jacqueline M. Hughes-Oliver always provides insightful and constructive feedbacks; Dr. Guoqing Diao guided me through survival analysis and provided computational resources.

I also want to thank the faculty and staff members in the Department of Statistics and researchers outside the department. By working with Dr. Clifton D. Sutton, I became confident in communicating statistical ideas with lay audience. I learned a lot of practical and statistical skills from him. His notes on statistical inference made the complex theory simple and elegant. Dr. James Gentle is extremely knowledgable. I benefited from his computational statistics course, and gained all the important computational skills through intensive exercises. Dr. Anand N. Vidiyashankar is also extremely knowledgable. I learned and refined my modeling skills by taking his longitudinal analysis course. Dr. Liansheng Tang provided me with access to his UNIX cluster. Mrs Liz Quigley is always supportive and efficient in helping me set up appointments and book rooms. She organized the annual department picnics. Dr. Alex Sverdlov, Dr. Sarah Zohar Dr. Sandrine Katsahian and Dr. Frederique Schortgen are pleasant and talented collaborators to work with. In particular, Dr. Zohar provided me with guidance on prior elicitation and pointed me to sequential testing stopping rules, which greatly refined my research. Dr. John M. White's lecture on JAGS was an eye-opening experience for me to enter the Bayesian world. Dr. Paul Gallo, Dr. Lanju Zhang, Dr. Yanqiong Zhang and Dr. Jinhui Li provided me with insightful and honest career advice.

I want to thank my friends at George Mason, who made the journey joyful and pleasant. Brad Patterson provided accommodation on my first day at school. I enjoyed the intelligent conversations with him. He provided me with wise advice, and set a great example of excellence. Xin Cao was my driving coach and is always a supportive and reliable friend. Zhi Li is a kind and caring meal buddy. Major Chuck Weko is an intelligent friend with great presentation and communication skills. The rest of the list include Yang Wang, Mengdie Yuan, Parwen Parhat, Ping Ren, Hui Shao, Daniel Saxton, Krista Heim, Xuan Ye, Seunghye Wilson, Gay Cororaton and Jun Xie.

Last but not least, I want to thank my supervisors and friends from my previous colleges, without whose support I cannot go this far. Special thanks go to Dr. Don Hadwin, Dr. Ernst Linder, Dr. Junhao Shen, Dr. Linyuan Li, Dr. Yintang Zhang, Qihui Li, Yanni Chen, Kewei Lu, Yibin Pan, Yi Xu, Zhangbo Guo, Liang Wang, Jinsong Wu, Wei Yuan, Chengwei Yuan, Zhaoyu Yin, Hyung Kim, Zhengwei Liu, Ping Zhong, Gang Chen, Qifeng Bai, Bintao Cao, Kelei Wang, Xiaowei Sun, Wei Wang, Jiaping Liu, Fuqiang Liu, Shisheng Zhen, Yongyi Sun, Yun Huang, Fei Han, Shulian Shang, Jian Guo, Liping Hou, Xiang Xiang, Jinyu Ma, Bin Sun, Jie Liu and Somchay Thavisouk.

Table of Contents

List of Tables

List of Figures

Abstract

BAYESIAN DOSE-FINDING PROCEDURE BASED ON INFORMATION CRITERION AND EFFICACY-TOXICITY TRADE-OFFS

Lei Gao, PhD

George Mason University, 2014

Dissertation Director: Dr. William F. Rosenberger

In dose-finding studies with toxicity-efficacy responses, utility functions and Bayesian procedures are used to find a single optimal dose with ethical toxicity-efficacy trade-offs. We demonstrate that the design can have convergence issues when the prior information is misspecified. We propose to incorporate an information penalty to obtain multiple-dose allocation with efficient ethical measures. A coefficient is introduced to control the trade-off between information gain and ethical gain. We conduct simulations using Markov chain Monte Carlo (MCMC) algorithms to examine the convergence of Bayesian dose finding designs and investigate their operating characteristics. Different stopping rules are considered and their benefits are demonstrated by simulation. Guidance on how to select design parameters are given in two hypothetical trial redesigns.

Chapter 1: Background and Literature Review

1.1 Introduction

A clinical trial is a carefully designed and controlled experiment to test the safety and efficacy of a drug. Phase I and II trials are generally small, and conducted on multiple doses. The goal is typically to identify a safe (phase I) and effective (phase II) dose that can be passed to a large-scale phase III randomized trial. A drug can be marketed only if it passes all the phases. According to Chuang-Stein and D'Agostino (2007), about 50% of drugs fail in phase III due to insignificant efficacy or underestimated toxicity, which causes ethical and economic losses. Therefore, to increase the chance to find effective drugs and reduce the risks, there is increasing demand for innovative and reliable designs in early phase trials. In this proposal we examine optimal designs for combined efficacy-toxicity outcomes.

The goal of dose finding studies in phase I trials, particularly in oncology, is to find the maximum tolerate dose (MTD) for further confirmatory studies. Researchers have recognized that the key in these studies is to understand the dose-response relationship (Gaydos et al., 2006). O'Quigley and Zohar (2006) surveyed several recent dose-finding studies and called for model-based dose-finding procedures. The information we learn from an experiment depends on the way patients are assigned to different doses: we wish to obtain maximum information about the dose-response relationship. Therefore, adaptive sequential designs are used in practice, in which the information learned from previous assignments is used to determine later assignments. Bayesian methods have become popular because they incorporate prior information and update it using data. This can lead to more efficient learning about the dose-response relationship.

In addition to maximizing the information gain, there are ethical considerations which

require that patients not be assigned to highly toxic doses. These may contradict the information goal, because in designs that have maximum information as their primary objective, a portion of participants would necessarily be exposed to high doses with high probability of unacceptable toxicity. Despite the arguable statement that it is the price to be paid for accruing accurate knowledge about a drug from which a large population can benefit later, it is unethical to assign participants to potentially toxic doses without even knowing the risk of toxicity.

The global theme of this dissertation research is to determine appropriate and robust procedures for dose-response studies from both the information and ethical perspectives. Specifically, we examine the convergence issue of best intention Bayesian methods through simulation. We propose an information-based Bayesian procedure. The new method is used to produce efficient and accurate parameter estimates. Various measures are introduced to evaluate the new method's operating characteristics. The rest of the dissertation is organized as follows. In Chapter 1, we review optimal design theory, the bivariate Gumbel model and commonly used dose-finding procedures. In Chapter 2, we introduce Bayesian dose-finding procedures based on the information criterion and examine its convergence property. In Chapter 3, we discuss an ill-behaved example with best intention Bayesian designs and propose information-based design as a robust alternative when there is prior misspecification. In Chapter 4, we discuss different stopping rules and demonstrate their advantages in simulation. In Chapter 5, we provide guidance on how to select design parameters by redesigning two hypothetical trials. We conclude the dissertation with discussion in Chapter 6.

1.2 Optimal designs for the dose-response relationship

Gaydos et al. (2006) argue that "insufficient exploration of the dose-response relationship is a shortcoming of clinical drug development, and failure to characterize dosing early is often cited as a key contributor to the high late-stage attrition rate currently faced by the industry". They mention that efficient learning of the dose-response relationship can

increase the chance to detect a successful drug and dose, reduce the late-stage attrition rate, and reduce the clinical development timeline. In this chapter we define the dose-response relationship and the information gain in the dose-finding context.

1.2.1 Dose-response relationship

We assume that patients from a population respond to a drug according to a identical distribution. Then the dose-response relationship is commonly defined as the relationship between doses and the probability of outcomes. Mathematically, it is a map from the dose space, a subset of $\mathbb{R}$, to $[0,1]^r$ when there are r outcomes. One can use parametric, semiparametric or nonparametric (Manukyan, 2009) models to represent the dose-response relationship. This dissertation will focus on parametric modeling.

The models and design methods depend on the type of toxicity measure. If one wants to determine the MTD, patients enrolled in the study are tested for toxicity. Depending on the measures, the repones can be binary– toxic or not (Atkinson et al. 2007), ordered–mildly toxic to severely toxic (Perevozskaya et al. 2003), or counts (Fedorov et al, 2007). Toxicity can also be time-to-event if the response is not immediate (Yuan and Yin, 2009).

For a binary toxicity response, we assume a Bernoulli distributed random variable, where $Y = 1$ denotes toxicity and $Y = 0$ otherwise. Then the dose-response relationship is a function that maps the dose x to the toxicity probability $P(Y = 1)$ at that dose. The function is often assumed to be non-decreasing. A commonly assumed model for the binary response is the location-scale family:

$$P(Y = 1; x) = F\left(\frac{x - \alpha}{\beta}\right), \tag{1.1}$$

where F is a distribution function. Typically, F is the logistic function. Sometimes the toxicity probability does not approach one, then one can use a modified logistic regression model with upper and lower limiting probabilities (Manukyan and Rosenberger, 2010).

1.2.2 Information in designs

We define a design to be a discrete measure on a transformed dose space $\mathcal{X} = \mathbb{R}$,

$$\xi = \begin{pmatrix} x_1 & x_2 & \cdots & x_{n-1} & x_n \\ w_1 & w_2 & \cdots & w_{n-1} & w_n \end{pmatrix}, \tag{1.2}$$

where x_i's are doses and w_i's are corresponding weights with sum one. The single point design at x is denoted δ_x. In practice, the doses are usually preselected and task remains to determine the proportion of patients assigned to each dose.

In optimal designs, we gauge the information about dose-response using efficiency of estimation for a specific model. The more accurate the estimation, the more information we learn about the dose-response relationship. The accuracy can be measured as a confidence ellipsoid dominated by the Fisher information matrix of the design ξ (1.2):

$$M(\xi;\theta) = \sum_{k=1}^{n} w_k I(x_k;\theta) \quad \text{with} \quad I_{ij}(x;\theta) = -\mathrm{E}\left\{ \frac{\partial^2 \log f(x,\theta)}{\partial\theta_i \partial\theta_j} \right\}. \tag{1.3}$$

Some investigators have suggested that the Fisher information matrix may not be good enough to characterize the variance structure with limited sample size. In this case, researchers have suggested investigating higher order terms in the Taylor expansion of the log-likelihood function (Dette and Grigoriev, 2000).

1.2.3 Design criterion and attainability

From now on we restrict ourselves to the Fisher information matrix as measure of information about the dose-response relationship. Then we determine a design, which minimizes the asymptotic variance ellipsoid, to gain maximum information. We have multiple objectives to minimize: the volume, the mean axis length, the maximum axis length, or the span along a specified direction. These lead us to the D, A, E and c optimality criteria (Atkinson

et al. 2007). We denote the optimal criterion function by $\Phi(\xi)$ for the design ξ (1.2) and summarize the criteria in Table 1.1. We will mainly use D-optimality in the dissertation research. However, we note that, by minimizing the volume of the asymptotic ellipsoid, the D-optimality criterion may result in a ellipsoid with disproportionate axes, which may not give accurate estimation for some parameter components.

Following Whittle (1973), we define the directional derivatives of $\Phi(M(\xi))$ along design η at ξ as

$$\phi(\xi, \eta; \theta) = \lim_{\alpha \to 0} \frac{1}{\alpha} \{\Phi(\xi) - \Phi(M([1-\alpha]\xi + \alpha\eta)\}. \tag{1.4}$$

Whittle (1973) show that the directional derivative can be obtained by the following formula:

$$\phi(\xi, \eta; \theta) = -Tr\left[\left(\frac{\partial \Phi}{\partial M}\right)^T\Big|_{\xi} M(\eta)\right] + Tr\left[\left(\frac{\partial \Phi}{\partial M}\right)^T\Big|_{\xi} M(\xi)\right].$$

We denote $\phi(\xi, \delta_x; \theta)$ by $\phi(\xi, x; \theta)$. Table 1.2 include the directional derivatives for the D, A, E and c optimal designs.

The relation of the optimum design and its derivatives are given by the generalized equivalence theorem (Manukyan and Rosenberger, 2010),

1. ξ^* is an Φ-optimal design, i.e., ξ^* minimizes Φ;
2. $\xi^* = \arg\min_\xi \max_x \phi(\xi, x) \leq 0$;
3. $\phi(\xi^*, x) \leq 0$. The equality is attained at the design points.

The third item can be used to verify if the optimum is attained.

1.2.4 Example: D-optimal design for the two-parameter logistic model

As an example, we investigate the optimal design for a two parameter logistic model, also known as the location-scale family model. The MTD can be considered as a quantile

corresponding to a target proportion of toxicity Γ. If μ is the quantile corresponding to Γ, it is computed by the equation $F((\mu-\alpha)/\beta) = \Gamma$. Thus $\mu = \alpha + \beta F^{-1}(\Gamma)$. The reason of using this canonical form is that

$$\mathrm{Var}(\hat{\mu}) = \mathrm{Var}(\hat{\alpha}) + \mathrm{Var}(\hat{\beta})[F^{-1}(\Gamma)]^2 + 2F^{-1}(\Gamma)\mathrm{Cov}(\hat{\alpha},\hat{\beta}) \tag{1.5}$$

has a closed form.

Now we define $p_i = P(\text{toxicity}|x)$ as $p_i = F((z-\alpha)/\beta)$ for $z \in \mathcal{X}$ some dose space in $\mathbb{R}$ and F is a distribution function. For $F(x) = e^x/(1+e^x)$, the Fisher information matrix for a single dose z is

$$M(\delta_z) = \frac{1}{\beta}\frac{[f(z)]^2}{F(z)[1-F(z)]}\begin{pmatrix} 1 & z \\ z & z^2 \end{pmatrix}.$$

Suppose the optimal design with n support points is given by (1.2) with the corresponding Fisher information matrix given by (1.3). The D-optimal design, ξ^*, minimizes the volume of the asymptotic variance ellipsoid, which is inversely proportional to the determinant of the Fisher information matrix over Ξ, the space of all discrete measures :

$$\xi^* = \arg\min_{\xi\in\Xi}\ \Phi(\xi) = \arg\min_{\xi\in\Xi}\ \log\det[M^{-1}(\xi)].$$

It is well known that local D-optimal designs are supported on two points at the 17.6th and 82.4th percentiles (Atkinson et al. 2007). For $\alpha = 0$ and $\beta = 1$ the optimal points are

$$\xi^* = \begin{pmatrix} z_1 = -1.5434 & z_2 = 1.5434 \\ 0.5 & 0.5 \end{pmatrix},$$

with $F(z_1) = 0.1760$ and $F(z_2) = 0.8240$.

The c-optimality criterion minimizes $[1, F^{-1}(\Gamma)]M^{-}[1, F^{-1}(\Gamma)]^T$, which approximates

(1.5). Here M^- is the generalized inverse of M when M is singular. Wu (1988) shows that when $z = [F^{-1}(\Gamma) - \alpha]/\beta$ and $\underline{z} \leq z \leq \bar{z}$, the c-optimal design puts weight on one point z. Otherwise, the design has two support points. To prove the result, Wu (1988) extends a geometric method, which is proposed by Elfving (1952) to solve the c-optimal problems for linear models. In particular, Wu (1988) uses the geometric method to determine the critical points $(\underline{z}, \bar{z})$ for a one point design to transition to a two point design. In the geometric method, they use a transform g and investigate the problem in the new coordinates $(V(z), zV(z))$ with $V(z) = \exp(0.5z)/(1 + \exp(z))$. Wu (1988) shows that $(\underline{z}, \bar{z})$ is the solution to the system of equations

$$\begin{aligned} r(\underline{z}) &= s(\bar{z}, \underline{z}); \\ r(\bar{z}) &= s(\underline{z}, \bar{z}), \end{aligned}$$

where $s(z_1, z_2) = [V(z_1)z_1 + V(z_2)z_2]/[V(z_1) + V(z_2)]$ is the slope of the line that connects $(V(z_1), z_1V(z_1))$ to $(V(z_2), z_2V(z_2))$, and $r = z + V(z)/V^{(1)}(z)$ is the tangent of the curve $(V(z), zV(z))$ at z. For example, if $\alpha = 0$ and $\beta = 1$, $c = (1, z)^T$. The resulting solution is given as follows:

1. When $\underline{z} \leq z \leq \bar{z}$, the optimal design puts weight 1 on z.

2. Otherwise, the optimal design is

$$\begin{pmatrix} \underline{z} & \bar{z} \\ 1 - w & w \end{pmatrix}$$

with

$$w = \frac{V(\underline{z})(\underline{z} - z)}{V(\underline{z})(\underline{z} - z) + V(\bar{z})(\bar{z} - z)}.$$

When $z = 2.8$, $w = 0.928$. The c-optimal design is

$$\begin{pmatrix} -2.399 & 2.399 \\ 0.072 & 0.928 \end{pmatrix}.$$

1.3 Bayesian sequential designs

As we note in Section 1.2.2, the local D-optimal design requires knowledge of the true parameter. In practice, we do not have such information. Therefore, we need to estimate the true parameter while conducting the experiment. This leads to the concept of a sequential design, in which patients are assigned sequentially according to an allocation rule. The information about the dose-response parameter is updated after every assignment, and the new information is used to evaluate the allocation rule for the next patient.

There are two approaches to sequential designs. In one approach, we estimate the parameter by the maximum likelihood estimator. In another approach, we estimate the parameter by a Bayesian estimator. In this research, we focus on the Bayesian sequential approach.

1.3.1 Bayesian D-optimal design

For the location-scale family (1.1), Chaloner and Larntz (1989) assume uniform distributions as the prior distributions for the parameters. They define the Bayesian Φ-optimality criterion:

$$\Psi(\xi) = \mathrm{E}_{\theta|\pi}\Phi(\xi, \theta) = \int \Phi(\xi, \theta) \cdot \pi(\theta|\mathcal{D}, \eta) d\theta, \tag{1.6}$$

where $\mathcal{D}$ is the data, $\Theta(\eta)$ is the prior distribution for θ, and $\pi(\theta|D, \Theta)$ is the posterior distribution given data $\mathcal{D}$. We use "$\mathrm{E}_{\theta|\pi}$" to denote the conditional expectation under the posterior distribution $\pi(\theta|\mathcal{D}, \eta)$ throughout the paper. This criterion is related to the

Shannon information (Atkinson et al., 2007).

Chaloner and Larntz (1989) define the corresponding directional derivative to be

$$D(x, \xi) = \mathrm{E}_{\theta|\pi} \phi(\xi, x; \theta),$$

where $\phi(\xi, x; \theta)$ is the local directional derivative at θ (1.4). Then the Bayesian generalized equivalence theorem for D-optimality is the equivalence of following three conditions:

1. $\xi^* = \arg\min_\xi \Psi$;

2. $\xi^* = \arg\min_\xi \max_x D(x, \xi)$;

3. $D(x, \xi^*) \leq 0, \forall x$. The equality is attained at the design points.

The last condition in the generalized equivalence theorem is used to determine if a design attains the optimum.

1.3.2 Bayesian D-optimal sequential design

The Bayesian D-optimal design is constructed iteratively to maximize the Bayesian D-optimal criterion (1.6). At the $n+1$th step, the new patient is assign to a dose x_{n+1} as follows (Haines, Perevozskaya, Rosenberger 2003):

$$x_{n+1} = \arg\min_x \quad E_{\theta|\pi_n} \Phi \left(\frac{\sum_{i=1}^n I(x_i; \theta) + I(x; \theta)}{n+1} \right), \tag{1.7}$$

where π_n is the posterior distribution given the data in the first n steps.

Roy, Ghosal and Rosenberger (2009) prove that the procedure in (1.7) converges whenever the dose space is restricted by an upper bound or not.

1.3.3 Best intention design

Fedorov et al. (2011) use the term "best intention design" to describe sequential designs that assign patients to the current best dose according to some measure. They contrast the

best intention designs to sequential optimal designs in the non-Bayesian context. In this dissertation, we say a design is a best intention design if it sequentially estimates the toxicity and/or efficacy probabilities and assigns patients to the dose at which a utility function of toxicity and efficacy probabilities is maximized or minimized. When the true parameter is known, the best intention designs assign patients to a unique dose, called a targeted dose, at which the utility function is maximized or minimized. In contrast, the local D-optimal design assigns patients to a combination of doses (Roy, Ghosal, Rosenberger, 2009).

The common choices for the utility function can involve the efficacy probability, the toxicity probability and the distance to the Γth toxicity quantile. For example, Fedorov et al. (2011) consider the following adaptive Robbins-Monro (ARM) design for finding the dose with 50% toxicity probability:

$$x_{n+1} = \arg\min_{x \in \mathcal{X}} \left| F(\hat{\theta}_{1n} + \hat{\theta}_{2n} x) - 0.5 \right|,$$

where $\hat{\theta}_{1n}$ and $\hat{\theta}_{1n}$ are the estimated parameters based on the data at step n. Although one expects the best intention design to converge to the targeted dose, Fedorov et al. (2011) show that these designs do not always converge.

1.4 Seamless phase I/II designs

Several issues arise when conducting separate phase I and II trials. First, although researchers carry out phase I studies to protect participants, it may not be efficient due to limited knowledge of the drug. Second, only limited information from phase I studies is carried to phase II studies. Third, it ignores the correlation between toxicity and efficacy probabilities.

Several researchers have made attempts to combine the two studies (Thall and Russell, 1998; Thall and Cook, 2004; Dragalin and Fedorov, 2006; Wang, 2007; Zohar, et al., 2009). They use bivariate parametric models to represent the dose-response relationship with the

two responses being probability of efficacy and probability of toxicity. The goals in these designs are to efficiently estimate the parameters and protect participants by controlling the toxicity probability. A trial using the Gumbel model with 52 patients can be found in de Lima et al (2008).

1.4.1 Bivariate D-optimal design

In the attempt to combine phase I and phase II trials, researchers are interested in modeling the toxicity and efficacy probabilities together. The Gumbel model and bivariate probit model are bivariate extensions of the logit and probit models.

In a Gumbel model, we have a pair of binary responses (Y, Z), where $Y = 1$ indicates efficacy and $Z = 1$ indicates toxicity. They equal zero otherwise. Let α be a parameter to characterize the correlation between toxicity and efficacy. We define the function (Dragalin and Fedorov 2006):

$$G(y, z) = F(y) \cdot F(z) \cdot \{1 + \alpha[1 - F(y)][1 - F(z)]\}, \quad -\infty < y, z < +\infty, \quad |\alpha| < 1,$$

and $F(y)$ is the logistic function $F(y) = 1/(1 + e^{-y})$.

The efficacy effects can be linear, $x_E = \mu_E + \beta_{E1}x$ (Dragalin and Fedorov, 2006), or quadratic, $x_E = \mu_E + \beta_{E1}x + \beta_{E2}x^2$ (Thall and Cook, 2004). The toxicity effect is usually monotone increasing and therefore is assumed to be linear, $x_T = \mu_T + \beta_T x$. Then the cell probabilities, $p_{yz} = \Pr(Y = y, Z = z)$, form a multinomial distribution, given by

$$\begin{aligned}
p_{11} &= \Pr(Y = 1, Z = 1|D = x; \theta) = G(x_E, x_T), \\
p_{10} &= \Pr(Y = 1, Z = 0|D = x; \theta) = G(x_E, +\infty) - G(x_E, x_T) = F(x_E) - G(x_E, x_T), \\
p_{01} &= \Pr(Y = 0, Z = 1|D = x; \theta) = G(+\infty, x_T) - G(x_E, x_T) = F(x_T) - G(x_E, x_T), \\
p_{00} &= \Pr(Y = 0, Z = 0|D = x; \theta) = 1 - G(x_E, +\infty) - G(+\infty, x_T) + G(x_E, x_T).
\end{aligned}$$

The Fisher information matrix $I(x; \theta)$ is calculated by Dragalin and Fedorov (2006). The

Fisher information matrix for the design ξ in (1.2) is

$$M(\xi;\theta) = \sum_i w_i I(x_i;\theta).$$

In a bivariate probit model (Dragalin, Fedorov, Wu, 2008; Wang, 2007), we need to model the efficacy and toxicity probabilities. Let $V = (V_E, V_T)$ follow a bivariate normal distribution with variance covariance matrix,

$$\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix},$$

where ρ is a parameter to characterize the correlation between toxicity and efficacy. We assume linear effects for both toxicity and efficacy:

$$\begin{cases} x_E = \mu_E + \beta_E x; \\ x_T = \mu_T + \beta_T x. \end{cases}$$

We denote $\theta = (\mu_E, \beta_E, \mu_T, \beta_T, \rho)$. We define the function (Dragalin, Fedorov, Wu, 2008),

$$p_{11}(x) = \int_{-\infty}^{x_E} \int_{-\infty}^{x_T} \frac{1}{2\pi|\Sigma|^{1/2}} \exp\left\{-\frac{1}{2} V \Sigma^{-1} V^T\right\} dV_T dV_E.$$

If we denote the cumulative distribution function of the standard normal distribution by F, then the toxicity and efficacy probabilities at dose level x are

$$\begin{aligned} &p_E = F(x_E), \quad p_{10} = F(x_E) - p_{11}, \quad p_{00} = 1 - p_{11} - p_{10} - p_{01}, \\ &p_T = F(x_T), \quad p_{01} = F(x_T) - p_{11}. \end{aligned}$$

All the cell probability p_{ij}'s follow a multinomial distribution.

1.4.2 Local D-optimal design

Recall that the local D-optimal design minimizes the following

$$\Phi(\xi, \theta) = \log \det \left\{ M^{-1}(\xi; \theta) \right\}, \tag{1.8}$$

over a design space Ξ. Note that once $p_{yz}(x; \theta)$ are known, one can derive the D-optimal design by the following steps:

First, the log-likelihood function of a single observation (Y, Z) at dose x with p parameters $\theta_1, \ldots, \theta_p$ is given by

$$l(y, z, x; \theta) = yz \log p_{11} + y(1-z) \log p_{10} + (1-y)z \log p_{01} + (1-y)(1-z) \log p_{00},$$

with $p_{00} = 1 - p_{11} - p_{10} - p_{01}$.

Second, the Fisher information matrix of a dose point x is

$$I(x; \theta) = \mathrm{E} \left[\frac{\partial l}{\partial \theta} \frac{\partial l}{\partial \theta^T} \right],$$

where

$$\frac{\partial l}{\partial \theta} = \begin{pmatrix} \frac{\partial l}{\partial \theta_1} \\ \vdots \\ \frac{\partial l}{\partial \theta_p} \end{pmatrix} = \begin{pmatrix} \frac{\partial p_{11}}{\partial \theta_1} & \frac{\partial p_{10}}{\partial \theta_1} & \frac{\partial p_{01}}{\partial \theta_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial p_{11}}{\partial \theta_p} & \frac{\partial p_{10}}{\partial \theta_p} & \frac{\partial p_{01}}{\partial \theta_p} \end{pmatrix} \begin{pmatrix} \frac{yz}{p_{11}} - \frac{(1-y)(1-z)}{p_{00}} \\ \frac{y(1-z)}{p_{10}} - \frac{(1-y)(1-z)}{p_{00}} \\ \frac{(1-y)z}{p_{11}} - \frac{(1-y)(1-z)}{p_{00}} \end{pmatrix} = \frac{\partial P}{\partial \theta} \cdot S_{yz}.$$

Thus,

$$I(x; \theta) = \mathrm{E} \left[\frac{\partial l}{\partial \theta} \frac{\partial l}{\partial \theta^T} \right] = \frac{\partial P}{\partial \theta} \mathrm{E}[S_{yz} S_{yz}^T] \frac{\partial P}{\partial \theta^T}.$$

Note that

$$\begin{aligned}
&\mathrm{E}[y(1-y)] = \mathrm{E}[z(1-z)] = 0, \quad \mathrm{E}(y^2(1-z)^2) = p_{10},\\
&\mathrm{E}((1-y)^2(1-z)^2) = p_{00}, \qquad \mathrm{E}((1-y)^2 z^2) = p_{01},\\
&\mathrm{E}(y^2 z^2) = p_{11}.
\end{aligned}$$

So

$$\begin{aligned}
V = \mathrm{E}(S^T S) &= \begin{pmatrix} \frac{1}{p_{11}} + \frac{1}{p_{00}} & \frac{1}{p_{00}} & \frac{1}{p_{00}} \\ \frac{1}{p_{00}} & \frac{1}{p_{10}} + \frac{1}{p_{00}} & \frac{1}{p_{00}} \\ \frac{1}{p_{00}} & \frac{1}{p_{00}} & \frac{1}{p_{01}} + \frac{1}{p_{00}} \end{pmatrix} \\
&= \mathrm{diag}^{-1}(p_{11}, p_{10}, p_{01}) + \frac{1}{1 - p_{11} - p_{10} - p_{01}} J(3,3,1) \\
&= P^{-1} + \frac{1}{1 - p_{11} - p_{10} - p_{01}} J(3,3,1),
\end{aligned}$$

where $P = \mathrm{diag}(p_{11}, p_{10}, p_{01})$ and $J(3,3,1) = (1,1,1) \cdot (1,1,1)^T$.

In summary,

$$I(x;\theta) = \frac{\partial P}{\partial \theta} V \frac{\partial P}{\partial \theta^T}.$$

One can obtain $\partial P / \partial \theta$ when P is specified by a model.

Note that $V = (P - pp^T)^{-1}$. When we approaches boundary, p_{00} or p_{11} will go to zeros, which makes the matrix $P - pp^T$ singular. For computational consideration, if $P - pp^T$ is singular, we regularized it by the following:

$$V = \begin{cases} (P - pp^T)^{-1}, & \text{if} \quad |P - pp^T| \geq \lambda; \\ (P - pp^T + \lambda I)^{-1}, & \text{otherwise.} \end{cases}$$

The regularized matrix is invertible because for any $n \times n$ matrix B, which may not be

invertible, we have

$$(\lambda I_n + B)^{-1} = \frac{1}{\lambda} I_n - \frac{1}{\lambda^2 + \mathrm{Tr}(\lambda B)} B,$$

where I_n is the $n \times n$ identity matrix. In practice, we take $\lambda = 0.000001$.

1.4.3 Penalized D-optimal design

The traditional D-optimal design assumes that any dose has the same cost, which could be ethical cost or monetary cost. However, this may not hold in practice. Dragalin and Fedorov (2006) introduce a cost function that penalizes doses with high toxicity or low efficacy. The cost function is

$$c(x, \theta; C_E, C_T) = [p_{10}(x, \theta)]^{-C_E} [1 - p_{01}(x, \theta) - p_{11}(x, \theta)]^{-C_T}, \tag{1.9}$$

where C_E and C_T are user defined constants. We plot the cost c against toxicity probability p_T and efficacy and non-toxicity probability p_{10} in Figure 1.1. For example, when $p_T = 0.9$ and $p_{10} = 0.1$, $c = 100$. We see that doses with high toxicity and low efficacy are expensive.

The total cost of the design ξ (1.2) is

$$C(\xi, \theta) = \sum_i w_i c(x_i, \theta). \tag{1.10}$$

The penalized D-optimality criterion minimizes

$$\Phi'(\xi, \theta) = \det \left\{ \left[\frac{M(\xi, \theta)}{C(\xi, \theta)} \right]^{-1} \right\}. \tag{1.11}$$

As noted in Dragalin and Fedorov (2006), the penalized D-optimal problem can be reduced to the traditional D-optimal design by applying the transform:

$$w'(x) = \frac{c(x, \theta)}{C(\xi, \theta)} w(x) \quad , \quad I'(x, \theta) = \frac{I(x, \theta)}{c(x, \theta)} \quad \text{and} \quad M'(\xi, \theta) = \frac{M(\xi, \theta)}{C(\xi, \theta)}.$$

The directional derivative is

$$\phi(\xi, x) = Tr[M_\xi^{-1} M_x] - \frac{c(x,\theta)}{C(\xi,\theta)} p,$$

where $M_x = I(x;\theta)$, $M_\xi = M(\xi;\theta)$, and $Tr\left[M_\xi^{-1} M_x\right] - p$ is the directional derivative for local D-optimal design defined in (1.4).

The generalized equivalence theorem for penalized optimal designs can be modified as the equivalence of the following three conditions (Dragalin and Federov, 2006):

1. $\xi^* = \arg\min_\xi \det\{[M(\xi,\theta)/C(\xi,\theta)]^{-1}\}$;

2. $\xi^* = \arg\min_\xi \max_x \left\{Tr[M_\xi^{-1} M_x] - \frac{c(x,\theta)}{C(\xi,\theta)} p\right\}$;

3. $\frac{1}{c(x,\theta)} Tr[M_\xi^{-1} M_x] - \frac{p}{C(\xi,\theta)} \leq 0, \forall x \in \mathcal{X}$. Equality is attained at the design points.

1.4.4 Bivariate best intention design

Thall and Cook (2004) consider desirability contours as trade-offs between efficacy and toxicity to protect patients. The aim of their design is to assign patients to the dose with the most desirability. By definition in Section 1.3.3, this is a best intention design.

Specifically, Thall and Cook (2004) suppose the efficacy and toxicity probability pair at dose x is $\kappa(x) = (p_E(x), p_T(x))$ and the clinician chooses three probability pairs to have the same desirability. Then the three pairs define a equivalent class $\mathcal{C}$ from the quadratic family $\{\kappa = (p_E, p_T) : p_T = a + b/p_E + c/p_E^2, 0 \leq p_E, p_T \leq 1\}$. Note that the ideal pair is $\kappa_0 = (p_E, p_T) = (1, 0)$. For any pair κ_p, draw a straight line L_p that joins π_0 and π_p in the p_E and p_T plane. Denote the intersection point of L_p and $\mathcal{C}$ by π_q. Denote the Euclidean distance between any point κ and κ_0 by $\rho(\kappa)$. Then the desirability of κ_p is defined by

$$\delta(\kappa_p) = \frac{\rho(\kappa_q)}{\rho(\kappa_p)} - 1. \tag{1.12}$$

Therefore, points with greater desirability are closer to κ_0. Also, points outside the convex hull of $\mathcal{C}$ and κ_0 have negative desirability. Denote the maximum toxicity probability by α_E and the minimum efficacy probability by α_T. Then the tradeoff method proceeds at each step by assigning the patients to the most desirable dose x with the safety and efficacy constraint:

$$Pr(\{\theta : p_E(x;\theta_i) > 0.2\}) > 0.9,$$

$$Pr(\{\theta : p_T(x;\theta_i) < 0.4\}) > 0.9. \tag{1.13}$$

In a Monte Carlo simulation with N samples, the acceptance region at x is approximated by the following:

$$\frac{1}{N}\sum_{i=1}^{N} 1\{p_E(x;\theta_i) > 0.2\} > 0.9,$$

$$\frac{1}{N}\sum_{i=1}^{N} 1\{p_T(x;\theta_i) < 0.4\} > 0.9,$$

where $p_E(\theta_i)$ and $p_T(\theta_i)$ are the efficacy and toxicity probabilities given the ith simulated θ of N total simulations.

After assigning $n-1$ cohort to doses $1, 2, \cdots, K$, denote the assigned dose level by $\mathcal{B}$ and denote the highest assigned dose level by J. The nth cohort is assign as follows:

1. If the highest assigned dose level J is less than K, let $\mathcal{A}_0 = \mathcal{B} \cup \{J+1\}$. Otherwise, let $\mathcal{A}_0 = \mathcal{B}$.

2. Denote the set of dose that satisfies (1.13) by $\mathcal{A}$.

3. (a) If the acceptable dose space $\mathcal{A} \cap \mathcal{A}_0$ is empty, then terminate the trial.

   (b) If the acceptable dose space $\mathcal{A} \cap \mathcal{A}_0$ is not empty, then assign the nth patient to the dose with the highest desirability in the acceptable dose space.

1.5 Penalized Bayesian optimal design

The penalized Bayesian D-optimality solves the following equation,

$$\xi^* = \arg\min_{\xi} E_{\theta|\pi} \det \left\{ \left[\frac{M(\xi,\theta)}{C(\xi,\theta)} \right]^{-1} \right\}. \tag{1.14}$$

1.5.1 Extension of the generalized equivalence theorem

In local designs, the generalized equivalence theorem can be used to assess the optimality of a design. Chaloner and Larntz (1989) extended the generalized equivalence theorem to the Bayesian realm for two parameter logistic models. Haines, Perevozskaya and Rosenberger (2003) considered a sequential Bayesian design for two parameter logistic models and derived the generalized equivalence theorem for designs on a restricted space. Dragalin and Fedorov obtained the generalized equivalence theorem for local penalized designs. To propose a Bayesian optimal design and its corresponding sequential design, the generalized equivalence theorem should be extended to the penalized case. This is stated as the equivalence of following three conditions:

1. $\xi^* = \arg\min_{\xi} \mathrm{E}_{\theta|\pi} \log\det \left\{ [M(\xi;\theta)/C(\xi;\theta)]^{-1} \right\}$;

2. $\xi^* = \arg\min_{\xi} \max_x \mathrm{E}_{\theta|\pi} \left\{ Tr[M_\xi^{-1} M_x] - \frac{c(x,\theta)}{C(\xi,\theta)} p \right\}$;

3. $\mathrm{E}_{\theta|\pi} \left\{ Tr[M_\xi^{-1} M_x] - \frac{c(x,\theta)}{C(\xi,\theta)} p \right\} \leq 0, \forall x \in \mathcal{X}$. Equality is attained at the design points.

Gao and Rosenberger (2013) choose the penalty function to be $c(x;\theta_\pi)$, by substituting θ_π for θ, where θ_π is the mean of the posterior distribution π of θ. Then the penalized optimal design is given by:

$$\xi^* = \arg\min_{\xi} E_{\theta|\pi} \det \left\{ \left[\frac{M(\xi,\theta)}{C(\xi,\theta_\pi)} \right]^{-1} \right\}. \tag{1.15}$$

We note that the assumption of a common penalty function is not necessary for one to obtain the generalized equivalence theorem.

1.5.2 Prior elicitation

Suppose we can obtain a rough knowledge about the parameters from a previous study. Then we can choose an informative prior such as the Gaussian distribution when we are confident about parameter estimation, or we can choose less informative priors such as the uniform distribution when we are not that confident. Once the prior family is chosen, the question becomes how to choose the associated prior parameters. We note that this is a problem about the trade-off between estimation bias and estimation variation, which depends on how much confidence one puts on priors. For example, as shown by Lehmann and Casella (1998, pp 233) suppose $X_1, \cdots, X_n$ are identically and independently distributed as $N(\theta, \sigma^2)$ with known σ. The prior is $N(\mu, b^2)$. Then the posterior mean and variance is given by

$$\mathrm{E}(\theta|x) = \frac{n\bar{x}/\sigma^2 + \mu/b^2}{n/\sigma^2 + 1/b^2};$$

$$\mathrm{Var}(\theta|x) = \frac{1}{n/\sigma^2 + 1/b^2}.$$

Thus, a large b, which means low confidence about the prior, has low bias towards the prior μ but high variance. A small b, which means high confidence about the prior, has high bias but low variance. Because we can approximate the bivariate nonlinear problem with a normal distribution, we obtain the same conclusion for trials with large sample. Therefore, the problem of proposing a good prior reduces to the problem of trade-off between bias and variation.

It is a good idea to address this problem in the context of hierarchical Bayesian modeling, because a misspecified hyper-parameter has less influence on posterior distributions than misspecified prior parameters (Theorem 4.5.7, Lehmann and Casella, 1998). Using the

hierarchical approach, Thall and Cook (2004) propose a method to reduce the influence of a misspecified prior parameter. Specifically, suppose the ith component of the parameter θ is normally distributed as $\pi_i(\theta_i) = N(\eta_{1i}, \eta_{2i})$. Then the mean and variance of p_E and p_T are defined as:

$$\begin{aligned} m_{Ej} &= \int p_E(x_j;\theta)\pi(\theta)d\theta, \\ s^2_{Ej} &= \int (p_E(x_j;\theta) - m_{Ej})^2\pi(\theta)d\theta, \\ m_{Tj} &= \int p_T(x_j;\theta)\pi(\theta)d\theta, \\ s^2_{Tj} &= \int (p_T(x_j;\theta) - m_{Tj})^2\pi(\theta)d\theta. \end{aligned}$$

Then Thall and Cook (2004) assume $s^2_{Ej} = 0.5m_{Ej}(1 - m_{Ej})$ and $s^2_{Tj} = 0.5m_{Tj}(1 - m_{Tj})$ by a vague beta-distribution argument. They ask the clinician to specify the mean and standard deviation of p_E and p_T, denoted $\hat{m}_{yj}$ and $\hat{s}_{yj}$ for $y = E, T$ and $j = 1, \cdots, K$. Then the prior parameter η is chosen by:

$$\eta^* = \arg\min_{\eta} \sum_{y=E,T} \sum_{j=1,2,\ldots,K} [m_{yj}(\eta) - \hat{m}_{yj}]^2 + [s_{yj}(\eta) - \hat{s}_{yj}]^2 + c_\eta \sum_{m\neq n} (\eta_{2m} - \eta_{2n})^2,$$

where c_η is a fixed tuning parameter. Thall, Cook and Estey (2006) also propose two criteria for a good prior:

1. The prior should be sufficiently non-informative so that the data dominates the design.
2. The prior should be sufficiently informative to avoid ill-behavior at the beginning of the trial.

In practice, clinicians are not familiar with distribution of p_E and p_T. Therefore, we need additional information to obtain the distributions. Chaloner et al. (1993) include interactive

graphical tools to help the clinician select the prior parameters. Kadane and Wolfson (1996) suggest asking clinicians a set of questions and then determine the prior parameters. However, the mean and standard derivation of p_T and p_E do not map directly to the parameters of the bivariate models. Alternatively, Rosenberger et al. (2005) demonstrate that one can map the Γth response quantile μ to the parameter of a logistic model, through the identity $\mu = \alpha + \beta F^{-1}(\Gamma)$. They suggest asking clinicians about the 25th and 50th quantiles.

1.5.3 Safety criterion

Suppose we want to protect the patients by avoiding highly toxic doses. Then we need to introduce a penalty function $c(x, \theta)$, which should satisfy the following properties:

Safe $c(x, \theta)$ peaks at high doses with high toxicity;

Economic $c(x, \theta)$ peaks at low doses with low efficacy;

Flexible Clinician can configure the penalty towards their preferences;

Consolidated Statisticians can consolidate the penalty function in D-optimal designs.

The penalty function (1.9) introduced by Dragalin and Fedorov (2006) fulfills the safe, economic and consolidated criteria. Extension of this penalty to the Bayesian realm is included in Gao and Rosenberger (2012). The desirability function (1.12) introduced by Thall and Cook (2004) fulfills the safe, economic and flexible criterions.

1.5.4 Bayesian sequential design with penalty

Suppose our dose space $\mathcal{X}$ is discrete. We identify the optimal weights of doses based on the prior information. To do this, we implement the penalized method and substitute the mean of the priors as the local parameter. After the initial n_0 doses are allocated, we need a sequential procedure to approximate the optimality in (1.15). Instead of using the first order algorithm to maximize $\phi(\xi; \theta)$ (Dragalin and Fedorov, 2006), we use the

idea proposed by Haines, Perevozskaya and Rosenberger (2003). Specifically, suppose the posterior distribution π_n can be inferred through previous responses at step n. Then we assign a new participant to the dose with maximum increment of penalized information:

$$x_{n+1} = \arg\min_{x} \quad E_{\theta|\pi_n} \Phi \left(\frac{\sum_{i=1}^{n} I(x_i;\theta) + I(x;\theta)}{\sum_{i=1}^{n} c(x_i;\theta_{\pi_n}) + c(x;\theta_{\pi_n})} \right);$$

As the above steps are iteratively processed, the adaptive design is expected to converge to its corresponding optimum design.

Table 1.1: List of design criteria and their geometric interpretation.

Criterion	$\Phi(M(\xi))$	Geometric Interpretation
D	$\log\det(M^{-1})$	Volume
A	$Tr(M^{-1})$	Average axes length
E	$\lambda_{\min}^{(*)}$	Maximum axis length
c	$c^T M^{-1} c$	Span along c

$(*)$ Here $\lambda_{\min}$ is the minimum eigenvalue of $M(\xi)$.

Table 1.2: List of directional derivatives for local D, A, E and c optimal designs.

Criterion	$\left(\frac{\partial \Phi}{\partial M}\right)^T$	$\phi(\xi, \eta; \theta)$
D	$-M^{-1}$	$Tr(M(\xi)^{-1}M(\eta)) - p$
A	$-M^{-2}$	$Tr(M(\xi)^{-2}M(\eta)) - Tr(M^{-1}(\xi))$
E	$rr^{T(*)}$	$r^T M(\eta))r + \lambda_{\min}$
c	$-M^{-1}cc^T M^{-1}$	$c^T M^{-1}(\xi)M(\eta)M^{-1}(\xi)c - c^T M^{-1}(\xi)c$

$(*)$ Here r is the eigenvector corresponding to $\lambda_{\min}$.

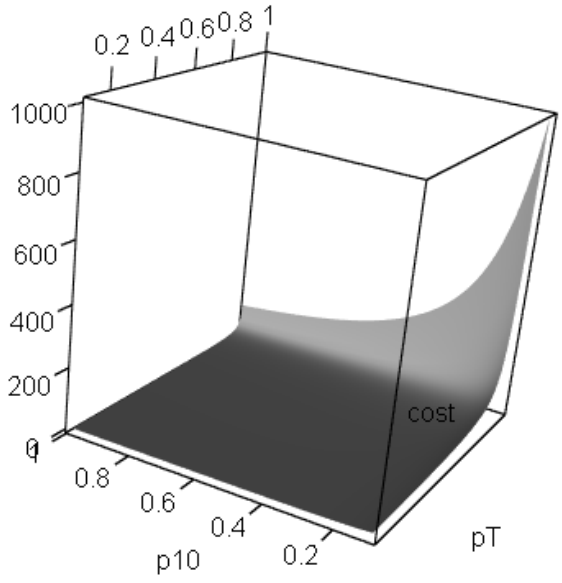

Figure 1.1: Cost c as a function of toxicity probability p_T and efficacy and non-toxicity probability p_{10}. Doses with high toxicity and low efficacy are expensive.

Chapter 2: Bayesian Sequential Designs

In a good sequential design, we want the parameter estimator $\hat{\theta}_n$ to converge to the true parameter θ. We also want the design ξ_n to converge to the optimal design $\xi(\theta)$. These two goals are interactive: on the one hand, assignment of patients depends on the current estimate of θ; on the other hand, design assignment determines estimation efficiency. The interaction makes it difficult to determine the convergence. In a bad design, patients may be always assigned to a single dose and therefore generate bad estimates, which in turn encourages single dose assignment. To address this issue, several researchers have examined the conditions to guarantee convergence of parameter estimates and designs for various models with different optimality. Specifically, Lai (1994), Chaudhuri and Mykland (1993, 1995) and Pronzato (2010a) considered the MLE in nonlinear models with D-optimality. Hu (1998) and Roy, Ghosal and Rosenberger (2009) considered Bayesian estimates in nonlinear model and two parameter logistic model with D-optimality respectively. Pronzato (2010b) considered the MLE in nonlinear models with penalized D-optimality. The convergence of the Bayesian sequential design for the Gumbel model is still unknown.

Despite the theoretical difficulties, questions related to how to determine an appropriate procedure for prior elicitation in sequential designs are very important. A unique question that arises in the context of bivariate response models is the selection of an appropriate prior for the correlation parameter. In this chapter, we address both the convergence and prior selection problems using simulation. Specifically, we explore the impact of the correlation parameter on various designs. Then we propose a prior elicitation procedure for bivariate models. Last, we examine the impact of different misspecified priors on sequential procedures.

2.1 Correlation

In this section, we explore the impact of correlation on the designs, especially on the dose-response relationship, information measures and optimal designs.

2.1.1 Dose-response relationship

Recall that we assume toxicity and efficacy responses, (Y, Z), are binary and follow a Gumbel model with linear effects: $x_E = \beta_E(x - \alpha_E)$ and $x_T = \beta_T(x - \alpha_T)$. Thall and Cook (2004) considered a quadratic efficacy effect. The cell probabilities are given by

$$\begin{aligned}
p_{11} &= \Pr(Y = 1, Z = 1|D = x; \theta) = G(x_E, x_T), \\
p_{10} &= \Pr(Y = 1, Z = 0|D = x; \theta) = F(x_E) - G(x_E, x_T), \\
p_{01} &= \Pr(Y = 0, Z = 1|D = x; \theta) = F(x_T) - G(x_E, x_T), \\
p_{00} &= \Pr(Y = 0, Z = 0|D = x; \theta) = 1 - p_{11} - p_{10} - p_{01},
\end{aligned}$$

where G is the Gumbel function:

$$G(y, z) = F(y) \cdot F(z) \cdot \{1 + \rho[1 - F(y)][1 - F(z)]\},$$

and $F(y) = 1/(1 + e^{-y})$.

Without loss of generality, we examine four scenarios with different toxicity and efficacy effects. Although the 50th percentile of efficacy are assumed to be smaller than those of toxicity in the first three scenarios, slopes are different across scenarios. The 50th percentile of efficacy is assumed to be greater than that of toxicity in scenario 4. Efficacy probabilities increase faster than toxicity probabilities in scenario 1 and 2. Toxicity probabilities increase faster than efficacy probabilities in scenario 3 and 4. We standardize the parameters by fixing the 50th percentile of efficacy and toxicity at -2 and 2 in the first three scenarios and at 2 and -2 in the last scenario. Figure 2.1 illustrates the toxicity, efficacy and "efficacy and no toxicity" probabilities for each scenario. The parameters $(\alpha_E, \beta_E, \alpha_T, \beta_T, \rho)$ are

$(-2, 1, 2, 0.5, 0.5)$, $(-2, 0.5, 2, 0.5, 0.5)$, $(-2, 0.2, 2, 0.5, 0.5)$ and $(-2, 0.2, -2, 0.2, 0.5)$ respectively. All other scenarios can be represented similarly by applying a linear transformation to the dose space.

We examine the impact of correlation on the dose-response relationship visually. Note that the toxicity and efficacy probabilities, $p_T = p_{11} + p_{01}$ and $p_E = p_{11} + p_{10}$, are not dependent on the correlation parameter ρ. The correlation parameter only determines the individual cell probabilities p_{11}, p_{10}, p_{01} and p_{00}. We plot p_{10}, the efficacy and no toxicity probability, for each scenario with different correlation parameters in Figure 2.2. We see there is not much difference among the p_{10}'s in all scenarios. The maximum difference of p_{10}'s is about 0.15 in scenario 3.

An interesting question is how many samples are required to obtain a precise estimate for ρ. We can obtain a crude estimate by treating p_{10} as an unknown parameter in a binomial distribution. Then to distinguish the correlation parameter -1 from 1, we need a sample of $(1.96/0.15)^2 \times 0.35 \times 0.65$, roughly 40 patients, given that we know the optimal dose for distinction is at -1. Therefore, we expect the estimation of the correlation parameter to be poor with a sample size less than 40.

2.1.2 Optimal designs

How much influence the correlation parameter has on optimal designs is an important problem. To explore this problem, we introduce the following D_A-optimal design, and compare the D-optimal designs to D_A-optimal designs.

The D_A criterion maximize the information about the toxicity and efficacy parameters with no intention to learn the correlation parameter. The D_A design minimize the following:

$$\Phi(\xi, \theta) = \log \det \left\{ A^T M^{-1}(\xi; \theta) A \right\}, \tag{2.1}$$

where

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

The directional derivative is

$$\phi(x, \xi; \theta) = Tr\left[(A^T M^{-1} A)^{-1} \cdot A^T M^{-1} I(x; \theta) M^{-1} A\right] - 4.$$

We summarize the D and D_A-optimal designs for four scenarios in Tables 2.1 – 2.4. For each parameter, we calculate the D_A information of the D-optimal design, and compute its relative efficiency with respect to the D_A-optimal design. We can see that the information loss due to the correlation parameter in D-optimal designs is fairly low (about 4%). This means D-optimal designs are almost as efficient as the D_A-optimal designs for the Gumbel model. Therefore, we can use D-optimal criterion alone and obtain near efficient estimation about the toxicity and efficacy parameters.

All optimal designs are supported on four points. The design points in D_A-optimal designs are more stretched out than D-optimal designs in scenario 1, 2 and 4, while they are more squeezed in D_A-optimal designs than D-optimal designs in scenario 3. The points do not change monotonically as ρ increases. For example, the points are stretched out when ρ gets to -0.5, and squeezed as ρ increases further. The middle points gains more weights as ρ increases. The D_A-optimal designs have even weights when ρ is -1. The D_A designs are similar across ρ's in scenario 1, 2 and 4.

2.2 Prior elicitation

There are two different scenarios for prior elicitation. In the first scenario, a group of experts agree on the range of reasonable percentiles of toxicity and efficacy probabilities. In the second scenario, each expert provides his/her own opinion. In this dissertation, we only consider the first scenario. The second is considered in Thall and Cook (2004).

Assume we want to select priors based on a single opinion. As we see in Section 2.1.1, the efficacy probability p_E and toxicity probability p_T are determined solely by the efficacy and toxicity parameters $\alpha_E, \alpha_T, \beta_E$ and β_T. The correlation parameter determines the efficacy and no toxicity probability p_{10}, which is shown in Figure 2.1.

It is easier to obtain information about toxicity and efficacy parameters. Suppose a Gumbel model for the dose-response relationship with linear toxicity and efficacy effects. Using Rosenberger et al.'s (2005) method, we ask the expert about the 50th and 25th percentiles of efficacy and toxicity, denoted by ED50, ED25, LD50 and LD25. Then the toxicity and efficacy parameters can be obtained by:

$$\alpha_E = ED50, \quad \beta_E = \frac{G50 - G25}{ED50 - ED25},$$

$$\alpha_T = LD50, \quad \beta_T = \frac{G50 - G25}{LD50 - LD25},$$

where G25 and G50 are the 25th and 50th percentiles of the standard logistic distribution. Using normal priors, these quantities are served as mean of the prior distributions of $\alpha_E, \beta_E, \alpha_T$ and β_T. The standard deviation of the priors are then decided by the experts preferably with the aid of an interactive graphical tool such as that used by Chaloner, et al. (1993), or by estimating the prior effective sample size (Morita et al, 2008). The probabilities of toxicity and efficacy are assumed to be monotone, thus the prior for β_E and β_T are restricted to be positive.

The correlation parameter cannot be obtained easily. We illustrate this in a special case shown in Figure 2.2. First, the mode of the efficacy and no toxicity probabilities

cannot be used to differentiate correlations. For example, the modes are the same across ρ's in scenarios 1, 2, and 4. Second, given the knowledge of the true toxicity and efficacy parameters and the knowledge of the optimal dose, it requires 40 patients to obtain a clear cut between $\rho = 1$ and $\rho = -1$ (Section 2.1.1). The expert is not likely to possess such accurate information. Therefore, we propose to use a non-informative prior for the correlation parameter such as the uniform distribution on $(-1, 1)$. In some cases, the expert is confident about whether the correlation is positive or negative based on clinical reasons. Then we can use a beta or truncated normal prior that peaks at -0.5 or 0.5 accordingly. The density function of a beta distribution defined on $[-1, 1]$ is proportional to $(1-x)^{\alpha-1}(1+x)^{\beta-1}$ with mean $(\alpha-\beta)/(\alpha+\beta)$ and mode $(\alpha-\beta)/(\alpha+\beta-2)$.

2.3 Sequential design

We conduct Monte Carlo simulation to examine the operating characteristics of the sequential designs as well as sensitivity to misspecified priors. We examine the aforementioned four scenarios. We propose to use a procedure based on the MLE as a benchmark to test our Bayesian sequential procedures. Also, we will use different priors for each scenario. We want to investigate the influence of misspecified priors on the designs and parameter estimation.

2.3.1 Design specification

We select four doses in the sequential procedure. The doses are selected with similar toxicity and efficacy probabilities as in Thall and Cook (2004). For each scenario, we compute the optimal designs that are restricted to four doses. The optimal weights and toxicity and efficacy probabilities for each dose are included in Table 2.5.

In the simulation, we consider five different priors listed in Table 3.11. Also, we consider a different and long dose range $(-3, -1, 1, 3)$. We want to see how the designs converge under different priors and dose ranges.

2.3.2 MCMC algorithm

Given the priors, there is no closed form expression for the posterior distribution,

$$\pi(\theta|\mathcal{D}) \propto L(\mathcal{D}|\theta)\pi(\theta).$$

Therefore for each assignment, we use the Gibbs sampler (Gelfand and Smith, 1990) to generate samples $(\theta^{(1)}, \ldots, \theta^{(K)})$ from the posterior distribution. Because the conditional distributions in Gibbs sampling do not come from any conjugate family, we use Metropolis-Hastings steps to obtain samples from the conditional distributions. Specifically, we use Gaussian univariate proposals (Roberts, Gelman, Gilks, 1997) in these steps. For each parameter, the standard deviation of the Gaussian proposal is tuned with an acceptance rate about 30%. After visual inspection with a few different initial values, we decided to use 10000 MCMC samples after 2000 burn-ins. Then the convergence is visually checked on a preselected sample.

Approximating integrals by averaging over the samples, we assign the $(n+1)$th participant to the dose:

$$x^{(n+1)} = \arg \max_{x \in \{x_1, \ldots x_m\}} \quad \sum_{k=1}^{K} \Phi(\xi \vee x_i, \theta^{(k)}).$$

We also monitor the maximum a posteriori (MAP) estimator of θ for the sample size of n,

$$\hat{\theta}_n = \arg \max_{\theta^{(k)}} L(\mathcal{D}_n|\theta^{(k)})\pi(\theta^{(k)}).$$

We will investigate these quantities in simulations in the next section.

2.3.3 Convergence measures

We define a few quantities to evaluate the performance of parameter estimation. This can be done from two different approaches. In the first approach, the estimates for each parameter

can be accessed separately. In the second approach, parameter estimation can be assessed using a scalar function of parameters. For example, the function can be deviance or the estimated optimal dose in the next chapter.

In the first approach, we denote the MAP estimate of θ at the jth step in the ith trial by $\hat{\theta}_{i,j}$. Denote the mean estimate at the jth step by $\bar{\theta}_j$. Then there are two interesting quantities to investigate. The first quantity is called "relative ratio", used to measure the bias of Bayesian parameter estimator, and defined by

$$RR_j = \left| \frac{\bar{\theta}_j}{\theta} \right|.$$

Also, the mean squared error (MSE) at the jth step is defined by

$$MSE(\hat{\theta}_j) = \frac{\sum_{i=1}^{1000} \left(\hat{\theta}_{i,j} - \theta \right)^2}{1000}.$$

Similar to Manukyan (2009), we define the relative root mean squared error (RRMSE) at the jth step to be

$$RRMSE(\hat{\theta}_j) = \frac{\sqrt{MSE(\hat{\theta}_j)}}{\theta}.$$

In the second approach, assume n patients are assigned and the response is $\boldsymbol{Y}_n$. Let $\hat{\boldsymbol{\theta}}_n$ be the parameter estimate. We define the deviance to be

$$D(\boldsymbol{Y}_n, \hat{\boldsymbol{\theta}}_n) = -2 \left[\log \left(L(\boldsymbol{Y}_n | \boldsymbol{\theta}) \right) - \log \left(L(\boldsymbol{Y}_n | \hat{\boldsymbol{\theta}}_n) \right) \right],$$

where $\boldsymbol{\theta}$ is the true parameter.

Note that a design converges to the optimal design when its dose allocation weights converge to the optimal weights. To evaluate the convergence of a design, we define a

quantity to evaluate the distance between a design to the corresponding optimum design. Suppose the design has a weight vector $\boldsymbol{w}$ and the optimal weight vector is $\boldsymbol{w}^*$. Then we define

$$d(\boldsymbol{w}, \boldsymbol{w}^*) = \sum_{i=1}^{4} |w_i - w_i^*|.$$

The reason we use L^1 norm is that $||\mathbf{w}||_{L^1} = 1$ for any $\boldsymbol{w}$.

In addition to the L^1 distance, we consider the information efficiency of a design ξ, given by

$$\text{eff}(\xi) = \exp\left\{\frac{\Phi(\xi^*, \theta) - \Phi(\xi, \theta)}{p}\right\},$$

where p is the number of parameter and ξ^* is the optimal design.

2.3.4 MLE benchmark

We want a benchmark with which we can compare our Bayesian sequential procedure. Manukyan (2009) conducted simulations and showed that sequential designs with Bayesian estimates outperform those using the MLE. In his MLE setting, assignments are determined by criteria using the MLE. Assuming the optimal allocation is known, we propose to assign patients to attain the optimal allocation and use the MLE to estimate the parameter. Therefore, the design is optimal regardless of the parameter estimates. We want to see how efficient the parameter estimates can be given the optimal allocation. This should set a upper bound for any MLE-type sequential designs because it requires time for these designs to find the optimal allocation. Then we use this as a benchmark to evaluate how the Bayesian sequential procedure performs.

The MLE procedure starts by assigning 5 patients to the lowest dose. Assume the optimal allocation is $\{w_1^*, \ldots, w_n^*\}$. The procedure proceeds iteratively as follow:

1. At step j, we have assigned j patients. Denote patient allocation weight by $\{w_1, \ldots, w_n\}$.

2. Determine next assignment x_{j+1} by $\arg\max_i\{w_i^* - w_i; i = 1, \ldots, n\}$.

3. Assign the $j+1$th patient to the x_{j+1}th dose, and obtain response y_{j+1}.

4. Obtain the MLE of θ given the data $(x_1, y_1), \ldots, (x_{j+1}, y_{j+1})$.

We compare the the MLE method with the Bayesian method with prior 1. We conduct 1000 trials and record relevant quantities. The graphical comparison of deviance is included in Figures 2.3. The MLE method performs uniformly worse than the Bayesian method because the deviance of the former is far from zero. The RR and RRMSE measures are included in Tables 2.7 – 2.10. Due to large RRMSE's, the MLE method produces unstable estimates at the beginning of a trial. Also, parameters α_E and α_T in the Bayesian method converge faster than those in the MLE method. Although the estimates of ρ in the MLE methods are closer to the true value, the MSE of the MLE method is greater than the Bayesian method. Overall, the Bayesian method performs better than the MLE method.

2.3.5 Bayesian sequential designs with different dose ranges

Dose range is another factor impacting parameter estimation. Wide dose ranges may lead to good parameter estimates. Here we compare designs operated on the narrow range $\{-3, -2, -1, 0\}$ with those operated on the wide range $\{-3, -1, 1, 3\}$. We use prior 1 for both dose ranges.

First, we look at individual parameter estimates. Tables 2.7 – 2.10 list the RR and MSE measures for the two methods. When sample size is 25, designs on the wide dose range perform worse than those on the narrow range for β_T. The former gets better than the latter when the sample size increases from 50 to 200. In scenarios 1–3, estimates for α_T from designs on the wide range are better than those from designs on the narrow range. In scenarios 2–4, estimates for β_E from the former are better than those from the latter. Both designs perform equally well in estimating α_E. Second, we evaluate the overall parameter estimation based on the deviance included in Figure 2.4. Designs operated on the narrow range perform uniformly better than those on the wide range. This contradicts

what we observe from the individual parameter estimates. Thus, there is a discrepancy between individual parameter estimates and overall parameter estimates. All the parameter estimates demonstrate convergence as sample size increases.

In addition to parameter estimation, we examine the convergence of dose allocation. Both the L^1 distance and design efficiency are included in Figure 2.5 for scenarios 1 and 4. Using the L^1 distances, we see that designs on the wide range performs better in scenario 1. It performs worse in scenario 4. We confirm these observations by examining design efficiency measures. Thus the deviance measure is consistent with the design efficiency measure rather than measures based on an individual parameter.

We summarize that the design operated on long dose ranges performs better by producing good individual parameter estimates. However, the advantage does not translate into better dose allocation or overall learning about the dose-response relationship, especially when the prior is misspecified in scenario 4.

2.3.6 Sequential design with different priors for correlation

By investigating designs with different priors (Table 3.11), we hope to answer three different questions: whether the design performs better with informative priors; whether a uniform prior is better than a normal prior; whether a misspecified correlation prior impacts the design.

Question 1: Does the design perform better with informative priors?

The deviances are plotted in Figure 2.6. Note that prior 1 is a normal prior and prior 2 is a more informative normal prior with smaller standard deviation. The design with prior 2 produces smaller deviances and performs better than that with prior 1 in scenario 1. However, when sample size is 50, the design with prior 2 performs worse than that with prior 1 in scenario 4. Both design perform equally well as sample size increases. Prior 4 is a uniform prior and prior 5 is a more informative uniform prior with smaller support. The former performs worse than the latter in scenario 1. However, the former performs much

better in scenario 4 when priors are misspecified.

We examine individual parameter estimates in Tables 2.7 – 2.10. Comparing priors 1 with 2, we see that estimates of β_T start poorly and converge faster in designs with informative prior. However, estimates of β_T are uniformly worse in designs with informative prior in scenario 4. Comparing priors 4 with 5 in scenario 1, we see that designs with informative prior perform better. In scenario 4 designs with informative prior perform better in estimating α_E and α_T, but perform worse in estimating β_E and β_T. Thus, we conclude that informative priors produce better parameter estimates when the priors are slightly misspecified. However, they can perform badly when the priors are severely misspecified.

We examine dose allocation measures. Comparing priors 1 with 2 in scenarios 1 and 4, the former produces dose allocations that are closer to the optimal dose allocation and thus performs better than the latter. The former also converges with faster speed and smaller variations. Comparing priors 4 and 5 in scenarios 1 and 4, the latter performs better with faster convergence rate and smaller dose allocation variations. Thus we observe a discrepancy between the overall parameter estimates measured by deviances and the dose allocation measures. Although the variation within the overall parameter estimates depends on information of prior and dose allocation variations, it seems that the former is more essential.

Question 2: Is a uniform prior better than a normal prior?

We compare uniform and normal priors using individual parameter estimates in Tables 2.7 – 2.10. The uniform prior 4 produces larger RRMSE's and thus performs worst than the normal prior 1 in estimating α_T and β_T in both scenarios 1 and 4. Priors 2 and 5 performs equally well in scenario 1. But normal prior 2 performs better in estimating β_E and β_T while uniform prior 4 performs better in estimating α_E and α_T. We see that one type of prior is not uniformly better than the other based on individual parameter estimates. We turn to the overall deviance measure.

We compare normal prior 1 with uniform prior 4 in Figure 2.6. In scenarios 1 and

4, the former performs uniformly better than the latter due to smaller deviances. Both priors have similar variation. Comparing normal prior 2 with uniform prior 5 in scenarios 1 and 4, the former performs uniformly better than the latter with smaller deviance and smaller variation. Therefore, we conclude that the normal prior performs better in overall parameter estimation.

We examine dose allocation measures in Figure 2.6. Comparing prior 1 with prior 4, we see that the normal prior performs slightly better with near optimal dose allocation. Comparing prior 2 with prior 5, the uniform prior performs much better than the normal prior with near optimal dose allocations and smaller variations. This is due to the more restricted parameter support of the uniform priors. We summarize that uniform priors produce stable dose allocation. This is desirable because it leads to better control of the design. However, normal priors with less informative prior may produce better parameter estimates and learning about the dose-response relationship.

Question 3: How does a misspecified correlation prior impact the design?

We compare prior 1 with prior 3. Note that prior 1 has correctly specified correlation while prior 3 has misspecified correlation. First, we examine individual parameter estimation measures in Tables 2.7 – 2.10. The two designs perform equally well in estimating α_E, α_T, β_E and β_T. Prior 1 outperforms prior 3 in estimating ρ. Second, we look at the overall parameter estimation measure, i.e., the deviance in 2.6. Prior 3 does not show convergence in both scenarios 1 and 4. Prior 1 performs better with smaller variation.

According to the L^1 distances in Figure 2.6, prior 1 produces better design that is close to the optimal dose allocation in scenario 1 and 4. However, the design with prior 3 does not converge to the optimal dose allocation in scenario 1. It appears to converge in scenario 4 with a slow rate. We obtain similar conclusion using the design efficiency measure. We summarize that design with misspecified correlation may have convergence issue form both the parameter estimation and the dose allocation perspectives.

2.4 Conclusion

In this chapter, we investigate optimal designs for bivariate Gumbel models and discuss the impact of correlation on the information criterion. To conduct Bayesian sequential designs, we discuss some practical problems such as prior elicitation. To assess the convergence of parameter estimation, we propose a procedure based on MLE and compare it with Bayesian procedures.

We examine the convergence of Bayesian sequential designs using simulation. Almost all designs show convergence with different convergence rates. It is interesting to see a discrepancy between individual parameter estimates and overall parameter estimates. There is also discrepancy between dose allocation and parameter estimates. This makes design comparison a challenging task. By comparing MLE and Bayesian procedures, we learn that Bayesian methods perform better by producing more accurate parameter estimates. By conducting designs on wide and narrow dose ranges, we see that the former produces better estimates for individual parameters at the cost of less stable dose allocation. By conducting designs with informative and less informative priors, we see that less informative priors produce better overall parameter estimates. Comparing normal priors to uniform priors, we see that normal priors are good at producing accurate parameter estimates while uniform priors may produce stable dose allocation. Finally, we observe non-convergence issues in some cases when the correlation is misspecified. Because of the inconsistency among all the measures, we strongly recommend the use of simulations to assess operating characteristics of adaptive designs and select informative or noninformative priors.

Table 2.1: Optimal designs for scenario 1, Gumbel model.

Scenario	ρ	D-optimal design	D_A-optimal design	Efficiency
1	-1	$\begin{pmatrix} -3.41 & -0.70 & 5.24 & 11.54 \\ 0.32 & 0.28 & 0.20 & 0.20 \end{pmatrix}$	$\begin{pmatrix} -3.59 & -0.54 & 5.29 & 11.55 \\ 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix}$	0.97
1	-0.5	$\begin{pmatrix} -3.23 & -0.72 & 5.07 & 11.50 \\ 0.27 & 0.33 & 0.20 & 0.20 \end{pmatrix}$	$\begin{pmatrix} -3.59 & -0.54 & 5.10 & 11.51 \\ 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix}$	0.96
1	0	$\begin{pmatrix} -3.35 & -0.78 & 4.83 & 11.45 \\ 0.24 & 0.36 & 0.20 & 0.20 \end{pmatrix}$	$\begin{pmatrix} -3.59 & -0.55 & 4.85 & 11.46 \\ 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix}$	0.96
1	0.5	$\begin{pmatrix} -3.44 & -0.77 & 4.5 & 11.39 \\ 0.22 & 0.38 & 0.20 & 0.20 \end{pmatrix}$	$\begin{pmatrix} -3.58 & -0.57 & 4.49 & 11.39 \\ 0.25 & 0.25 & 0.26 & 0.25 \end{pmatrix}$	0.96
1	1	$\begin{pmatrix} -3.47 & -0.69 & 3.97 & 11.30 \\ 0.21 & 0.38 & 0.21 & 0.20 \end{pmatrix}$	$\begin{pmatrix} -3.57 & -0.64 & 3.84 & 11.29 \\ 0.25 & 0.24 & 0.27 & 0.25 \end{pmatrix}$	0.95

Table 2.2: Optimal designs for scenario 2, Gumbel model.

Scenario	ρ	D-optimal design				D_A-optimal design				Efficiency
2	-1	$\begin{pmatrix} -5.89 & -1.00 & 2.59 & 7.44 \\ 0.24 & 0.28 & 0.26 & 0.22 \end{pmatrix}$				$\begin{pmatrix} -6.29 & -1.08 & 2.76 & 7.77 \\ 0.27 & 0.24 & 0.24 & 0.25 \end{pmatrix}$				0.99
2	-0.5	$\begin{pmatrix} -6.23 & -1.27 & 2.54 & 7.90 \\ 0.20 & 0.31 & 0.29 & 0.20 \end{pmatrix}$				$\begin{pmatrix} -6.50 & -1.47 & 2.85 & 8.09 \\ 0.24 & 0.25 & 0.26 & 0.25 \end{pmatrix}$				0.97
2	0	$\begin{pmatrix} -6.22 & -1.23 & 2.51 & 7.91 \\ 0.18 & 0.32 & 0.30 & 0.20 \end{pmatrix}$				$\begin{pmatrix} -6.45 & -1.56 & 2.84 & 8.08 \\ 0.23 & 0.25 & 0.28 & 0.25 \end{pmatrix}$				0.97
2	0.5	$\begin{pmatrix} -6.06 & -1.15 & 2.48 & 7.90 \\ 0.17 & 0.33 & 0.30 & 0.20 \end{pmatrix}$				$\begin{pmatrix} -6.27 & -1.58 & 2.83 & 8.05 \\ 0.22 & 0.24 & 0.30 & 0.25 \end{pmatrix}$				0.96
2	1	$\begin{pmatrix} -5.80 & -1.02 & 2.40 & 7.84 \\ 0.17 & 0.32 & 0.31 & 0.20 \end{pmatrix}$				$\begin{pmatrix} -5.98 & -1.52 & 2.76 & 7.98 \\ 0.20 & 0.24 & 0.31 & 0.25 \end{pmatrix}$				0.96

Table 2.3: Optimal designs for scenario 3, Gumbel model.

Scenario	ρ	D-optimal design	D_A-optimal design	Efficiency
3	-1	$\begin{pmatrix} -9.91 & -0.48 & 3.93 & 10.55 \\ 0.19 & 0.29 & 0.33 & 0.19 \end{pmatrix}$	$\begin{pmatrix} -9.88 & -0.60 & 4.03 & 11.14 \\ 0.24 & 0.26 & 0.27 & 0.23 \end{pmatrix}$	0.98
3	-0.5	$\begin{pmatrix} -10.32 & -0.86 & 3.63 & 10.76 \\ 0.18 & 0.31 & 0.35 & 0.17 \end{pmatrix}$	$\begin{pmatrix} -10.15 & -1.09 & 3.89 & 11.00 \\ 0.23 & 0.27 & 0.29 & 0.22 \end{pmatrix}$	0.98
3	0	$\begin{pmatrix} -11.05 & -1.47 & 3.35 & 10.80 \\ 0.16 & 0.32 & 0.37 & 0.15 \end{pmatrix}$	$\begin{pmatrix} -10.91 & -1.95 & 3.64 & 10.97 \\ 0.20 & 0.30 & 0.31 & 0.20 \end{pmatrix}$	0.97
3	0.5	$\begin{pmatrix} -11.67 & -1.98 & 3.11 & 11.01 \\ 0.13 & 0.34 & 0.38 & 0.15 \end{pmatrix}$	$\begin{pmatrix} -11.56 & -2.62 & 3.45 & 11.16 \\ 0.16 & 0.33 & 0.32 & 0.19 \end{pmatrix}$	0.97
3	1	$\begin{pmatrix} -11.97 & -2.22 & 2.92 & 11.12 \\ 0.12 & 0.34 & 0.39 & 0.15 \end{pmatrix}$	$\begin{pmatrix} -11.83 & -2.93 & 3.30 & 11.22 \\ 0.14 & 0.34 & 0.32 & 0.20 \end{pmatrix}$	0.96

Table 2.4: Optimal designs for scenario 4, Gumbel model.

Scenario	ρ	D-optimal design	D_A-optimal design	Efficiency
4	-1	$\begin{pmatrix} -14.78 & -2.96 & 4.85 & 16.58 \\ 0.24 & 0.28 & 0.26 & 0.23 \end{pmatrix}$	$\begin{pmatrix} -15.72 & -3.15 & 5.17 & 17.43 \\ 0.27 & 0.24 & 0.231 & 0.26 \end{pmatrix}$	0.99
4	-0.5	$\begin{pmatrix} -16.48 & -3.69 & 5.06 & 18.45 \\ 0.19 & 0.31 & 0.30 & 0.20 \end{pmatrix}$	$\begin{pmatrix} -16.94 & -4.25 & 5.77 & 18.88 \\ 0.24 & 0.25 & 0.27 & 0.24 \end{pmatrix}$	0.97
4	0	$\begin{pmatrix} -16.38 & -3.56 & 5.14 & 18.7 \\ 0.18 & 0.32 & 0.31 & 0.19 \end{pmatrix}$	$\begin{pmatrix} -16.85 & -4.42 & 5.99 & 19.09 \\ 0.22 & 0.26 & 0.28 & 0.24 \end{pmatrix}$	0.97
4	0.5	$\begin{pmatrix} -15.88 & -3.37 & 5.29 & 18.92 \\ 0.17 & 0.33 & 0.32 & 0.19 \end{pmatrix}$	$\begin{pmatrix} -16.42 & -4.43 & 6.2 & 19.28 \\ 0.21 & 0.26 & 0.30 & 0.24 \end{pmatrix}$	0.96
4	1	$\begin{pmatrix} -15.01 & -3.1 & 5.42 & 19.12 \\ 0.15 & 0.33 & 0.33 & 0.19 \end{pmatrix}$	$\begin{pmatrix} -15.66 & -4.31 & 6.34 & 19.47 \\ 0.19 & 0.26 & 0.32 & 0.23 \end{pmatrix}$	0.96

Table 2.5: Design specification for four scenarios, including dose level, toxicity and efficacy probabilities, D-optimal weights and optimal Φ.

θ	Φ	Dose	(p_E, p_T)	weight
Scenario 1				
$(-2, 1, 2, 0.5, 0.5)$	16.03	**-3**	(0.27, 0.08)	0.25
		-2	(0.50, 0.12)	0.19
		-1	(0.73, 0.18)	0.22
		0	(0.88, 0.27)	0.34
Scenario 2				
$(-2, 0.5, 2, 0.5, 0.5)$	19.51	**-3**	(0.38, 0.08)	0.31
		-2	(0.50, 0.12)	0.29
		-1	(0.62, 0.18)	0.00
		0	(0.73, 0.27)	0.40
Scenario 3				
$(-2, 0.2, 2, 0.5, 0.5)$	20.69	**-3**	(0.45, 0.08)	0.33
		-2	(0.50, 0.12)	0.20
		-1	(0.55, 0.18)	0.09
		0	(0.60, 0.27)	0.38
Scenario 4				
$(2, 0.2, -2, 0.2, 0.5)$	28.49	**-3**	(0.45, 0.27)	0.32
		-2	(0.50, 0.31)	0.29
		-1	(0.55, 0.35)	0.00
		0	(0.60, 0.40)	0.39

Table 2.6: Prior distributions used in the simulation. The prior distributions for β_E and β_T are assumed to be positive and they are not normalized in the table.

Prior	α_E	β_E	α_T	β_T	ρ
1	$N(-2.5, 1)$	$N(1, 1)^+$	$N(2.5, 1)$	$N(1, 1)^+$	$U[0, 0.7]$
2	$N(-2.5, 0.2^2)$	$N(1, 0.2^2)^+$	$N(2.5, 0.2^2)$	$N(1, 0.2^2)^+$	$U[0, 0.7]$
3	$N(-2.5, 1)$	$N(1, 1)^+$	$N(2.5, 1)$	$N(1, 1)^+$	$U[-0.9, -0.7]$
4	$U[-5, 0]$	$U[0, 4]$	$U[0, 5]$	$U[0, 4]$	$U[0, 0.7]$
5	$U[-3.1, -1.9]$	$U[0.4, 1.6]$	$U[1.9, 3.1]$	$U[0.4, 1.6]$	$U[0, 0.7]$

Table 2.7: RR and RRMSE for different number of patients, scenario 1, with different priors or methods. Simulation based on 1000 replications.

		α_E		β_E		α_T		β_T		ρ	
Setting	n	RR	RRMSE	RR	RRMSE	RR	RRMSE	RR	RRMSE	RR	RRMSE
Prior 1	25	1.05	0.21	1.11	0.42	1.16	0.24	1.05	0.41	0.87	0.59
Prior 1	50	1.03	0.17	1.08	0.32	1.12	0.27	1.05	0.39	0.84	0.59
Prior 1	100	1.02	0.12	1.05	0.23	1.10	0.28	1.02	0.30	0.84	0.58
Prior 1	200	1.01	0.09	1.02	0.16	1.08	0.28	1.01	0.23	0.89	0.52
Prior 2	25	1.21	0.22	0.96	0.10	1.22	0.22	1.37	0.46	0.76	0.64
Prior 2	50	1.19	0.19	0.95	0.12	1.22	0.22	1.19	0.30	0.75	0.63
Prior 2	100	1.15	0.16	0.95	0.12	1.21	0.21	1.05	0.18	0.79	0.59
Prior 2	200	1.11	0.12	0.96	0.11	1.21	0.21	0.97	0.12	0.85	0.53
Prior 3	25	1.05	0.20	1.11	0.42	1.16	0.25	1.06	0.42	1.50	2.50
Prior 3	50	1.04	0.17	1.07	0.33	1.13	0.27	1.03	0.36	1.46	2.47
Prior 3	100	1.02	0.12	1.03	0.22	1.10	0.28	1.02	0.30	1.43	2.43
Prior 3	200	1.01	0.09	1.01	0.14	1.09	0.29	1.01	0.25	1.41	2.41
Prior 4	25	1.04	0.28	1.24	0.68	1.05	0.90	1.79	1.87	0.83	0.61
Prior 4	50	1.01	0.18	1.11	0.36	1.12	0.85	1.35	0.99	0.82	0.60
Prior 4	100	1.01	0.13	1.06	0.24	1.16	0.75	1.13	0.56	0.84	0.58
Prior 4	200	1.00	0.09	1.02	0.16	1.13	0.61	1.06	0.39	0.87	0.54
Prior 5	25	1.08	0.18	1.06	0.36	1.12	0.24	1.06	0.33	0.83	0.59
Prior 5	50	1.05	0.14	1.04	0.28	1.12	0.24	1.00	0.22	0.81	0.60
Prior 5	100	1.03	0.10	1.03	0.22	1.12	0.24	0.97	0.17	0.83	0.58
Prior 5	200	1.02	0.07	1.02	0.16	1.12	0.23	0.96	0.14	0.88	0.52
Wide	25	1.05	0.22	1.15	0.43	1.11	0.29	1.18	0.57	0.84	0.59
Wide	50	1.03	0.19	1.10	0.34	1.09	0.27	1.09	0.37	0.81	0.61
Wide	100	1.01	0.14	1.06	0.23	1.07	0.23	1.03	0.24	0.83	0.59
Wide	200	1.01	0.10	1.03	0.16	1.04	0.19	1.02	0.18	0.86	0.56
MLE	25	0.80	7.59	0.34	33.10	2.61	3.93	4.10	11.28	0.70	1.62
MLE	50	1.02	0.21	1.06	0.36	2.19	3.16	1.44	3.27	0.74	1.41
MLE	100	1.01	0.14	1.03	0.25	1.67	2.10	1.05	0.56	0.84	1.12
MLE	200	1.00	0.09	1.02	0.17	1.28	1.12	1.03	0.39	0.96	0.84

Table 2.8: RR and RRMSE for different number of patients, scenario 2, with different priors or methods. Simulation based on 1000 replications.

		α_E		β_E		α_T		β_T		ρ	
Setting	n	RR	RRMSE	RR	RRMSE	RR	RRMSE	RR	RRMSE	RR	RRMSE
Prior 1	25	1.08	0.27	1.22	0.72	1.15	0.22	1.04	0.38	0.79	0.62
Prior 1	50	1.06	0.24	1.12	0.51	1.12	0.26	1.03	0.36	0.81	0.59
Prior 1	100	1.04	0.20	1.04	0.33	1.10	0.29	1.03	0.31	0.82	0.57
Prior 1	200	1.02	0.15	1.01	0.24	1.09	0.29	1.00	0.25	0.87	0.50
Prior 2	25	1.21	0.21	1.60	0.64	1.22	0.22	1.38	0.47	0.66	0.68
Prior 2	50	1.19	0.19	1.43	0.49	1.22	0.22	1.18	0.29	0.71	0.64
Prior 2	100	1.17	0.18	1.26	0.34	1.21	0.21	1.05	0.17	0.77	0.58
Prior 2	200	1.15	0.16	1.13	0.23	1.21	0.21	0.97	0.12	0.82	0.53
Prior 3	25	1.10	0.27	1.21	0.69	1.15	0.24	1.06	0.42	1.48	2.48
Prior 3	50	1.08	0.24	1.11	0.47	1.12	0.27	1.04	0.38	1.44	2.45
Prior 3	100	1.06	0.20	1.05	0.34	1.11	0.28	1.01	0.30	1.42	2.42
Prior 3	200	1.04	0.16	1.02	0.24	1.09	0.29	1.00	0.25	1.41	2.41
Prior 4	25	1.12	0.58	1.21	0.84	1.05	0.89	1.70	1.73	0.85	0.59
Prior 4	50	1.06	0.42	1.08	0.53	1.12	0.82	1.30	0.92	0.83	0.59
Prior 4	100	1.04	0.29	1.04	0.36	1.11	0.70	1.13	0.54	0.87	0.55
Prior 4	200	1.02	0.17	1.02	0.25	1.08	0.55	1.05	0.34	0.90	0.50
Prior 5	25	1.11	0.23	1.20	0.54	1.11	0.23	1.06	0.35	0.81	0.60
Prior 5	50	1.09	0.20	1.11	0.38	1.11	0.23	1.01	0.24	0.83	0.58
Prior 5	100	1.06	0.16	1.05	0.26	1.12	0.23	0.97	0.17	0.86	0.54
Prior 5	200	1.04	0.12	1.01	0.19	1.11	0.23	0.97	0.14	0.87	0.51
Wide	25	1.10	0.28	1.22	0.63	1.09	0.28	1.25	0.63	0.80	0.61
Wide	50	1.07	0.26	1.10	0.40	1.06	0.26	1.13	0.42	0.81	0.60
Wide	100	1.05	0.23	1.04	0.25	1.04	0.23	1.06	0.28	0.81	0.57
Wide	200	1.04	0.19	1.02	0.17	1.03	0.18	1.03	0.19	0.88	0.51
MLE	25	1.78	16.11	16.89	496.77	2.63	3.96	3.87	7.09	0.62	1.59
MLE	50	1.14	1.20	1.07	0.48	2.19	3.14	1.58	3.05	0.79	1.28
MLE	100	1.03	0.41	1.03	0.35	1.48	1.74	1.09	1.02	0.91	0.96
MLE	200	1.02	0.18	1.00	0.25	1.18	0.83	1.04	0.33	0.99	0.70

Table 2.9: RR and RRMSE for different number of patients, scenario 3, with different priors or methods. Simulation based on 1000 replications.

		α_E		β_E		α_T		β_T		ρ	
Setting	n	RR	RRMSE	RR	RRMSE	RR	RRMSE	RR	RRMSE	RR	RRMSE
Prior 1	25	1.14	0.29	1.62	1.53	1.15	0.24	1.04	0.39	0.80	0.61
Prior 1	50	1.14	0.29	1.35	1.07	1.11	0.27	1.05	0.38	0.80	0.60
Prior 1	100	1.13	0.28	1.17	0.75	1.09	0.27	1.02	0.29	0.86	0.55
Prior 1	200	1.11	0.26	1.06	0.51	1.08	0.28	1.00	0.23	0.91	0.49
Prior 2	25	1.20	0.21	3.43	2.51	1.22	0.22	1.38	0.47	0.61	0.70
Prior 2	50	1.20	0.20	2.77	1.88	1.21	0.22	1.19	0.30	0.67	0.66
Prior 2	100	1.19	0.20	2.11	1.23	1.21	0.21	1.05	0.17	0.75	0.60
Prior 2	200	1.19	0.20	1.59	0.74	1.21	0.21	0.97	0.12	0.85	0.51
Prior 3	25	1.16	0.29	1.62	1.61	1.16	0.24	1.05	0.41	1.47	2.47
Prior 3	50	1.15	0.28	1.30	1.07	1.14	0.27	1.02	0.34	1.44	2.44
Prior 3	100	1.13	0.28	1.18	0.75	1.11	0.29	1.02	0.33	1.42	2.42
Prior 3	200	1.11	0.27	1.08	0.54	1.09	0.29	1.00	0.24	1.41	2.41
Prior 4	25	1.12	0.80	1.64	1.68	0.97	0.87	1.91	1.99	0.77	0.63
Prior 4	50	1.10	0.74	1.33	1.05	1.07	0.82	1.42	1.11	0.82	0.59
Prior 4	100	1.10	0.65	1.14	0.74	1.10	0.71	1.17	0.57	0.85	0.55
Prior 4	200	1.10	0.54	1.03	0.53	1.09	0.57	1.07	0.37	0.88	0.50
Prior 5	25	1.08	0.20	2.26	1.40	1.11	0.23	1.07	0.37	0.75	0.62
Prior 5	50	1.05	0.16	2.12	1.16	1.12	0.23	1.01	0.23	0.77	0.60
Prior 5	100	1.02	0.12	2.05	1.06	1.11	0.22	0.97	0.17	0.84	0.54
Prior 5	200	1.00	0.08	2.02	1.02	1.09	0.20	0.97	0.14	0.87	0.49
Wide	25	1.17	0.29	1.20	0.86	1.11	0.27	1.27	0.69	0.79	0.61
Wide	50	1.15	0.28	1.09	0.56	1.08	0.24	1.13	0.43	0.82	0.58
Wide	100	1.12	0.29	1.04	0.38	1.06	0.22	1.06	0.28	0.85	0.54
Wide	200	1.10	0.28	1.01	0.26	1.03	0.18	1.03	0.18	0.92	0.48
MLE	25	1.86	3.11	1.35	1.35	2.63	3.94	4.80	10.98	0.64	1.49
MLE	50	1.46	2.32	1.17	0.95	2.12	3.02	1.61	4.26	0.77	1.22
MLE	100	1.45	1.96	1.07	0.73	1.53	1.81	1.05	0.49	0.87	0.95
MLE	200	1.26	1.39	1.01	0.54	1.21	0.95	1.02	0.33	0.95	0.70

Table 2.10: RR and RRMSE for different number of patients, scenario 4, with different priors or methods. Simulation based on 1000 replications.

		α_E		β_E		α_T		β_T		ρ	
Setting	n	RR	RRMSE	RR	RRMSE	RR	RRMSE	RR	RRMSE	RR	RRMSE
Prior 1	25	1.15	0.29	1.62	1.53	1.21	0.25	1.02	0.63	0.80	0.60
Prior 1	50	1.15	0.28	1.28	1.02	1.20	0.25	0.98	0.47	0.83	0.56
Prior 1	100	1.13	0.27	1.12	0.71	1.18	0.26	0.96	0.35	0.87	0.52
Prior 1	200	1.10	0.26	1.05	0.52	1.15	0.28	0.96	0.26	0.92	0.43
Prior 2	25	1.20	0.21	3.44	2.52	1.22	0.22	1.92	1.11	0.63	0.68
Prior 2	50	1.20	0.20	2.78	1.88	1.22	0.22	1.45	0.62	0.70	0.63
Prior 2	100	1.19	0.20	2.10	1.22	1.23	0.23	1.16	0.33	0.77	0.57
Prior 2	200	1.19	0.20	1.62	0.75	1.23	0.23	1.02	0.19	0.88	0.46
Prior 3	25	1.15	0.29	1.67	1.63	1.20	0.26	1.08	0.79	1.45	2.45
Prior 3	50	1.14	0.27	1.31	1.01	1.21	0.26	0.97	0.44	1.42	2.42
Prior 3	100	1.13	0.27	1.13	0.70	1.19	0.26	0.95	0.31	1.41	2.41
Prior 3	200	1.12	0.26	1.03	0.52	1.16	0.28	0.96	0.27	1.40	2.40
Prior 4	25	1.11	0.81	1.49	1.44	1.12	0.99	1.45	1.38	0.81	0.60
Prior 4	50	1.09	0.74	1.25	0.97	1.20	0.97	1.25	0.87	0.84	0.57
Prior 4	100	1.10	0.64	1.10	0.70	1.22	0.93	1.13	0.61	0.87	0.51
Prior 4	200	1.07	0.53	1.03	0.52	1.22	0.86	1.06	0.48	0.91	0.45
Prior 5	25	1.06	0.17	2.27	1.38	0.97	0.06	2.04	1.05	0.73	0.60
Prior 5	50	1.04	0.14	2.15	1.18	0.96	0.04	2.01	1.01	0.75	0.57
Prior 5	100	1.01	0.10	2.07	1.08	0.96	0.04	2.01	1.01	0.78	0.54
Prior 5	200	0.99	0.06	2.03	1.03	0.95	0.05	2.00	1.00	0.83	0.46
Wide	25	1.18	0.29	1.17	0.81	1.17	0.28	1.25	1.00	0.83	0.58
Wide	50	1.17	0.29	1.09	0.54	1.15	0.28	1.10	0.58	0.85	0.56
Wide	100	1.15	0.30	1.03	0.37	1.13	0.28	1.03	0.36	0.89	0.52
Wide	200	1.11	0.28	1.00	0.26	1.10	0.28	1.00	0.27	0.91	0.45
MLE	25	3.40	4.90	1.49	2.86	1.60	3.00	1.41	1.44	0.77	1.30
MLE	50	3.28	4.70	1.18	0.99	1.47	2.41	1.18	1.00	0.87	1.05
MLE	100	2.75	3.98	1.07	0.72	1.42	2.08	1.05	0.73	0.98	0.78
MLE	200	2.19	3.11	1.02	0.56	1.31	1.52	1.01	0.52	0.99	0.60

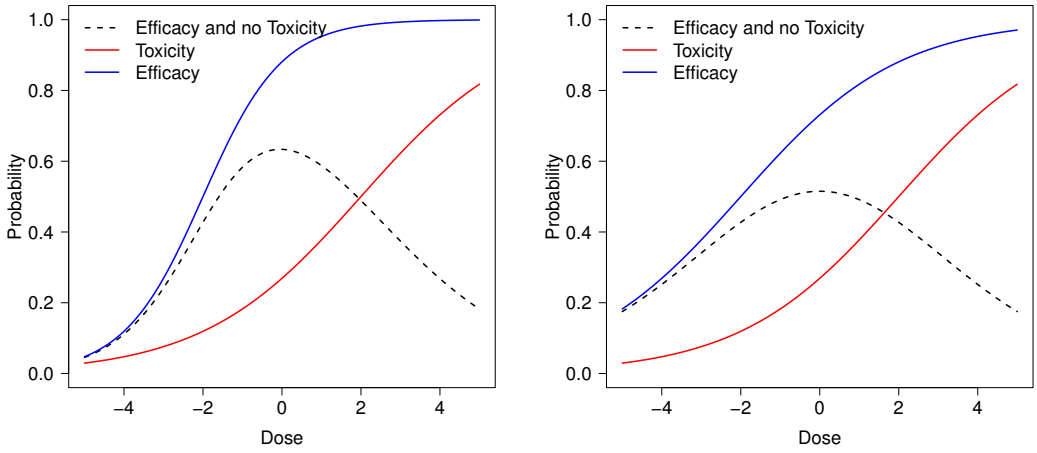

(a) Scenario 1

(b) Scenario 2

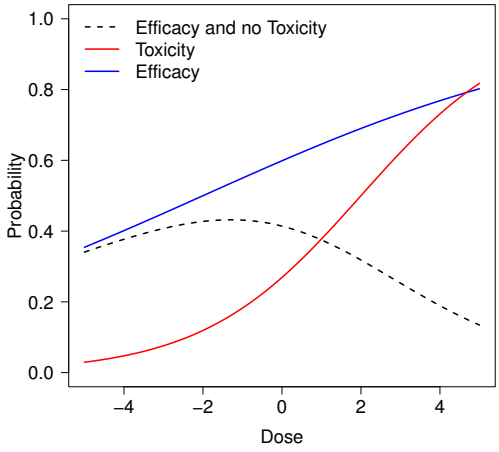

(c) Scenario 3

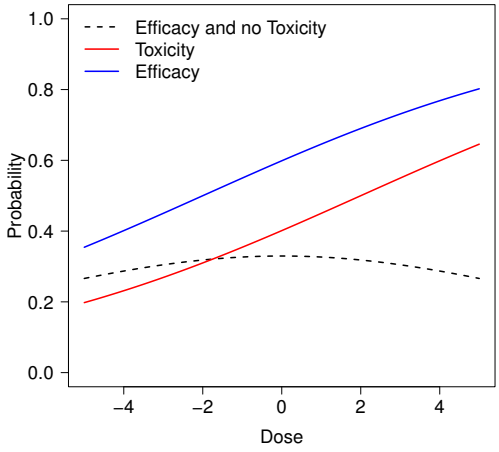

(d) Scenario 4

Figure 2.1: Toxicity, efficacy and "efficacy and no toxicity" probabilities for four scenarios. The parameters are $(-2, 1, 2, 0.5, 0.5)$, $(-2, 0.5, 2, 0.5, 0.5)$, $(-2, 0.2, 2, 0.5, 0.5)$ and $(-2, 0.2, -2, 0.2, 0.5)$ respectively.

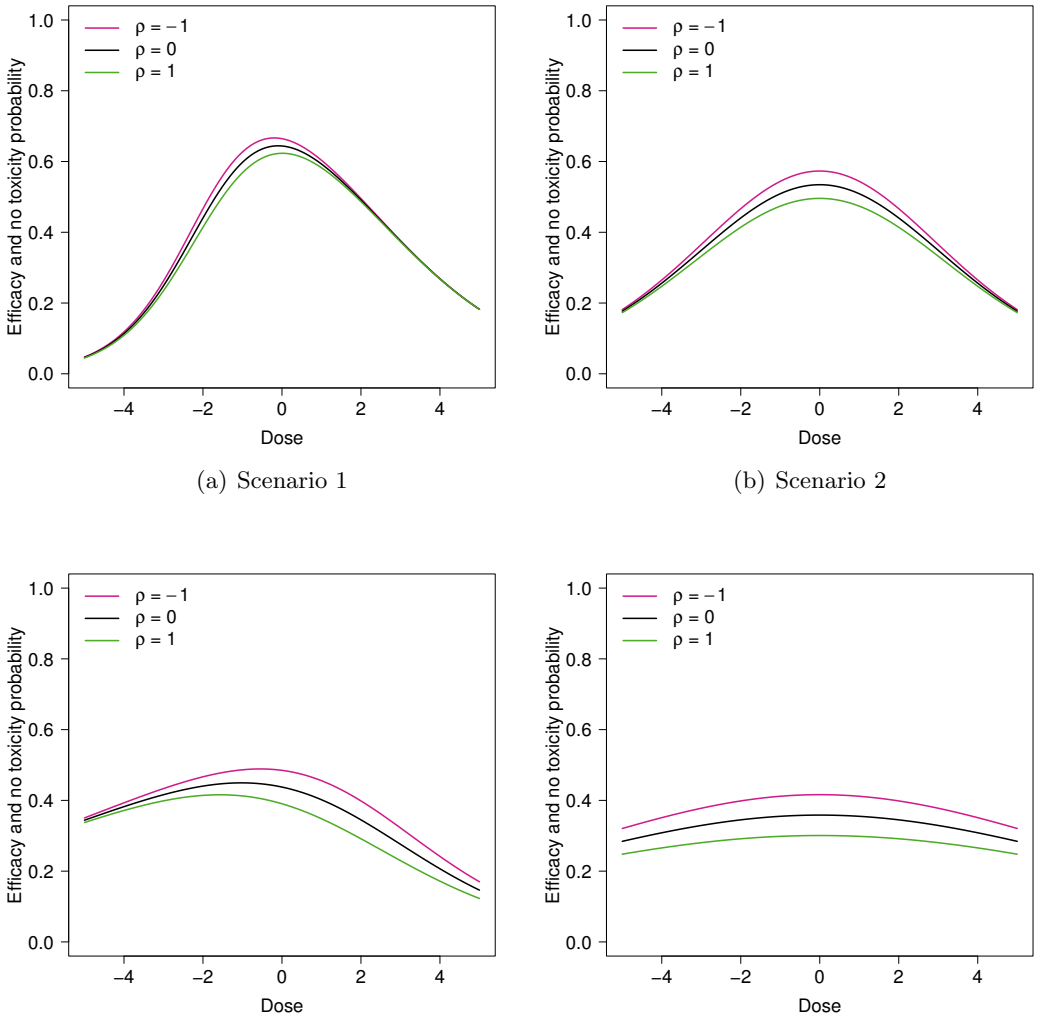

(a) Scenario 1

(b) Scenario 2

(c) Scenario 3

(d) Scenario 4

Figure 2.2: Efficacy and no toxicity probabilities for four scenarios. The correlation parameter is assumed to be -1, 0 and 1 for each scenario.

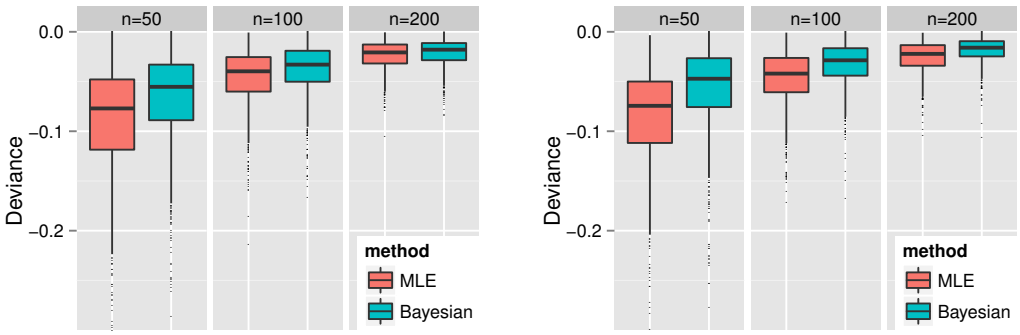

(a) Scenario 2

(b) Scenario 4

Figure 2.3: The deviance measure in scenarios 2 and 4 for both Bayesian and MLE methods.

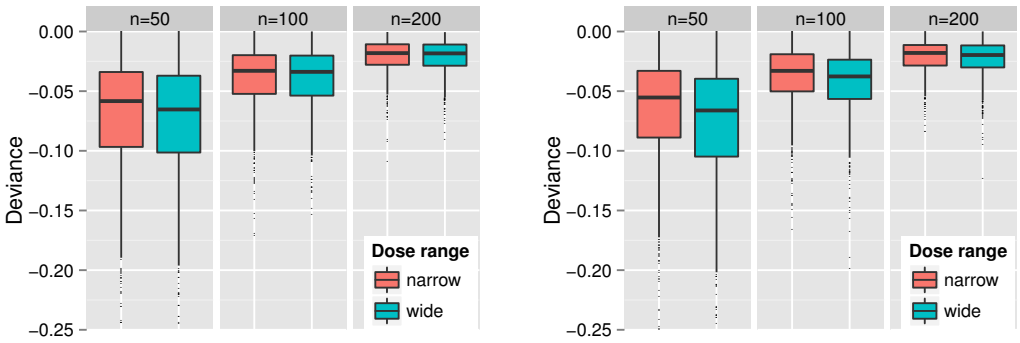

(a) Scenario 1

(b) Scenario 2

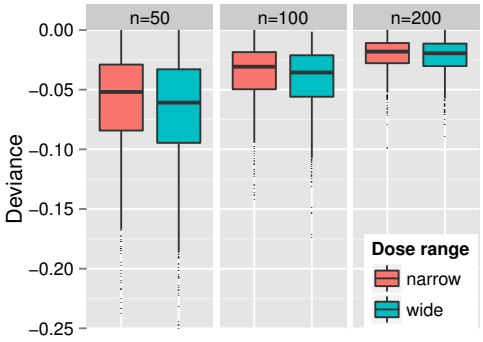

(c) Scenario 3

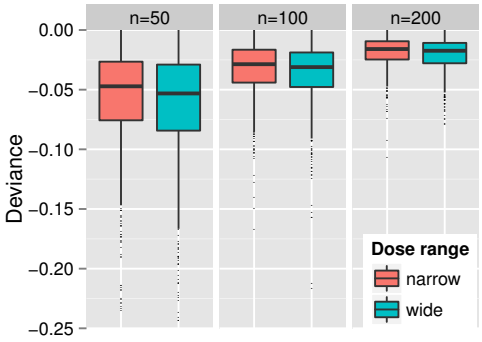

(d) Scenario 4

Figure 2.4: Deviance measure in scenarios 1 – 4 for designs operated on wide and narrow dose ranges.

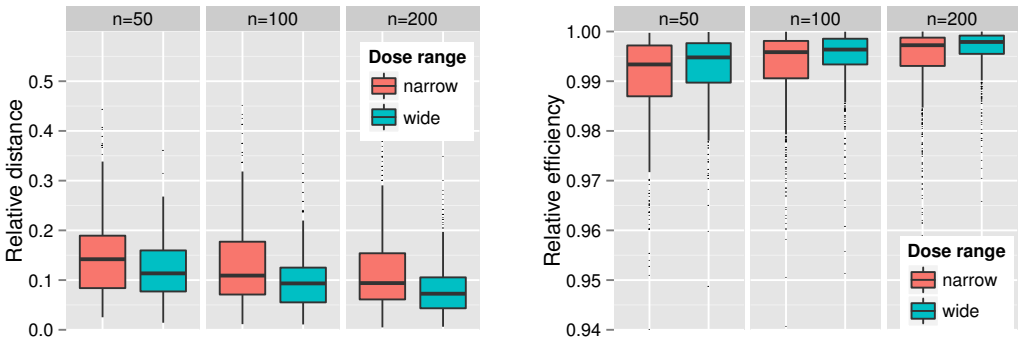

(a) Scenario 1

(b) Scenario 1

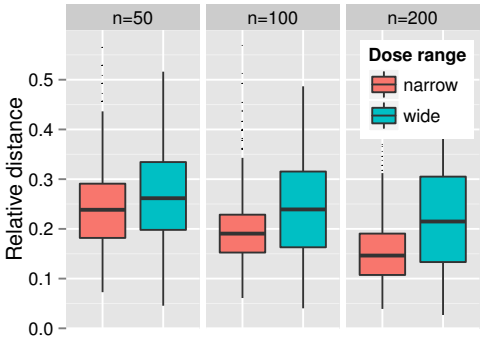

(c) Scenario 4

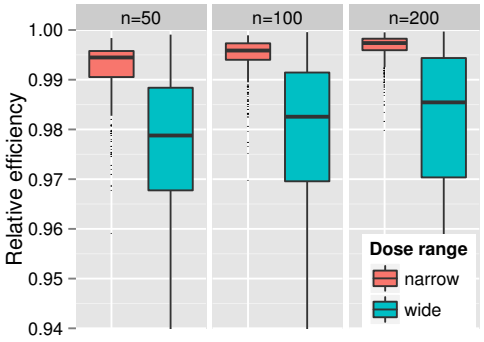

(d) Scenario 4

Figure 2.5: Dose allocation measures in scenarios 1 and 4 for designs operated on wide and narrow dose ranges.

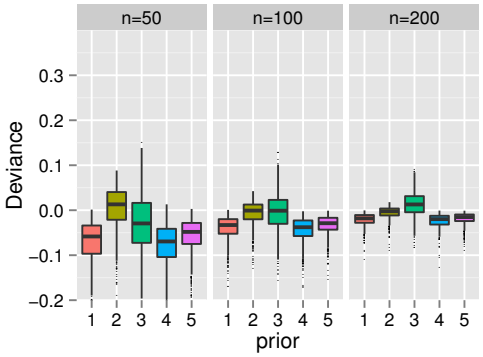

(a) Scenario 1

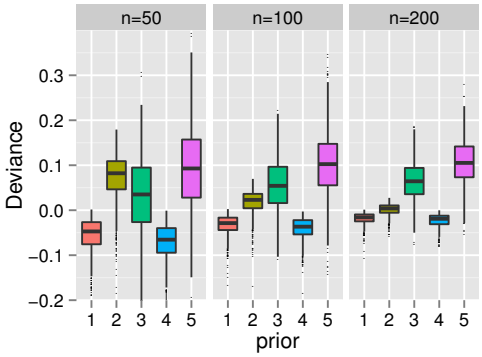

(b) Scenario 4

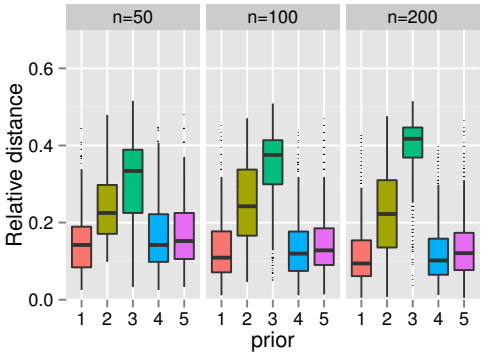

(c) Scenario 1

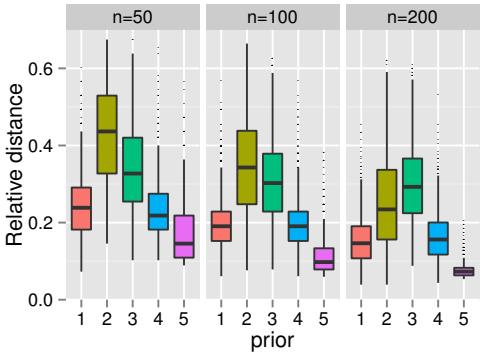

(d) Scenario 4

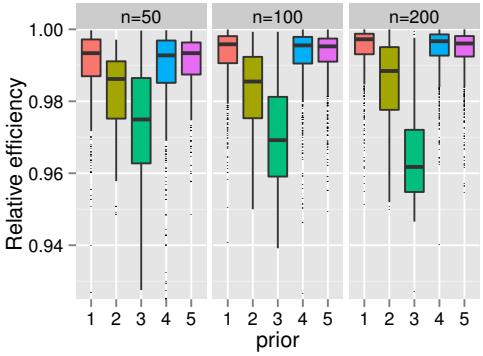

(e) Scenario 1

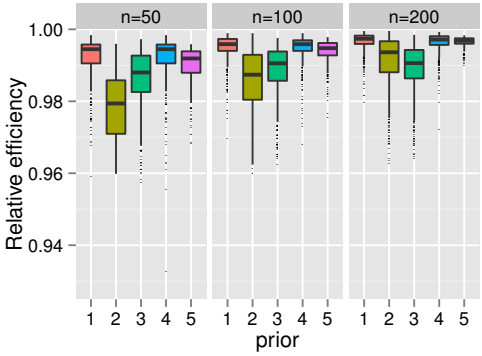

(f) Scenario 4

Figure 2.6: Design measures for scenarios 1 and 4 using different priors.

Chapter 3: Sequential Designs with Penalty

Model-based Bayesian methods are used to incorporate experts' opinion, identify an optimal dose with high efficacy and low toxicity, and reduce the sample size (Thall and Cook, 2004). In particular, Thall and Cook (2004) consider desirability contours as trade-offs between efficacy and toxicity to protect patients. The doses with high desirabilities are those with low toxicity and high efficacy. Their philosophy is to maximize the desirability. To find the optimal dose, patients are sequentially assigned to the dose with the most desirable utility, while the utility function is updated with new parameter estimates. We call these "best intention" Bayesian designs. The term "best intention" was first used by Fedorov et al. (2011) who examine best intention designs in the frequentist setting.

Although great focus has been put on inventing novel utility functions and stopping rules in literature, few have examined the convergence of best intention Bayesian designs. On the one hand, these designs target a single dose, and the targeting depends on estimation of parameters. However, single dose designs are typically poor at producing good estimates. It is possible that the selected dose has the least utility. On the other hand, one cannot fully utilize the dose-response relationship without accurate parameter estimates. Better doses beyond the preselected dose range may be inferred from the dose-response relationship, which leads to a more powerful phase III trial. Biostatistician should be aware and cautious about these facts when devising best intention Bayesian designs.

From a different point of view, we conducted sequential designs in the last chapter to maximize the information about the dose-response relationship. In practice, difficulty arises from ethical considerations. A design may be optimal for efficiency or precision, but it may also allocate participants at a high dose level with high toxicity. Therefore, achieving a balance between the information gain and participants' protection becomes important in these designs. There are some methods to address this problem: the restricted dose space

method (Mats, Rosenberger and Flournoy, 1998) sets an upper bound for the dose; the local penalty function method (Dragalin and Fedorov, 2006) penalizes high doses with high toxicity and low doses with low efficacy through a penalty function. The philosophy in the information-based methods is to maximize the information gain over a restricted space. So far we have two different philosophies for best intention design and information-based design. Therefore, an interesting question is which design philosophy has better performance with respect to important outcome measures.

In this chapter, we first introduce Thall and Cook's (2004) method and compare it with an information-based design. Then we extend the information-based method by considering a class of penalty functions. Next, we conduct simulations to illustrate possible problems with the best intention design using the proposed penalty function. Last, we examine the performance of best intention designs and information-based designs under different scenarios. In particular, we examine how the penalty function impacts the estimation convergence rate in Bayesian sequential designs and how the estimation impacts the penalty function.

3.1 Common Bayesian sequential procedure

During the design stage of a trial, we do not have information on the true parameter value for the given model. But the clinician may supply vague information about these parameters, either from preclinical studies or clinical studies with a similar compound. Therefore it is worthwhile to incorporate this prior information into the design, so that we can efficiently achieve the clinical goal – identifying a safe and efficacious dose. Also from an ethical perspective, patients should not be assigned on highly toxic or ineffective doses. Therefore we have to implement a design to achieve the clinical and ethical goals at the same time. Bayesian adaptive sequential designs can serve this purpose very well. Instead of assigning all patients at the same time, we assign them one by one or in groups as we accumulate information about the true parameters. With more information about the dose-response relationship, we can have a better chance of selecting doses and protecting participants.

Suppose the prior distribution density function $\pi(\theta)$ for θ is defined on Θ. Also a penalty function $\delta(x;\theta)$ is chosen to achieve certain clinical and ethical goals, so that the optimal dose x^* minimizes $\delta(x;\theta)$. After we initialize n_0 patients to the lowest dose x_1, a Bayesian sequential design proceeds iteratively through the following steps:

1. Suppose n participants are assigned to the preselected doses. Their bivariate responses are collected and denoted by $\mathcal{D}_n$;

2. Obtain the posterior distribution $\pi_n(\theta|\mathcal{D})$, and apply some stopping rule. If the stopping rule is triggered or a preseclected maximal sample size N is reached, terminate the trial;

3. Assign the next participant to the dose that minimizes the expected penalty:

$$x^{(n+1)} = \arg \max_{x \in \{x_1, \ldots x_m\}} \int_\Theta \delta(x, \theta) \pi_n(\theta|\mathcal{D}_n) d\theta;$$

4. Obtain the new response (y_{n+1}, z_{n+1}).

It is commonly assumed that, as iterations increase, the Bayesian sequential procedure converges by assigning most patients to the optimal dose x^* with accurate parameter estimates. We argue that it is not always the case when the design aims at a single optimal dose. Throughout the paper, we emphasize parameter estimation and exclude stopping rules. Technical problems such as prior elicitation and Markov chain Monte Carlo (MCMC) algorithms are deferred to later sections.

3.1.1 Computational issues

The simulation programs are coded in C and run on a Unix cluster with an Intel Xeon CPU E5440 (2.84 GHz) and Red Hat 4.3.0 operating system. We use *gcc version 4.3.0* as the C compiler. The run time is estimated as 11 hours for $\lambda = 1$ and 17 hours for $\lambda < 1$ with a sample size of 200. All the graphics are done in R with the *ggplot2* packages version 0.9.3.1 (Wickham, 2009).

The main algorithm for assignment is included in Section 3.1. It is divided into three components – generating responses according to the Gumbel model, generating MCMC samples, and calculating the determinant of the Fisher information matrix.

To generate the response from the Gumbel model, we generate responses from a multinomial distribution with cell probabilities given in Section 1.4.1. To generate samples from a uniform distribution, we use the random number generator *r8_uniform_01* developed by Burkardt (2006) in the NORMAL library. The seed for the random number generator is fixed at 1234.

All the mathematical functions can be found in the *math.h* library included in most C compilers. Matrices and vectors are coded in the array structure in C. The determinant of a matrix is obtained by using the Gauss-Jordan elimination method (Fraleigh and Beauregard, 1995). The inverse of a three by three matrix is obtained using the inverse formula (Fraleigh and Beauregard, 1995). The expression of the Fisher information matrix is given in Section 1.4.2. The MCMC algorithm is given in Section 2.3.2.

3.2 Thall and Cook's best intention design

Similar to Fedorov et al. (2011), we define a design to be best intention when all patients are assigned to a single dose given the true parameter. Here penalty functions are used to measure the trade-off between efficacy and toxicity. An ethical penalty function takes high values on doses that are toxic or ineffective. An dose is ethically the most desirable when the penalty is minimized. With only toxicity in tradition phase I trials, the penalty function measures the distance between a dose and the maximum tolerated dose. Substituting ethical penalty functions in the common Bayesian sequential procedure, researchers have introduced many Bayesian designs (Thall and Cook, 2004; Yin et al., 2006; etc.). They are all best intention Bayesian designs because they are aiming at assigning most patients to the (single) most desirable dose. In this section, we conduct simulation to compare Thall and Cook's (2004) best intention design with our information-based information.

Recall that Thall and Cook (2004) considered the following toxicity and efficacy effects in a Gumbel model:

$$x_E = \mu_E + \beta_{E1}x + \beta_{E2}x^2$$

$$x_T = \mu_T + \beta_T x.$$

The desirability function is defined in Chapter 1 by:

$$\delta(\pi_q) = \frac{\rho(\pi_p)}{\rho(\pi_q)} - 1.$$

The algorithm starts from the lowest dose and escalates through steps. At each step, a cohort of three patient are assigned to the dose with the highest desirability according to the posterior. If the expected toxicity is higher than a critical value, the algorithm stops. They showed that the trade-off algorithm assigns the most patients to the most desirable dose. However, because efficiency of estimation is not their primary concern, yet their algorithm relies on precise parameter estimation, the trade-off method may not be robust to model misspecification. As a result, the algorithm may not be efficient compared to our algorithm with respect to participant protection.

3.2.1 Practical criteria

To evaluate the performance of the penalized Bayesian sequential design and Thall and Cook's trade-off design, we record the measures from two categories, one related to information gain and the other related to toxicity and efficacy trade-off desirability:

1. Expected number of toxicities, i.e. events of $Z = 1$, denoted by "Toxicity";
2. Expected number of efficacies, i.e. events of $Y = 1$, denoted by "Efficacy";
3. Expected number of efficacies but not toxicities, i.e. events of $Y = 1$ and $Z = 0$, denoted by "Eff-Tox";

4. Expected average desirabilities across four doses, denoted by "Desirability" (larger is better);

5. Expected number of patients assigned to dose with the highest desirability, denoted by "DesHigh";

6. Expected precision, $\log |M(\xi, \theta_0)|^{-1/6}$, denoted by "Precision" (smaller is better);

7. Expected total cost, $C(\xi, \theta_0)$, denoted by "Cost" (smaller is better);

8. Expected information loss per cost, $|M(\xi, \theta_0)|^{1/6}/C(\xi, \theta_0)$, denoted by "Info/Cost" (larger is better).

The last three measures are used in Dragalin and Fedorov (2006). The trade-off design (Thall and Cook, 2004) is expected to assign most participants to the most desirable dose. On the other hand, the penalized design (Dragalin and Fedorov, 2006) is expected to achieve a high Info/Cost ratio, which balances the parameter estimation precision and the ethical cost.

3.2.2 Simulation with severely misspecified prior

We present an example in which the least desirable dose is assigned for Thall and Cook's method. We assume the prior is misspecified in a way that it overestimates the toxicity

probability. In the Gumbel model, the prior is

$$\mu_E \sim N(-0.5, 1.1^2),$$

$$\beta_{E1} \sim N(0.5, 0.9^2),$$

$$\beta_{E2} \sim N(0.1, 1.1^2),$$

$$\mu_T \sim N(0.25, 0.9^2),$$

$$\beta_T \sim N(3, 1.7^2),$$

$$\alpha \sim U[-1, 1],$$

where the true parameter is $(0.5, 0.5, 0.1, -1.5, 0.25, 0)$. The corresponding desirability under the true parameter is included in Table 3.6.

In a simulation of 300 iterations, Thall and Cook's method always assigns patients to dose level 1. Specifically, it assigns 300 patients to the lowest dose. However, although they claim the method assigns most patients to the most desirable dose, it is not the case in our setting. The method also gives poor parameter estimation. The parameter are estimated as $(0.05, 0.40, 0.97, -0.90, 1.51, -0.57)$ at the 300th iteration.

Dragalin and Fedorov's methods do not have this property. It assigns only 93 patients to the lowest dose and assigns 123 patients to the highest dose with the most desirability (Table 3.6). It also yields better estimators after the 300th iteration. The estimation from two methods are summarized in Table 3.2.

3.2.3 Simulation with mildly misspecified prior

We conduct a simulation study to compare the penalized method with the trade-off method. We choose a similar setting as Thall and Cook (2004) in which the design sample size is 36 and the initial sample size is 3. We assume the dose-response curve follows a Gumbel model with true parameter $\theta_0 = (-1.5, 3, 0.5, -0.5, 0.5, 0.5)$ in scenario 1 and $(-1, 1, 0.5, -1, 1, 0.5)$

in scenario 2. The prior distribution is chosen almost the same as in TD2004: $\mu_E, \beta_{E1}, \beta_{E2}, \mu_T$ and β_T are independently normally distributed with means and standard deviations (-1.496, 1.113) for μ_E, (1.180, 0.069) for β_{E1}, (0.149, 1.192) for β_{E2}, (-0.619, 0.941) for μ_T and $(0.587, 1.659)$ for β_T. The different setting is the correlation parameter α, which is assumed to follow a uniform distribution on $[-1, 1]$. All parameters are independently distributed. Because there are six parameters, it is computationally infeasible to compute six dimensional integrations over the posterior distribution. Thus we employ Monte Carlo methods to approximate the expectations in a six dimensional space. Specifically, we generate two chains, each with 5000 θ's, and approximate the expectation by the average.

The raw dose d is transformed to x in the same way as Thall and Cook (2004): $x_i = \log(d_i) - \text{mean}(d)$. Therefore, the dose space $\mathcal{X}$ consists of -0.79, -0.10, 0.30 and 0.59 instead of the four raw doses, 0.25, 0.5, 0.75 and 1. We use the same desirability function defined in Thall and Cook (2004) and obtain their desirabilities as -0.194, -0.137, 0.021 and 0.161 in scenario 1 and 0.002, 0.031, 0.080 and 0.109 in scenario 2. Their probabilities of efficacy and toxicity are $(0.027, 0.290)$, $(0.142, 0.366)$, $(0.368, 0.414)$ and $(0.611, 0.449)$ in scenario 1 and $(0.186, 0.143)$, $(0.250, 0.249)$, $(0.343, 0.332)$ and $(0.442, 0.399)$ in scenario 2. The true parameters are chosen so that the efficacy increases faster than toxicity in scenario 1, and they have comparable increasing rate in scenario 2.

We implement 5000 simulations for each scenario. One simulation of the penalized algorithm starts by assigning 3 patients to the lowest dose for a fair game. Then we begin the iteration in a few steps. Specifically, we estimate the posterior distribution based on the responses and the prior using MCMC. Then we assign the next patient to an optimal dose. Also we obtain the bivariate response for each patient according to the true model. We continue the iterations until all 36 patients are assigned. The trade-off algorithm is implemented in the same manner as Thall and Cook (2004).

As for the simulation result, we have 5498 simulations in scenario 1 and 5271 simulations in scenario 2 until the trade-off algorithm produces 5000 complete trials in each scenario. Table 3.3 includes the results of the two methods for scenario 1 and Table 3.4 for scenario

2. They perform similarly with respect to protecting the patients, for they result in the similar Toxicity, Efficacy and Eff-Tox. For desirability considerations, the trade-off method is expected to outperform the penalized method. However, the penalized method performs better by putting more patient on the dose with highest desirability in both scenarios. The penalized method also achieves a higher average desirability than the trade-off method in scenario 1. When it comes to the precision measure, the penalized method performs better with more estimation information (or low Precision value) and less cost as expected. As a result the penalized method has larger and better Info/Cost ratios. Note that our penalized method has smaller Monte Carlo variations and performs more consistently due to a more efficient parameter estimation. Overall, the penalized method performs more stably and provides more information than the trade-off method in our settings. The effect is amplified when efficacy increases faster than toxicity.

These results have been published in Gao and Rosenberger (2013) and have been reported by Gao at the statistics in Biomedical research conference at National Institute of Health (NIH) in 2012.

3.3 Penalty function

In this section, we consider a class of penalty function introduced by Dragalin and Fedorov (2006). Also, we examine local designs with the penalty functions and show how penalty impacts local designs.

As shown in the introduction, there are two types of penalty function, given by Dragalin and Fedorov (2006) and Thall and Cook (2004). The former is more complex in the sense that it uses p_{10}, so it depends on the correlation parameter. Consider a situation where the correlation is high for drug A and low for drug B, and the drugs share common toxicity and efficacy parameters. The drugs will have equal penalty using Thall and Cook's (2004) method, while drug A has more penalty than drug B using Dragalin and Fedorov's (2006)

method. In this dissertation, we use Dragalin and Fedorov's (2006) penalty function:

$$c(x, \theta; C_E, C_T) = [p_{10}(x, \theta)]^{-C_E}[1 - p_{01}(x, \theta) - p_{11}(x, \theta)]^{-C_T}.$$

They did not specify how to choose the constant C_E and C_T. As in Thall and Cook's (2004) trade-off method, we propose to fix C_E at 1 and determine C_T with experts' input. Specifically, we ask experts to provide two probability pairs, (p_{10}^A, p_T^A) and (p_{10}^B, p_T^B), with equivalent penalty (desirability). Then C_T is determined by

$$C_T = \frac{\log\left(p_{10}^A\right) - \log\left(p_{10}^B\right)}{\log\left(1 - p_T^B\right) - \log\left(1 - p_T^A\right)}.$$

For example, if (0.3, 0) and (0.5, 0.2) are selected to be equivalent, then C_T is 2.23.

Recall that the penalized D optimality in Dragalin and Fedorov (2006) is:

$$\Phi'(\xi, \theta) = \det\left\{\left[\frac{M(\xi, \theta)}{C(\xi, \theta)}\right]^{-1}\right\}.$$

Pronzato (2010b) considered a more flexible criterion:

$$\Phi''(\xi, \theta) = \log\det\left\{[M(\xi, \theta)]^{-1}\right\} + \lambda C(\xi, \theta),$$

where λ is introduced to control the compromise between information gain and ethical penalty: the design criterion puts more weights on ethical goals as λ increases. We rewrite Dragalin and Fedorov's (2006) penalized D optimality criterion as:

$$\Phi_\lambda(\xi, \theta) = (1 - \lambda)\log\det\left\{[M(\xi, \theta)]^{-1}\right\} + p\lambda\log C(\xi, \theta). \tag{3.1}$$

When $\lambda = 0.5$, it reduces to the criterion of Dragalin and Fedorov (2006). The corresponding

directional derivative is given by

$$\phi_\lambda = (1-\lambda)\left\{Tr[M_\xi^{-1}M_x] - p\right\} - \lambda\left\{\frac{c(x,\theta)}{C(\xi,\theta)}p - p\right\}.$$

Cook and Wong (1994) and Clyde and Chaloner (1996) showed optimal designs from weighted criteria are equivalent to designs from a constrained criterion for different models. Also they considered the relative information efficiency and ethical efficiency:

$$\text{Eff}_\lambda(\xi) = \exp\left(\frac{\Phi_\lambda(\xi_\lambda^*) - \Phi_\lambda(\xi)}{p}\right),$$

where ξ_λ^* is the optimal designs that minimizes the penalized criterion Φ_λ.

We examine the penalized local D-optimal design with different λ for our four scenarios in Table 3.5–3.8 with relative information and ethical efficiencies. We see that information efficiency decreases and ethical efficiency increases as λ increases. The penalty coefficient λ can be determined by examining the operational characteristics of a series of λ. One of the operational characteristics is the weight of the most desirable dose, denoted by w_c^* in Table 3.5–3.8. For example, if the design is required to put at least 50% patients on the most desirable dose, then we can choose $\lambda = 0.6, 0.8, 0.4$ and 0.4 respectively for scenario 1,2,3 and 4.

3.4 Problems with best intention designs when $\lambda = 1$

Although it is claimed that best intention Bayesian designs assign most patients to the most desirable dose, there is a chance that the designs perform poorly by assigning all patients to the least desirable doses. This can be understood by examining their dynamic nature. In these designs, assignments are determined by an ethical penalty function, and the quality of parameter estimates is important: good estimates lead to good assignments, which in turn lead to more accurate parameter estimates. However, good quality of parameter estimates rely on accurate prior information and balanced dose allocation. In the worst case scenario,

a trial starts with a misspecified prior that gives rise to unbalanced dose allocation. The information from unbalanced data is not sufficient to correct the misleading prior knowledge. In the end, all patients are assigned to only a few doses, which turn out to be the least desirable doses. Although the toxicity and efficacy probabilities can be learned at the allocated doses, much more could have been learned about the dose-response relationship. The loss is the chance to assign better doses and make better decisions with respect to further studies.

In this section, we consider best intention designs with $\lambda = 1$ in (3.1). We consider two examples where information-based designs with $\lambda < 1$ outperform best intention designs. Also we examine the convergence issue of best intention Bayesian methods through simulations.

In this section, we conduct simulations to examine the convergence of parameter estimates for the best intention Bayesian design and the information-based design introduced in Section 3. Two different design settings are given in Examples 1 and 2. Under each design setting, we run 1000 simulations for both model. In each simulation, we assign 200 patients sequentially according to the Bayesian sequential procedure. Then we examine the operating characteristic for both designs with different sample sizes. The most interesting sample size is 50, which is typical for phase I trials. Recall that given the compound penalty function (3.1), $\lambda = 1$ corresponds to a best intention design, while $\lambda < 1$ corresponds to an information based design.

Example 1

Best intention designs can cease to gain new information about the dose-response relationship once it begins assigning participators to a single dose. Given incorrect parameter estimates, the dose is probably not optimal. We demonstrate this convergence issue for the best intention designs here. Four doses are chosen. The prior toxicity is assumed to increase faster with dose than the true value. The prior efficacy is the same as the true value. Specifically, the prior mean is (-1.5, 1, -0.25, 3) for $(\alpha_E, \beta_E, \alpha_T, \beta_T)$ with standard

deviation 1. The prior distribution for ρ is a uniform distribution on [-1,1]. The true parameter is $(-1.5, 1, , 3, .6, -0.9)$. The efficacy, toxicity, "efficacy and no toxicity" probability and ethical penalty are included in Table 3.9.

Setting $\lambda = 1$ and 0.5, we obtain a best intention design and an information-based design. The dose allocation distributions are included in Figure 3.1(a). In most simulations, all 200 patients are assigned to the lowest dose level in the best intention design, which is the least desirable dose. There are more balanced assignments in the information-based design. Because of this, the information-based design has a better ethical performance for the set of parameters. Given a sample size of 50, the median proportion of toxicity in Figure 3.1(b) is 0.08 for the best intention design and 0.14 for the information-based design. However, the information-based design has higher proportions for efficacy and "efficacy and no toxicity" (0.78 and 0.66) than the best intention design (0.64 and 0.60).

The mean MAP estimates $\hat{\theta}_n$ are plotted against the sample size n for each parameter in Figure 3.2. For the best intention design, the estimates remain about the same as the misspecified prior for each n. Unlike in the best intention design, the estimates in the information-based design tend to converge to the true value. The slope parameter β_T converges faster than the intercept α_T. The correlation parameter converges with the slowest rate.

To illustrate the learning about the dose-response relationship, we estimate the optimal dose using MAP estimates. We include the distributions of the estimated optimal dose with n patients for n = 25, 50, 100, and 200 in Figure 3.1(c). The true optimal dose at 0.0 is not included in any interquartile range. However, the information-based design has closer estimates than the best intention designs. The estimates from the latter scatter around the prior optimal dose at -1.1.

In this example, the estimates from the best intention design converge to wrong values when the design starts assigning patients on a single dose, given the severely misspecified

prior.

Example 2

In this example, we assume that sponsors decide to launch an early phase trial for they believe a drug is promising based on non-human experiments. Then this belief can be translated to an optimistic prior, which assumes strong efficacy and weak toxicity. Therefore, we will use an optimistic prior to examine the operating characteristics of the designs. We choose four doses. The prior efficacy is assumed to increase faster with dose than the true value. Slightly higher toxicity is assumed in the prior. Specifically, the prior mean is (-2.5, 1, 2.5, 1) for $(\alpha_E, \beta_E, \alpha_T, \beta_T)$ with standard deviation 1. The prior distribution for ρ is a uniform distribution on [-1,1]. The true parameter is $(-2, 0.2, 2, 0.5, 0.5)$. The efficacy, toxicity, "efficacy and no toxicity" probability and ethical penalty are included in Table 3.9.

Setting $\lambda = 1$, 0.8 and 0.5, we obtain one best intention design and two information-based designs. All designs give similar MAP estimates for α_E, α_T and ρ. However, the estimates for β_E and β_T are different. We include the root of mean square error for the two parameters in Table 3.10. The best intention design ($\lambda = 1$) has poor estimates for β_E, because the RMSE does not change much as the sample size increases from 50 to 200. This is possibly caused by a small-range dose allocation. The information-based designs ($\lambda = 0.8$ and 0.5) have poor estimates for β_T with a sample size of 50.

For dose allocation weights, we note that the first dose is the optimal dose. In terms of simulation medians in Figure 3.3(a), the best intention design assigns most patients (80%) to the second dose, given a sample size of 50. In contrast, the information-based designs ($\lambda = 0.5$ and 0.8) assign the most participants (42% and 46%) to the optimal dose. The information-based designs have smaller variation for the first and second dose weights, meaning that they perform more stably than the best intention design. This is due to more accurate parameter estimates caused by multiple-dose allocations.

We examine ethical gains by looking at the median proportions of efficacy, toxicity and "efficacy and no toxicity" in Figure 3.3(b). The information-based design with $\lambda = 0.5$

has the highest efficacy (0.56), followed by the other two designs (0.54 for both). Due to multiple-dose allocations, the design with $\lambda = 0.5$ has higher median toxicity proportion (0.24) than the other two designs (0.18 for both). All the designs have the same "efficacy and no toxicity" proportion (0.40).

The optimal doses are estimated using the MAP estimates with n patients for $n = 25$, 50, 100, and 200. Their distributions are included in Figure 3.3(c). The true optimal dose at -2.45 is included in all interquartile ranges. However, the information-based designs have more confidence about the true optimal dose than the best intention design.

We see that the information-based designs have better and more stable dose-allocations than the best intention design by assigning most participants to the most desirable dose. They also provide better estimates about the optimal dose given the MAP estimates. The choice between the two information-based designs requires additional considerations. When toxicity is a major concern, we prefer to use the design with $\lambda = 0.8$ because it produces smaller number of toxicities. When performance stability is a major concern, we prefer the design with $\lambda = 0.5$.

3.5 Sequential designs with different priors

In this section, we simulate sequential designs to investigate several problems: how information-based designs compare with best intention designs on different measures; how design operating characteristics change with the support of the prior distribution; how designs with uniform priors perform compared to normal priors; how do the two types of design perform with misspecified correlation.

To answer these questions, we conduct simulations using the same scenarios as in the last chapter. We set λ to be 0.5, 0.8 and 1 for each scenario. We also select five different priors listed in Table 3.11. We conduct Monte Carlo simulations with 1000 replications for each scenario. We assign 1000 patients in one replication. After 50 patients are assigned, we monitor quantities from several aspects: information criteria, event proportions, estimated

optimal doses, parameter estimates and dose allocation weights.

Question 1: Do the designs converge to the optimal designs respectively?

Convergence can be measured in two ways. The first measure is the L^1 distance between the dose allocation weight of the optimal design and those in simulations shown in Figure 3.4. Here we use prior 1. As the sample size increases, the L^1 distance tends to zero and convergence is better. With a sample size of 50 in scenario 1, the information-based design with $\lambda = 0.5$ has the best and the smallest distance measures. All the designs show convergence because the distance measures tend to zero as the sample size increases. The best intention design with $\lambda = 1$ converges the fastest. The information-based design with $\lambda = 0.8$ has the slowest convergence rate. In scenario 4, there is a convergence issue with the best intention design, because the distance measure gets far away from zero when the sample size increases. However, the information-based designs still show convergence.

The second measure is the penalized criterion Φ_λ. We expect this measure to converge to one. We include the distribution of this measure in Figure 3.4. In scenario 1, all three designs show convergence. We also observe that the best intention design has the lowest penalized criterion value, due to different scales between the information criterion and the ethical criterion. In scenario 4, information-based designs show convergence, while the best intention design appears to converge with the slowest rate.

We conclude that information-based designs converge to the penalized optimal design regardless of prior misspecification. The best intention converges to the optimal design when the prior is mildly misspecified, and does not when the prior is severely misspecified.

Question 2: How does the best intention design compare with information-based designs with respect to operating characteristics?

We investigate this question with a fixed sample size of 50. The operating characteristics for scenarios 1 and 4 are included in Figure 3.5. We use prior 1. First, we investigate dose allocation weights. In scenario 1, all designs assign most patients to the last two doses

that have the least penalty. The best intention design also identifies the optimal dose by assigning most patients to dose 3. In scenario 4, the design with $\lambda = 0.5$ converges to the optimal dose allocation by assign more patients to optimal dose 1. The design with $\lambda = 0.8$ is unstable – although it assigns more patients to dose 1, it also assigns a great number of patients to dose 4. The best intention design still assigns the most patients to dose 3, which means the design is the so influenced by the prior information that it is not adapted well to data.

When it comes to the learning about the dose-response relationship, we investigate the distribution of the estimated optimal dose in Figure 3.5. In scenario 5, all designs perform well by having estimates close to the true optimal dose. In scenario 4, the best intention design performs the worst, by underestimating the optimal dose.

The proportions of events are included in Figure 3.5. In scenario 1, the three designs perform similarly. We see that designs with $\lambda = 0.5$ and 1 have lower toxicity than that with $\lambda = 0.8$. The design with $\lambda = 0.8$ has the highest efficacy. As a result, designs with $\lambda = 1$ and 0.8 have higher "efficacy and no toxicity" than the rest. In scenario 4, the three designs perform very similar with similar "efficacy and no toxicity". The design with $\lambda = 0.5$ has the least toxicity, while the design with $\lambda = 1$ has the highest efficacy.

In conclusion, the best intention design performs better when the prior is mildly misspecified. However, it may not adapt well to data. When the prior is severely misspecified, it may cause problems by assigning patients to highly toxic doses. The learning about the dose-response relationship is also poor in this case. In contrast, the robustness of the information-based design is less influenced by the prior information, and can provide better learning about the dose-response relationship and thus better protection.

Question 3: How do different priors affect the operating characteristics of designs?

We consider sequential design with priors 1, 2, 4 and 5 in Table 3.11. In particular, we examine the dose allocation weights, event proportions, estimated optimal doses and

parameter estimates in scenario 4. The dose allocation weights for scenario 1 are included in Figure 3.6. We see that weights of the design with $\lambda = 0.5$ are similar across priors, which shows robustness. The best intention design changes dramatically under different priors. For example, with the informative prior 5, it incorrectly assigns almost all patients to dose 3, which is the optimal dose in priori. The assignment becomes more and more unstable as the prior becomes less and less informative from 2 and 5 to 1 and 4. The design with $\lambda = 0.8$ acts as a compromise between the other two designs – it has balanced dose allocation, but weights becomes unstable as prior becomes noninformative.

The parameter estimates are included in Figure 3.7. Designs with the prior 1 have very good estimates about the slopes β_E and β_T. However, the estimates for α_E and α_T are not good, although the information-based designs have slightly better estimates for α_E. Designs with the prior 4 have good estimates for β_E and β_T. Although the true α_E and α_T are contained in the interquartile ranges of the estimates, the ranges are extreme large. This results in unstable estimates. Designs with the prior 2 overestimate β_E. The estimates for α_E and α_T are poor for they are stuck at prior values. Designs with the prior 5 have very good estimates for α_E and α_T. However, accurate estimates about the slope cannot be obtained because the true value is below the designated prior support.

The estimated optimal doses are included in Figure 3.8. Because this measure is based on parameter estimates, we can draw similar conclusion as the last paragraph. The information-based designs perform better than the best intention design with noninformative priors 1 and 4. This is due to worse parameter estimates for the best intention designs. However, there is no difference among the designs with informative priors 2 and 5. The estimated optimal doses are close to the prior optimal dose with the prior 2. The estimated optimal doses get better with the prior 5.

We include the event proportions in Figure 3.9. One should immediately notice the great variation for the best intention design when the prior is 2 or 5. The two sources contributing to the variation are dose allocation and response variations. Although the best intention design has little dose allocation variation, for it assigns almost all patients to dose 3, there

are still great variation from the response. Comparing the prior 4 with prior 5, we see that dose allocation contributes a little to the event variation. Designs with $\lambda = 0.5$ 0.8 and 1 perform similarly with different priors.

We conclude that the information-based design with $\lambda = 0.5$ is more robust to different priors with consistent dose allocation weights. Priors are the most influential for the best intention design where unstable dose allocation is observed. The great variation of responses makes it difficult to produce precise parameter estimates. Although the estimation variation can be reduced by imposing informative priors, these priors introduce estimation bias.

Question 4: How does misspecified correlation affect the operating characteristics of designs?

We investigate this problem in scenario 8. We compare prior 1 with prior 3. Note that the support of ρ is $[0, 0.7]$ in prior 1 and $[-0.9, -0.7]$ in prior 3, while the true ρ is 0.5. Thus, prior 1 has correctly specified prior correlation while prior 3 has misspecified prior correlation. First, designs with different priors perform equally well in estimating all paramors except the correlation. This confirms what we observe for information-based designs in Chapter 2. Second, we are interested in how the correlation parameter affects dose allocation. Note that patient assignments depend on correlation through the penalty function. Dose allocation weights are included in Figure 3.10. When $\lambda = 0.8$ or 1, We see that prior 3 leads to more stable dose allocation despite its poor estimates of correlation. Also, designs with prior 3 assign less patients to the toxic dose 4. However, when $\lambda = 1$, we see that prior 3 prior 1 leads to more stable dose allocation. This supports what we observed in Chapter 2, where misspecified correlation leads to unstable information-based designs. The estimated optimal doses are included in Figure 3.10. We see that prior 3 leads to uniformly better learning about the dose-response relationship due to smaller variation. This is due to the more informative correlation prior used in prior 3. We summarize that when the penalty coefficient is close to one, the use of an informative prior for correlation may be beneficial even if it is misspecified, because an informative prior leads to stable dose

allocation. However, this dose not apply to information-based designs with small λ.

3.6 Conclusion

In this chapter, two design philosophies are compared. In the best intention design philosophy, patients are assigned to minimize the ethical penalty. In the information-based design philosophy, patients are assigned to maximize the learning about the dose-response relationship. The ideal outcome in these designs is to estimate an optimal dose that can be passed on to future study. A good estimated optimal dose indicates good learning about the dose-response relationship. Given misspecified priors, we identify the non-convergence issue with best intention designs using both Thall and Cook's (2004) or Dragalin and Fedorov's (2006) ethical penalty functions. In these cases, information-based designs perform better with more stable estimates of the optimal dose. However, information-based designs may have ethical issues by assigning patients to highly toxic doses. To solve this problem, we introduce a penalty coefficient to balance the information goal and ethical goal. It is shown that a compromise between the two goals can be achieved by tuning the penalty coefficient. In addition, we examine the impact of different priors to parameter estimation and dose allocation. We find that informative prior performs well when priors are not misspecified. Non-informative priors results in more robust designs when prior misspecification becomes severe. In our experiment with misspecified correlation, we find that designs with informative prior leads to stable dose allocation.

Table 3.1: The assignment after 300 iterations for Thall and Cook (TC)'s method and Dragalin and Fedorov (DF)'s method.

dose level	p_E	p_T	Desirability	No. assigned by TC	No. assigned by DF
1	0.54	0.15	0.68	300	93
2	0.61	0.18	0.87	0	84
3	0.66	0.19	1.01	0	1
4	0.70	0.21	1.09	0	123

Table 3.2: The parameter estimation after 300 iterations for Thall and Cook (TC)'s method and Dragalin and Fedorov (DF)'s method.

	μ_E	β_{E1}	β_{E2}	μ_T	β_T	α
True	0.50	0.50	0.10	-1.50	0.25	0
TC	0.05	0.40	0.97	-0.90	1.51	-0.57
DF	0.33	0.13	0.11	-1.48	0.17	0.43
Prior	-0.50	0.50	0.10	0.25	3	0

Table 3.3: Summary of Monte Carlo simulations for scenario 1. The mean and standard deviation (sd) of the measures defined across the 5000 complete trials are listed for each method.

Measure	Penalized		Trade-off	
	mean	(sd)	mean	(sd)
Toxicity	14.69	(2.80)	14.51	(2.61)
Efficacy	15.36	(3.33)	14.40	(4.89)
Eff-Tox	7.78	(2.92)	7.24	(2.58)
Desirability	0.012	(0.005)	0.0079	(0.019)
DesHigh	22.04	(5.16)	17.55	(10.03)
Precision	29.48	(2.68)	33.78	(3.57)
Cost	14.13	(0.198)	18.83	(9.12)
Info/Cost	0.0024	(0.0002)	0.0018	(0.0005)

Table 3.4: Summary of Monte Carlo simulations for scenario 2. The mean and standard deviation (sd) of the measures defined across the 5000 complete trials are listed for each method.

Measure	Penalized		Trade-off	
	mean	(sd)	mean	(sd)
Toxicity	10.68	(2.47)	11.03	(2.69)
Efficacy	11.85	(2.96)	12.07	(3.37)
Eff-Tox	7.1	(2.44)	7.34	(3.54)
Desirability	0.016	(0.004)	0.017	(0.006)
DesHigh	16.40	(6.49)	14.9	(10.30)
Precision	57.46	(4.56)	68.83	(7.01)
Cost	7.17	(1.11)	7.40	(0.20)
Info/Cost	0.0025	(0.0004)	0.0020	(0.0002)

Table 3.5: Penalized optimal designs for scenario 1, Gumbel model. In the optimal designs, design points are included in the first row, weights are included in the second row, and penalty is included in the third row.

Scenario	λ	Penalized optimal design	Φ_λ	Φ_0	Φ_1	Eff_0	Eff_1	w_c^*
$\theta = (-2, 0.5, 2, 0.5, 0.5)$								
1	0	$\begin{pmatrix} -6.06 & -1.15 & 2.48 & 7.9 \\ 0.17 & 0.33 & 0.3 & 0.2 \\ 10.98 & 4.57 & 8.85 & 30.06 \end{pmatrix}$	11.38	11.38	22.19	1	0.03	0.33
1	0.2	$\begin{pmatrix} -5.69 & -1.45 & 1.82 & 5.01 \\ 0.17 & 0.38 & 0.35 & 0.11 \\ 10.22 & 4.63 & 7.34 & 17.26 \end{pmatrix}$	11.72	12.2	9.79	0.85	0.35	0.38
1	0.4	$\begin{pmatrix} -4.99 & -1.5 & 1.45 & 4.35 \\ 0.14 & 0.45 & 0.36 & 0.06 \\ 8.85 & 4.65 & 6.64 & 14.72 \end{pmatrix}$	10.87	13.19	7.38	0.7	0.57	0.45
1	0.6	$\begin{pmatrix} -4.46 & 1.1 & 4.03 & -1.46 \\ 0.1 & 0.33 & 0.03 & 0.54 \\ 7.89 & 6.07 & 13.57 & 4.63 \end{pmatrix}$	9.44	14.47	6.09	0.54	0.74	0.54
1	0.8	$\begin{pmatrix} -4.06 & 0.7 & 3.85 & -1.33 \\ 0.05 & 0.26 & 0.01 & 0.68 \\ 7.23 & 5.53 & 12.94 & 4.6 \end{pmatrix}$	7.49	16.56	5.22	0.35	0.88	0.68
1	1	$\begin{pmatrix} -0.98 \\ 1 \\ 4.56 \end{pmatrix}$	4.56	$+\infty$	4.56	0	1	1

Table 3.6: Penalized optimal designs for scenario 2, Gumbel model. In the optimal designs, design points are included in the first row, weights are included in the second row, and penalty is included in the third row.

Scenario	λ	Penalized optimal design	Φ_λ	Φ_0	Φ_1	Eff_0	Eff_1	w_c^*
$\theta = (-2, 1, 2, 0.5, 0.5)$								
2	0	$\begin{pmatrix} -3.44 & -0.77 & 4.5 & 11.39 \\ 0.22 & 0.38 & 0.2 & 0.2 \\ 9.03 & 3.63 & 15.03 & 47.04 \end{pmatrix}$	2.12	2.12	39.01	1	0	0.38
2	0.2	$\begin{pmatrix} -3.16 & -0.87 & 2.39 & 6.21 \\ 0.25 & 0.43 & 0.26 & 0.06 \\ 8.03 & 3.65 & 8.03 & 22.2 \end{pmatrix}$	5.94	4.7	10.87	0.6	0.23	0.43
2	0.4	$\begin{pmatrix} -2.58 & -0.86 & 1.52 & 5.23 \\ 0.25 & 0.42 & 0.3 & 0.03 \\ 6.24 & 3.65 & 5.98 & 17.97 \end{pmatrix}$	6.49	6.35	6.7	0.43	0.54	0.42
2	0.6	$\begin{pmatrix} 0.89 & -2.03 & -0.85 & 4.78 \\ 0.35 & 0.28 & 0.35 & 0.02 \\ 4.86 & 4.94 & 3.65 & 16.13 \end{pmatrix}$	6.24	7.86	5.16	0.32	0.74	0.35
2	0.8	$\begin{pmatrix} -1.38 & 0.26 & 4.47 \\ 0.52 & 0.48 & 0.01 \\ 3.98 & 4.07 & 14.91 \end{pmatrix}$	5.41	9.97	4.27	0.21	0.88	0.52
2	1	$\begin{pmatrix} -0.63 \\ 1 \\ 3.62 \end{pmatrix}$	3.62	$+\infty$	3.62	0	1	1

Table 3.7: Penalized optimal designs for scenario 3, Gumbel model. In the optimal designs, design points are included in the first row, weights are included in the second row, and penalty is included in the third row.

Scenario	λ	Penalized optimal design	Φ_λ	Φ_0	Φ_1	Eff_0	Eff_1	w_c^*
$\theta = (-2, 0.2, 2, 0.5, 0.5)$								
3	0	$\begin{pmatrix} -11.67 & -1.98 & 3.11 & 11.01 \\ 0.13 & 0.34 & 0.38 & 0.15 \\ 10.36 & 4.89 & 12.06 & 45.69 \end{pmatrix}$	12.36	12.36	36.1	1	0	0.34
3	0.2	$\begin{pmatrix} -12.27 & -2.15 & 2.65 & \\ 0.18 & 0.46 & 0.36 & \\ & 10.88 & 4.87 & 10.78 \end{pmatrix}$	12.11	12.92	8.87	0.89	0.45	0.46
3	0.4	$\begin{pmatrix} -10.65 & -2.41 & 1.98 & \\ 0.15 & 0.54 & 0.31 & \\ & 9.49 & 4.85 & 9.15 \end{pmatrix}$	11.08	13.56	7.36	0.79	0.61	0.54
3	0.6	$\begin{pmatrix} -9.38 & -2.58 & 1.4 & \\ 0.11 & 0.64 & 0.25 & \\ & 8.45 & 4.85 & 7.96 \end{pmatrix}$	9.63	14.63	6.3	0.64	0.75	0.64
3	0.8	$\begin{pmatrix} -8.34 & -2.63 & 0.83 & \\ 0.06 & 0.77 & 0.17 & \\ & 7.65 & 4.85 & 7.01 \end{pmatrix}$	7.71	16.56	5.5	0.43	0.88	0.77
3	1	$\begin{pmatrix} -2.49 \\ 1 \\ 4.85 \end{pmatrix}$	4.85	$+\infty$	4.85	0	1	1

Table 3.8: Penalized optimal designs for scenario 4, Gumbel model. In the optimal designs, design points are included in the first row, weights are included in the second row, and penalty is included in the third row.

Scenario	λ	Penalized optimal design	Φ_λ	Φ_0	Φ_1	Eff_0	Eff_1	w_c^*
$\theta = (2, 0.5, -2, 0.5, 0.5)$								
4	0	$\begin{pmatrix} -7.9 & -2.48 & 1.15 & 6.06 \\ 0.2 & 0.3 & 0.33 & 0.17 \\ 25.42 & 18.62 & 23.71 & 41.39 \end{pmatrix}$	11.38	11.38	33.17	1	0.05	0.3
4	0.2	$\begin{pmatrix} -7.1 & -2.5 & 0.98 & 2.89 \\ 0.2 & 0.42 & 0.28 & 0.09 \\ 23.74 & 18.61 & 23.31 & 28.75 \end{pmatrix}$	14.23	12.03	23.03	0.88	0.41	0.42
4	0.4	$\begin{pmatrix} -6.27 & -2.54 & 1.03 & \\ & 0.19 & 0.52 & 0.29 \\ & 22.14 & 18.61 & 23.42 \end{pmatrix}$	16.13	12.8	21.13	0.75	0.6	0.52
4	0.6	$\begin{pmatrix} -5.65 & -2.63 & 0.54 & \\ & 0.17 & 0.62 & 0.21 \\ & 21.08 & 18.59 & 22.32 \end{pmatrix}$	17.59	13.87	20.07	0.61	0.74	0.62
4	0.8	$\begin{pmatrix} -5.1 & -2.71 & 0.11 & \\ & 0.12 & 0.76 & 0.12 \\ & 20.26 & 18.58 & 21.46 \end{pmatrix}$	18.57	15.87	19.24	0.41	0.88	0.76
4	1	$\begin{pmatrix} -2.81 \\ 1 \\ 18.58 \end{pmatrix}$	18.58	$+\infty$	18.58	0	1	1

Table 3.9: Design settings for Examples 1 and 2. The efficacy, toxicity, "efficacy and no toxicity" probabilities and ethical penalty $c(x, \theta)$ are included for the prior and true models.

		Prior				True			
level	log(dosage)	p_E	p_T	p_{10}	penalty	p_E	p_T	p_{10}	penalty
				Example 1					
1	-1	0.62	0.10	0.58	1.9	0.62	0.08	0.59	1.9
2	-0.5	0.73	0.32	0.54	2.8	0.73	0.11	0.67	1.7
3	0.5	0.88	0.90	0.09	113	0.88	0.18	0.73	1.7
4	1	0.92	0.98	0.02	1921	0.92	0.23	0.72	1.8
				Example 2					
1	-3	0.38	0.00	0.38	2.67	0.45	0.08	0.41	2.7
2	0	0.92	0.08	0.85	1.27	0.60	0.27	0.41	3.3
3	1	0.97	0.18	0.79	1.55	0.65	0.38	0.38	4.3
4	3	1.00	0.62	0.38	7.05	0.73	0.62	0.25	10.5

Table 3.10: The root of mean square error for the MAP estimators of β_E and β_T. The sample sizes are 50 and 200.

		RMSE		
Parameter	n	$\lambda = 1$	$\lambda = 0.8$	$\lambda = 0.5$
$\beta_E = 0.2$	50	0.29	0.18	0.16
	200	0.22	0.11	0.07
$\beta_T = 0.5$	50	0.16	0.24	0.28
	200	0.10	0.11	0.11

Table 3.11: The priors used in the adaptive designs in Section 3.5.

Prior	α_E	β_E	α_T	β_T	ρ
1	$N(-2.5, 1)$	$N(1, 1)^+$	$N(2.5, 1)$	$N(1, 1)^+$	$U[0, 0.7]$
2	$N(-2.5, 0.2^2)$	$N(1, 0.2^2)^+$	$N(2.5, 0.2^2)$	$N(1, 0.2^2)^+$	$U[0, 0.7]$
3	$N(-2.5, 1)$	$N(1, 1)^+$	$N(2.5, 1)$	$N(1, 1)^+$	$U[-0.9, -0.7]$
4	$U[-5, 0]$	$U[0, 4]$	$U[0, 5]$	$U[0, 4]$	$U[0, 0.7]$
5	$U[-3.1, -1.9]$	$U[0.4, 1.6]$	$U[1.9, 3.1]$	$U[0.4, 1.6]$	$U[0, 0.7]$

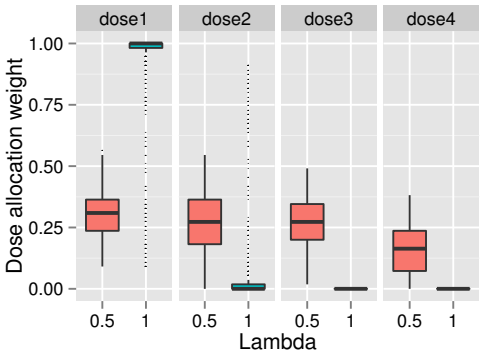

(a) Dose allocation weights

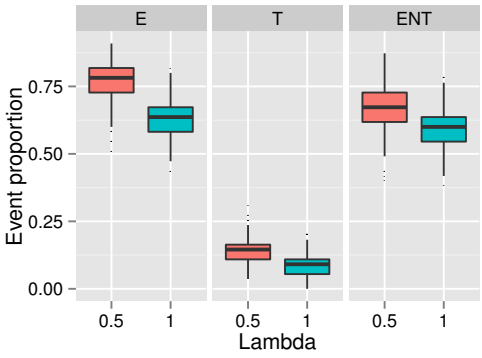

(b) Event proportion

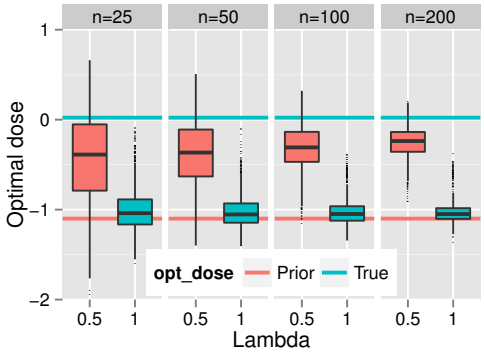

(c) Optimal Doses

Figure 3.1: Operating characteristics of best intention design ($\lambda = 1$) and and information-based design ($\lambda = 0.5$) in Example 1. The events E, T and Et denote efficacy, toxicity and "efficacy and no toxicity" respectively.

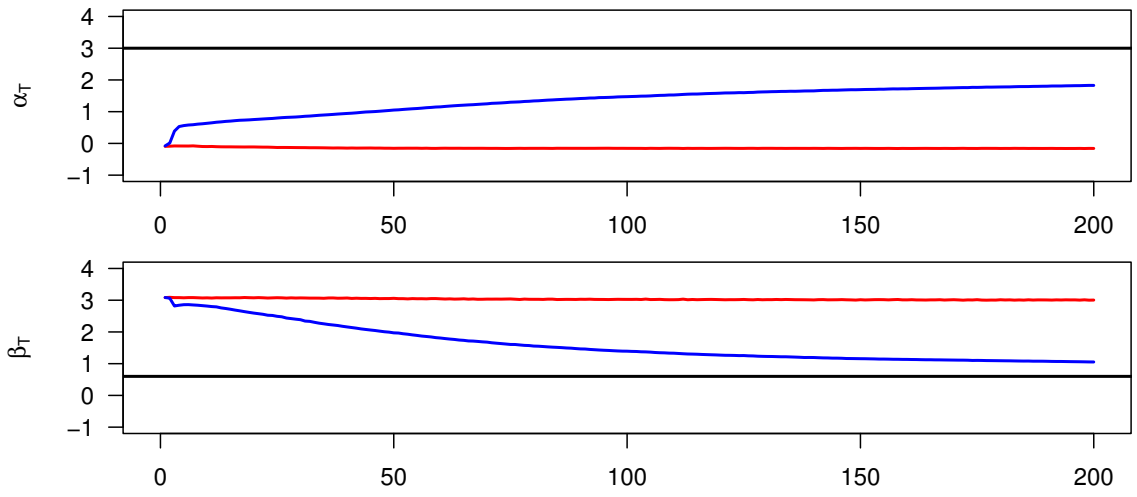

Figure 3.2: The mean MAP estimates across 1000 simulations for the information-based design ($\lambda = 0.5$) and the best intention design ($\lambda = 1$). The estimates of α_T and β_T in the best intention design do not converge to the true parameters.

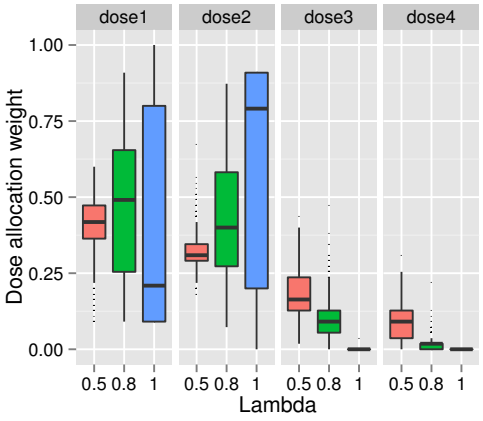

(a) Dose allocation weights

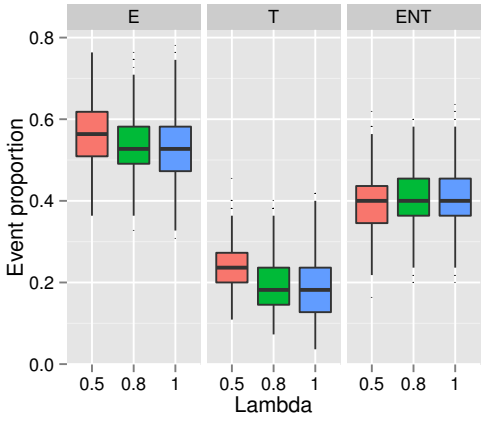

(b) Event proportion

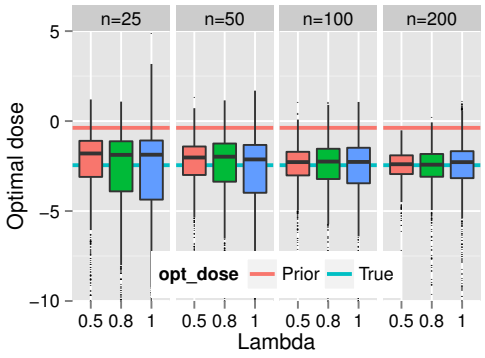

(c) Optimal Doses

Figure 3.3: Operating characteristics of the best intention design ($\lambda = 1$) and and information-based designs ($\lambda = 0.5$ and 0.8.) in Example 2. The events E, T and ENT denote efficacy, toxicity and "efficacy and no toxicity" respectively.

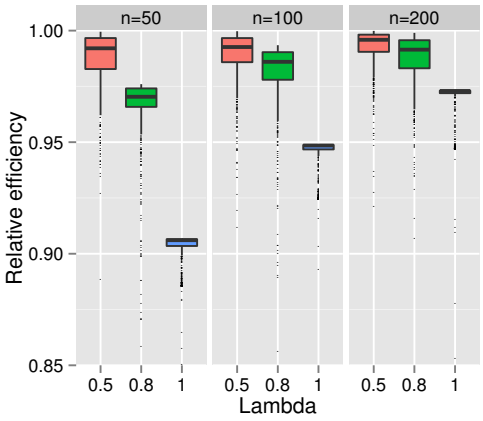

(a) Efficiency of penalized information, scenario 1.

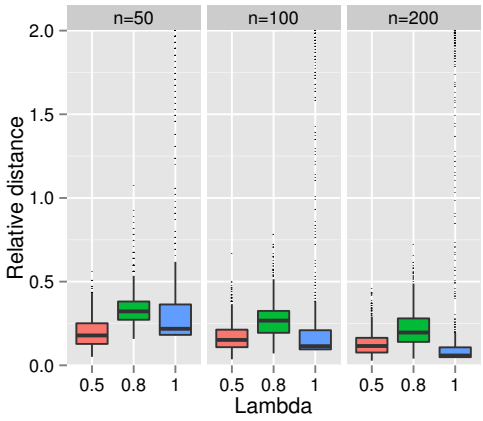

(b) L^1 distance of dose allocation weights, scenario 1.

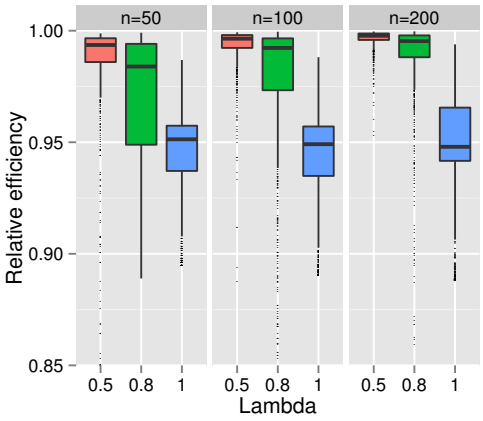

(c) Efficiency of penalized information, scenario 4.

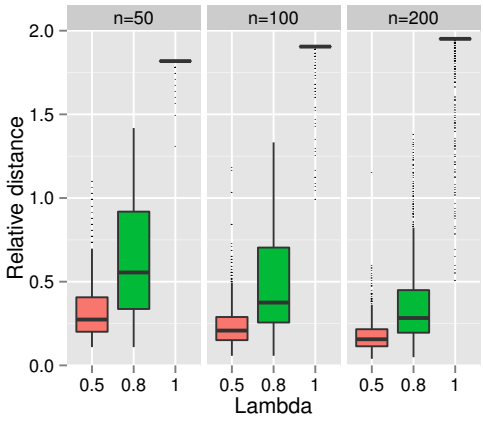

(d) L^1 distance of dose allocation weights, scenario 4.

Figure 3.4: Convergence measure for scenario 1 and 4. Relative efficiency is based on the optimal Φ_λ and the criterion value in simulations. The distance measure is calculated as the L^1 distance between the dose allocation weight of the optimal design and those in simulations.

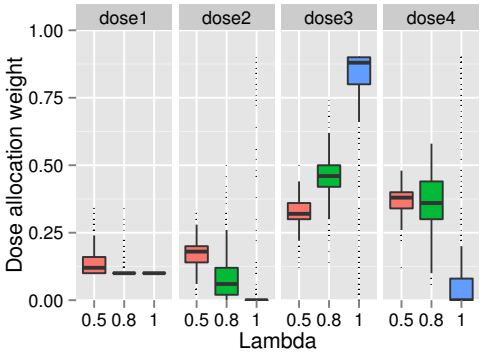

(a) Dose allocation for scenario 1.

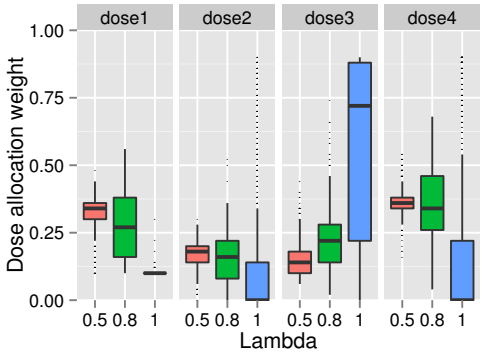

(b) Dose allocation for scenario 4.

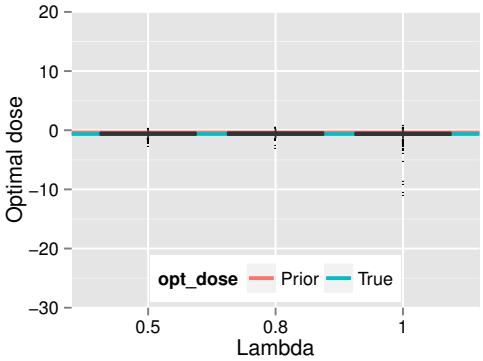

(c) Optimal doses for scenario 1.

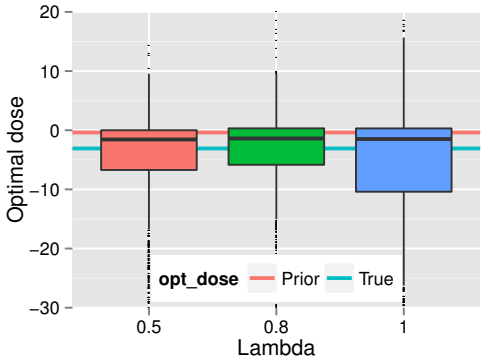

(d) Optimal doses for scenario 4.

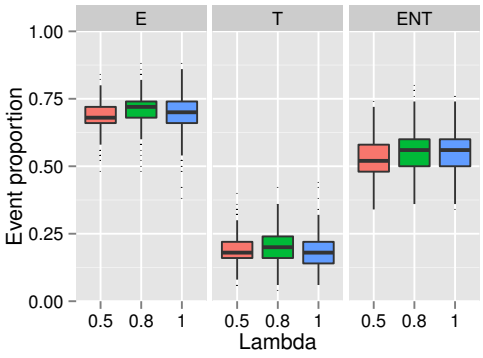

(e) Event proportions for scenario 1.

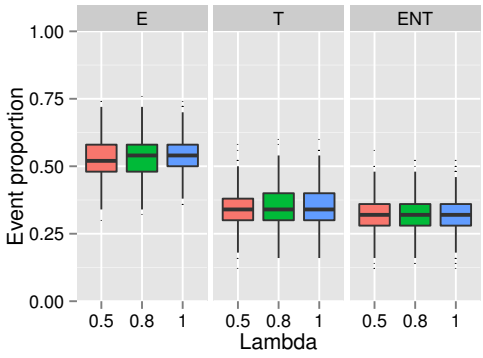

(f) Event proportions for scenario 4.

Figure 3.5: Operating characteristics for scenario 1 and 4.

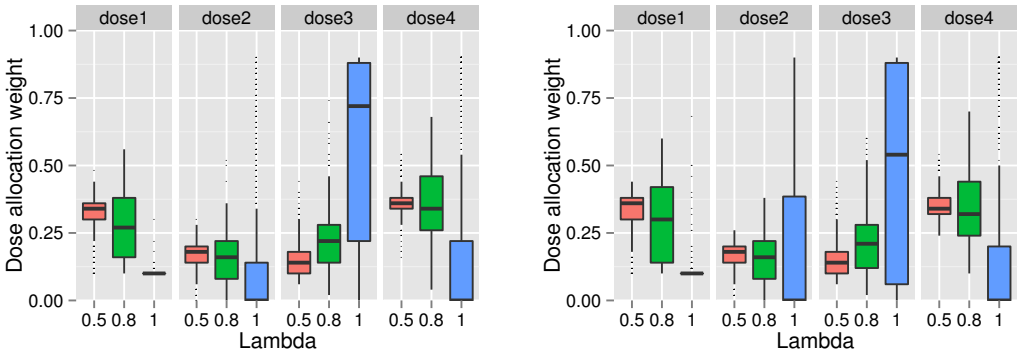

(a) Prior 1

(b) Prior 4

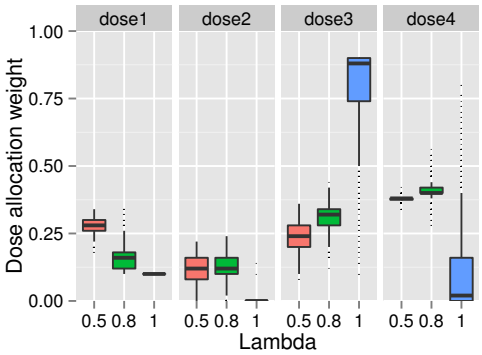

(c) Prior 2

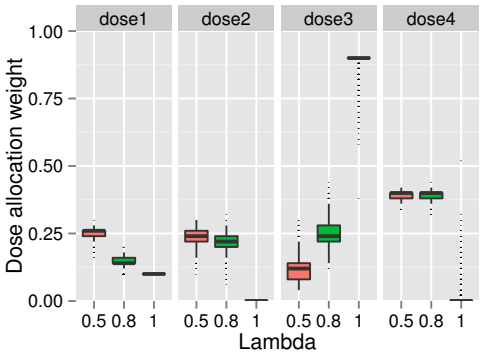

(d) Prior 5

Figure 3.6: Dose allocation weights in scenario 4 with priors 1, 2, 4 and 5.

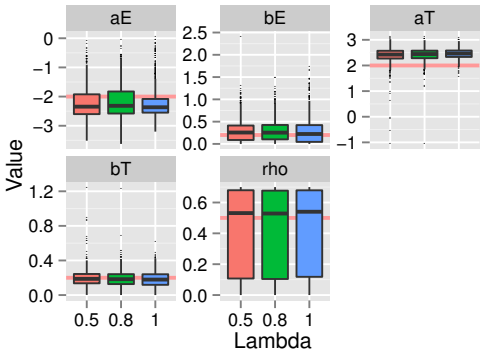

(a) Prior 1

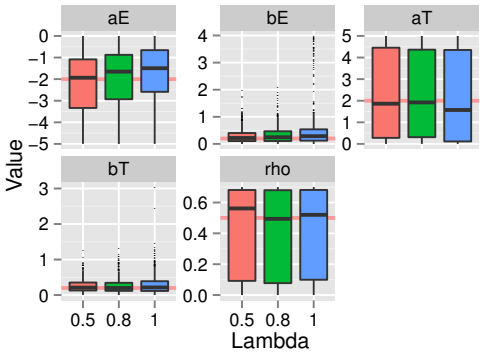

(b) Prior 4

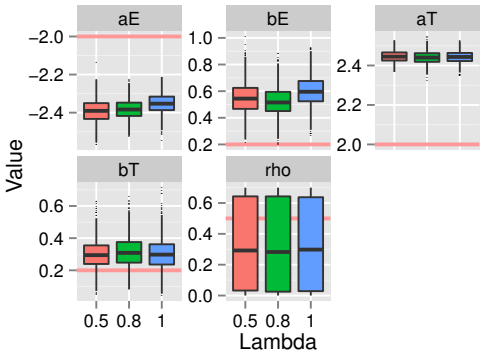

(c) Prior 2

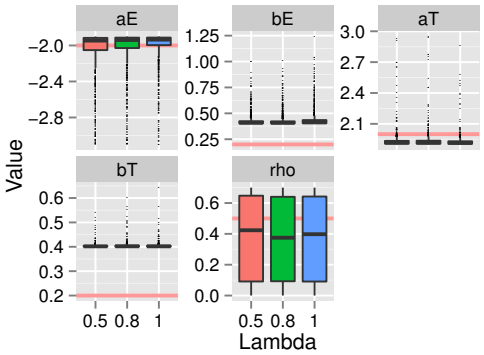

(d) Prior 5

Figure 3.7: Parameter estimates in scenario 4 with priors 1, 2, 4 and 5.

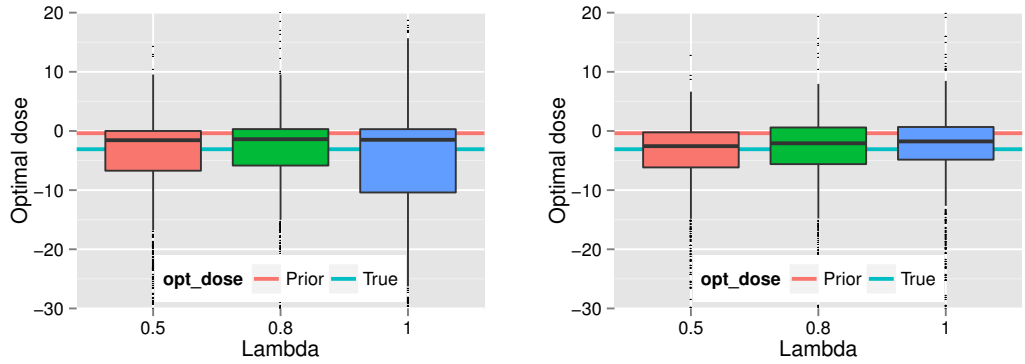

(a) Prior 1

(b) Prior 4

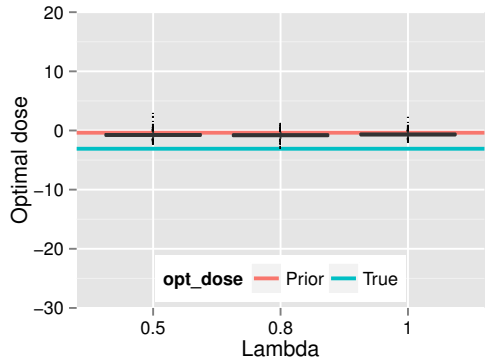

(c) Prior 2

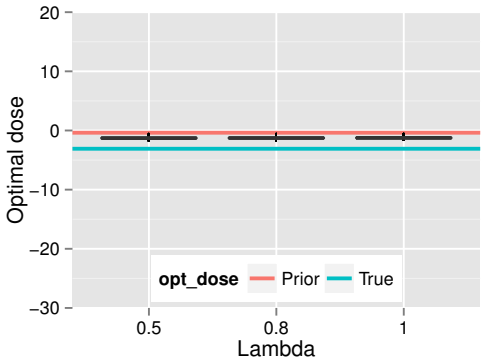

(d) Prior 5

Figure 3.8: Estimated optimal dose in scenario 4 with priors 1, 2, 4 and 5.

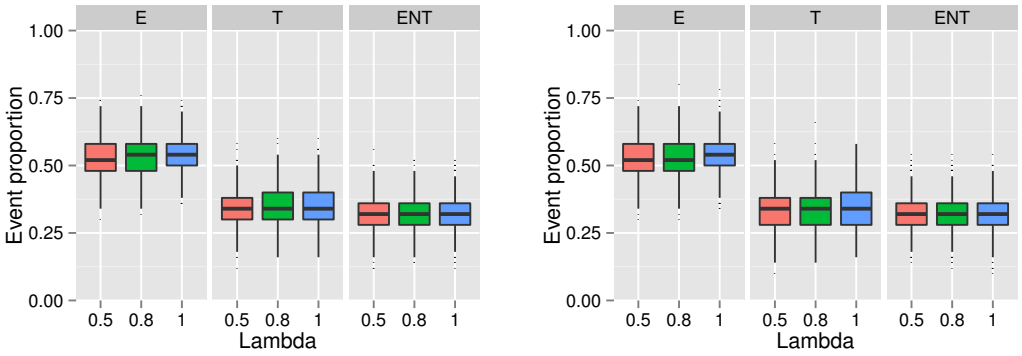

(a) Prior 1

(b) Prior 4

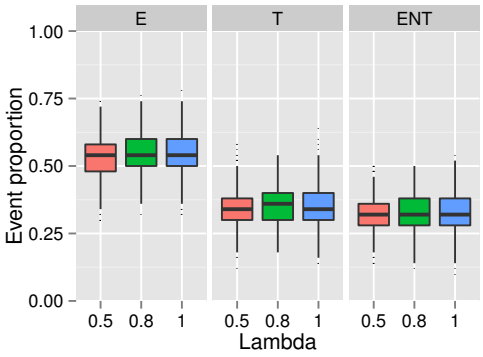

(c) Prior 2

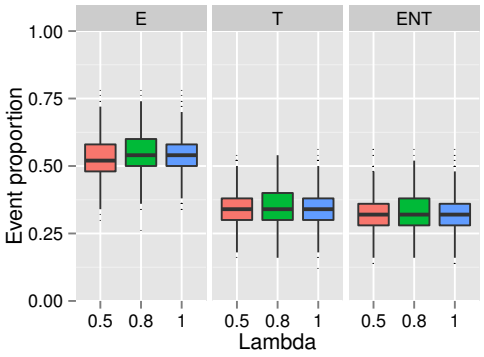

(d) Prior 5

Figure 3.9: Event proportions in scenario 4 with priors 1, 2, 4 and 5.

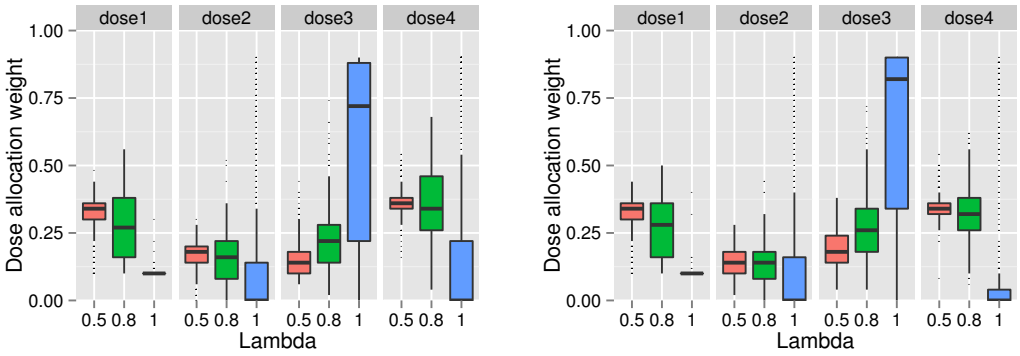

(a) Dose allocation for prior 1.

(b) Dose allocation for prior 3.

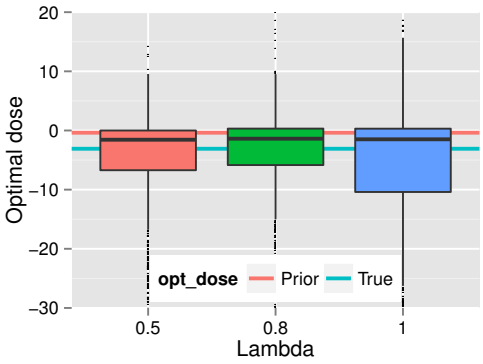

(c) Optimal doses for prior 1.

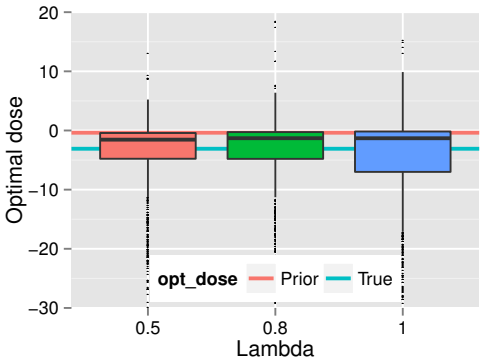

(d) Optimal doses for prior 3.

Figure 3.10: Operating characteristics for priors 1 and 3.

Chapter 4: Stopping Rules

In adaptive designs, traditional stopping rules are used to stop a seemingly unsuccessful trial or a trial with strong enough evidence for efficacy earlier than planned. The rules are devised to protect the patients from excessive toxicity and insufficient efficacy. For example, Thall and Cook (2004) introduced rules based the posterior probability of toxicity and efficacy. The design philosophy of adaptively updating the dose-response curve leads us to consider a different type of stopping rule. Also, the concept of trade-off between efficacy and toxicity leads us to consider stopping a trial when the desired trade-off is not attained.

In this chapter, we propose to stop a trial when there is enough learning about the dose-response relationship or when the trial exceeds an efficacy and toxicity trade-off threshold. We examine stopping rules in simulation.

4.1 Information stopping rules

One information stopping rule is proposed by Zhu (2012):

$$\max_j \left| \frac{\hat{\theta}_j^{(n)} - \hat{\theta}_j^{(n-1)}}{\hat{\theta}_j^{(n-1)}} \right| < \delta_0,$$

where $\hat{\theta}_j^{(n)}$ is the jth component of parameter estimate with n patients and δ_0 is a threshold. In our case, we may have near-zero estimates, therefore we propose to stop a trial when the maximum parameter ratio is small.

$$\textbf{Ratio criterion:} \qquad \max_j \left| \frac{\hat{\theta}_j^{(n)} - \hat{\theta}_j^{(n-1)}}{a_j} \right| < \delta_0, \tag{4.1}$$

where a_j's are preselected parameter weights.

Another information stopping rule is based on the information criterion Φ_0. We stop a trial with n patients if the information criterion is small.

$$\textbf{Information criterion:} \qquad \Phi_0(\xi_n; \hat{\theta}^{(n)}) - \log pn < m_0, \tag{4.2}$$

where $\Phi_0(\xi_n)$ is the information criterion for design with n patients, p is number of parameters, and m_0 is an information threshold. We call the left hand side the the above inequality the scaled information criterion.

4.1.1 Example with parameter ratio stopping rules

To illustrate the stopping rules, we examine the dynamics of the adaptive design for scenario 4. We set $a = (2.5, 1, 2.5, 1, 1)$ whose first four components are the same as the prior means while the last component is amplified. The history of the estimates of α_E is plotted in Figure 4.1 (a). We see that the estimates scatter more in the beginning than in the end, thus sample size will be reduced if we stop a trial when its parameter estimates stabilize. The history of the maximum parameter ratio of the first three simulations is plotted in Figure 4.1. The original ratio is messy due to random noise. We calculate the moving average of the ratios and plotted them in Figure 4.1 (c).

We apply the stopping rules in Table 4.1. We note that some trials may attain the stopping criteria in the beginning, therefore we also introduce the minimum sample size $t_{\min}$ beyond which the stopping rule is applied. Also we introduce the maximum sample size $t_{\max}$ to force the trial to stop when the stopping rule is not triggered. The stopping threshold δ_0 is chosen so that the the average stopping sample size is maintained at 55. The simulated stopping times are summarized in Figure 4.2 (a). Rules R01 and R03 are not restricted so they have extreme stopping times, either too large or too small. R05 and R06 are fixed design with sample sizes 55 and 205.

As we mentioned earlier, the goal of early phase trial is to learn about the dose-response

relationship. The accuracy of this relationship can be measured by the knowledge of the estimated optimal dose for given parameter estimates. The closer the estimated optimal dose to the true one, the better the dose-response relationship is learned. We summarize the estimated optimal dose of trials with different stopping rules in Figure 4.2 (c). In the optimal dose plot, we see that R03 has a smaller interquartile range than R05. The average design with stopping rule R03 performs better than the fixed sample size design R05. This is likely to be caused by increased sample size of poor trials in R05. Admittedly, R06 with a fixed sample size of 205 performs the best with the smallest interquartile range. It is the only one that differentiates the prior optimal dose at -0.376 from the true optimal dose at -3.086. Comparing R01 with R03 or R02 with R04, one can see that the smoothing of the ratio curves leads to smaller interquartile range and thus better estimates about the optimal dose. Thus, a better stopping rule could be developed if one has better smoothing. The introduction of minimum and maximum stopping time has mixed effects. The restricted R02 performs better than the unrestricted R01, while the restricted R04 performs worse than the unrestricted R03.

Interestingly, these results have little to do with the individual parameter estimates, as shown in Figure 4.2 (b). For example, all the designs R01–R05 have similar estimate interquartile ranges for individual parameters. However, the (nonlinear) combination of these estimates gives rise to different optimal dose estimates.

Another question is how the stopping rule affects dose allocation. We summarize the dose allocation proportions in Figure 4.2 (d). The dose allocations of the trials with stopping rules resemble those of their fixed sample size counterpart. The only difference is that the latter, R05, has less extreme values than the former. The proportions of R06 are significant different from the rest. This is caused by better parameter estimates from the data rather than from the misspecified prior knowledge.

We conclude from the above observations that stopping rules change the operating characteristics. With a good stopping rule such as R03, better dose-response relationship is learned for a given average sample size across simulations. The price is some extreme dose

allocations.

4.1.2 Example with information criterion stopping rules

We now turn to stopping rules based on the information criterion given in (4.2). We examine the same cases as the last example. The stopping rules are detailed in Table 4.2. The thresholds are chosen so that the average sample sizes are about 50 across simulations. These stopping rules are expected to have different properties from last set of stopping rules, because they emphasize the interplay between parameter estimates and dose-allocation.

We summarize the most interesting results in Figure 4.3. First, moving averages of Φ_0 are calculated within windows of size five for the first three simulations (Figure 4.3 (a)). The mean information criterion decreases as the sample size increases. Therefore, the stopping rule will be effective when enough information is learned from the data. After applying the stopping rules in Table 4.2, we include the distributions of stopping time in Figure 4.3 (b). The distributions for the unrestricted rules M01 and M03 are highly skewed. They have long right tails which almost reach the maximum 205. It indicates slow convergence rate for the criterion. However, their medians are around 35. Thus half of the trials do not have sufficient learning about the dose-response relationship, which is shown in Figure 4.3 (c). The estimated optimal doses of M01 and M03 are dominated by the prior information. The moving average version M03 has better learning than M01. All the other stopping rules perform better than these two. Although M02 and M04 have smaller interquartile ranges than the fixed M05, it is difficult to decide which are better, because the median of M05 is closer to the true optimal dose than the other two. Comparing with the last example, there are fewer extreme values in this example. This may be due to less extreme dose allocations shown in Figure 4.3 (d). For example, the most extreme proportion at dose 1 is 0.7 in this example versus 0.8 in last example. There are less extreme values above the medians at dose 4 in this example than in last example. However, the span of the proportions in this example are larger than the previous example, especially for the unrestricted rules M01 and M03.

We conclude that the information criterion stopping rules result can reduce extreme dose allocations and thus reduce extreme estimates about the optimal doses. However, the price is large variation of dose-allocation.

4.2 Ethical stopping rule

In Bayesian sequential designs, stopping rules are introduced to stop a trial when toxicity is too high. We argued in the last section that efficient protection cannot be obtained when parameter estimates are poor. Therefore, it is a good practice to combine ethical stopping rules with information rules. In this section, we examine the operating characteristics of stopping rules with ethical stopping and with combined stopping.

We propose to use the following the MAP criterion for ethical stopping in a trial with n patients.

$$\textbf{MAP criterion:} \qquad \min_{1\leq i\leq k} \Phi_1(\delta_{x_i}; \hat{\theta}^{(n)}) > e_0, \tag{4.3}$$

where $\Phi_1(\delta_{x_i})$ is the ethical penalty at single dose x_i and e_0 is an ethical threshold.

A commonly used posterior ethical stopping rule can be proposed.

$$\textbf{Posterior criterion:} \qquad Pr_{\Theta}(\min_{1\leq i\leq k} \Phi_1(\delta_{x_i}; \theta^{(n)}) > e_1 | \mathcal{D}_n) > p_1, \tag{4.4}$$

where δ_{x_i} is a single dose allocation at x_i, and $\theta^{(n)}$ is a sample from the posterior distribution supported on Θ. The second term incorporates the belief that the longer the trials, the great exposure of patients to drugs and thus greater ethical costs.

4.2.1 Example with ethical MAP threshold

We illustrate the ethical stopping rule (4.3) in the same case as the last example. We use three λ's 1, 0.8 and 0.5. We assume the clinician decide that the maximal tolerated efficacy-toxicity trade-off is 4. However, because the true trade-offs at all doses are greater than 4,

the smaller the average stopping time, the better a design is. We summarize the stopping times for the three designs in Figure 4.4(a). Surprisingly, there is not much difference among the three designs. Designs with λ's 0.8 and 0.5 perform equally well because they have similar distribution for stopping times. The design with $\lambda = 1$ performs slightly worse than the other two due to a larger interquartile range. Similar conclusion can be inferred in Figure 4.4(b). All three designs have almost identical distribution for estimated optimal dose, except that the design with $\lambda = 1$ has a median that is closer to the prior belief. We consider dose allocation shown in Figure 4.4(c). Recall that the last dose is the most desirable given the prior information while it is the least desirable given the true parameter. We observe that most patients are assigned to the truly most desirable dose – dose 1. The second most assigned patients for the information-based design is dose 4 while it is dose 3 for the best intention design. We also observe greater and greater variation among assignments as λ increases. Because of the dose-allocation, we observe that the best intention design has the least proportions for efficacy, toxicity and efficacy and no toxicity.

If we allow the ethical threshold to vary in (4.3), and calculate the mean stopping sample size for each threshold and each design, we will obtain the threshold characteristics curves in Figure 4.3(d). The curves are sigmoid shaped. When the sample size is smaller than 25, the prior dominates the estimates when all the ethical penalties are thought to be below 4. When the threshold increases, less and less MAP's will satisfy the criterion because the true minimum ethical penalty around 4.48. This can be seen in the plot that after the critical value is greater than 5, the increasing curve begins to decelerate. However, the three designs cannot be distinguished in the plot.

In conclusion, the MAP stopping rules can stop a trial early when all doses are not desirable. When the trials are stopped, the three designs with λ's 1, 0.8, and 0.5 have similar average sample sizes. At stop times, the information-based design demonstrates better dose-response learning than the best intention design because of the better median estimated optimal dose.

4.2.2 Example with ethical posterior threshold

We examine the last stopping rule for the same scenario used in previous example. Here we set the ethical threshold e_1 to be 4 and allow the posterior probability threshold p_1 to vary. This produces the stopping time characteristic curve in Figure 4.5(a). The three curves for different designs are distinguishable. The best intention design performs the worst and requires the largest sample size. Surprisingly, the information-based design with $\lambda = 0.8$ performs better than that with $\lambda = 0.5$.

We set the posterior threshold to be 0.5 and plot the distributions of stopping time in Figure 4.6(a). The design with $\lambda = 0.8$ performs the best with the smallest interquartile range. The estimated optimal doses at stopping times are summarized in Figure 4.6(c). The design with $\lambda = 0.5$ performs the best with a median that is closest to the Turk optimal dose. It is followed by the design with $\lambda = 0.8$. The best intention design performs still the worst with a median that is close to the prior optimal dose and the largest distribution span. The dose allocation is shown in Figure 4.6(e). Still, the best intention design performs the worst by assign most patients to the dose 3. It also has the largest allocation variation. The information-based design performs similar in terms of median proportions of assigned patients on each dose. However, the design with $\lambda = 0.5$ has the smallest variation. The design with $\lambda = 0.8$ provides better protection of patients by assigning more patients to dose 2 and 3 instead of dose 4.

To demonstrate the influence of the correlation parameter on design operating characteristics, we apply the same stopping rule in a different scenario. In this scenario, the correlation parameter is misspecified and restricted in $[-0.9, -0.7]$. We call this scenario U2 and the previous scenario U1. The stropping time characteristic curve for U2 is included in Figure 4.5(b). The difference between U1 and U2 are seen immediately in the plot. Because of the misspecification, it requires more than 100 patients for U2 to stop while it requires about 50 patients in U1. Also, the curve of $\lambda = 1$ is more distinguishable in U2 than in U1.

To obtain comparable sample sizes, we set the posterior threshold to be 25 for U2, so that both U1 and U2 have sample sizes around 50. We summertime the stopping times in

Figure 4.6(b). We observe that there are much grater variation for the design with $\lambda = 1$ in U2 than in U1. The information-based designs suffer less from the misspecified correlation when one compares U1 with U2 and show more robustness. Also, the information-based designs performs much better than the best intention designs. The estimated optimal dose plot in Figure 4.6(b) shows that the U2 has better estimated optimal doses than U1. This is probably due to the more spread out stopping time in U2 and the fact that correlation misspecification keeps the orders of Φ_1. The design with $\lambda = 0.5$ performs slightly better than the other two designs in U2. In the dose allocation plot in Figure 4.6(f), there is slightly difference between the designs in U1 and U2. For example, the variations in U2 are greater than those in U1. Also, there is a small amount of patients assigned on the highest dose in U1 while there is no assignments on that dose in U2.

4.3 Conclusion

We examine different stopping rules in this Chapter. We notice the discrepancy between individual parameter estimates and the overall estimated represented by the estimated optimal dose. We observe that appropriate stopping rules improve operating characteristics of designs. With information-based stopping rules, a better learning about the dose-response relationship is obtained without increasing the average sample size. By using a ethical stopping rule, an unsuccessful trial can be stopped early. We also observe that information-based designs allow better learning about the dose-response relationship than best intention designs. To assess the performance of different designs with stopping rule, we propose to use a characteristic curve. Using this curve, we observe that information-based design may performs better by stopping earlier when no dose is desirable.

Table 4.1: Stopping rule settings with ratio criterion (4.1).

Rule	moving average	δ_0	$t_{\min}$	$t_{\max}$	average sample size
R01	No	0.032	5	205	55.2
R02	No	0.032	25	75	54.7
R03	Yes	0.069	5	205	55.48
R04	Yes	0.068	25	75	55.12
R05	Sample size fixed at 55				
R06	Sample size fixed at 205				

Table 4.2: Stopping rule settings with information criterion (4.2).

Rule No.	moving average	m_0	$t_{\min}$	$t_{\max}$	average sample size
M01	No	2.7	5	205	55.41
M02	No	2	25	75	54.45
M03	Yes	4.1	5	205	55.42
M04	Yes	2.7	25	75	54.54
M05	Sample size fixed at 55				
M06	Sample size fixed at 205				

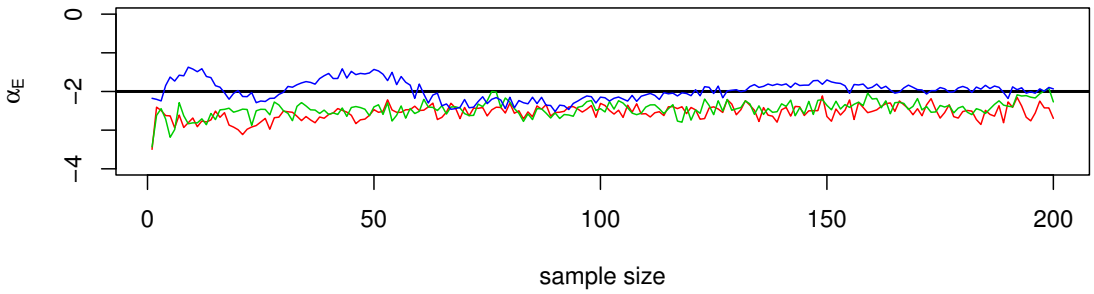

(a) Parameter estimates

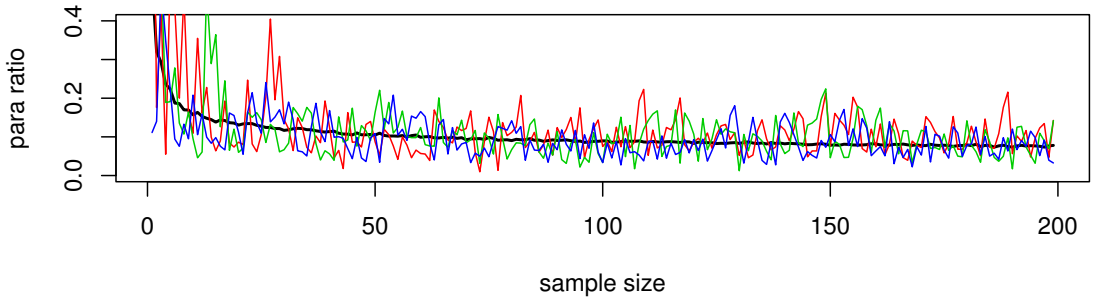

(b) Parameter ratios

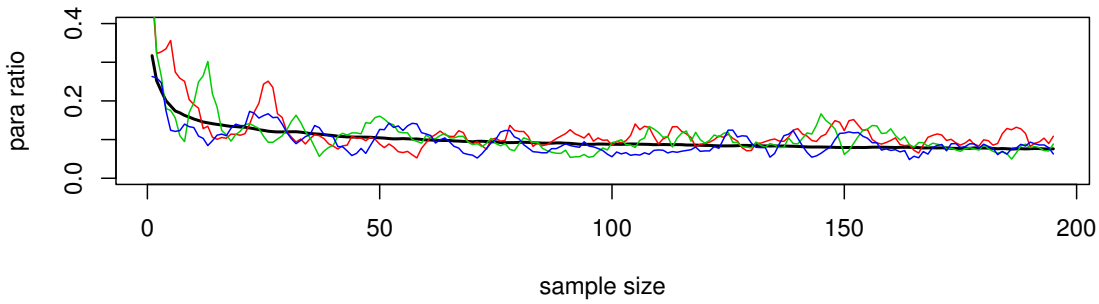

(c) Parameter ratios moving average

Figure 4.1: The history of parameter ratios for the first three simulations. The original parameter estimates and ratios are plotted in (a) and (b) and the moving average ratios in (c) are calculated in a window of size five.

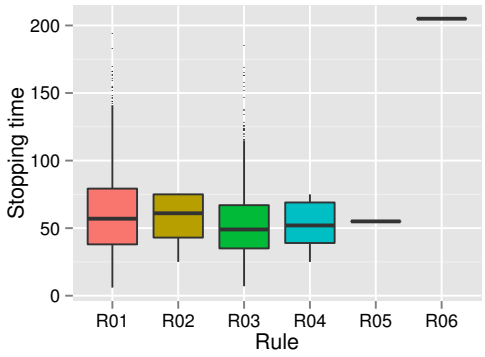

(a) Stopping time

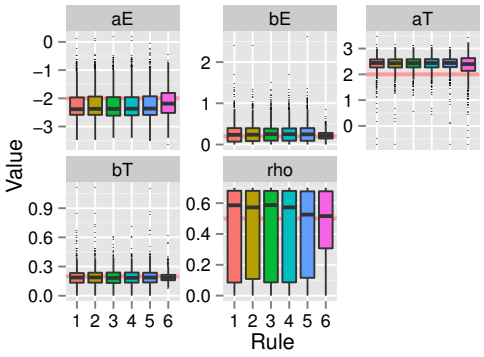

(b) Parameter estimates

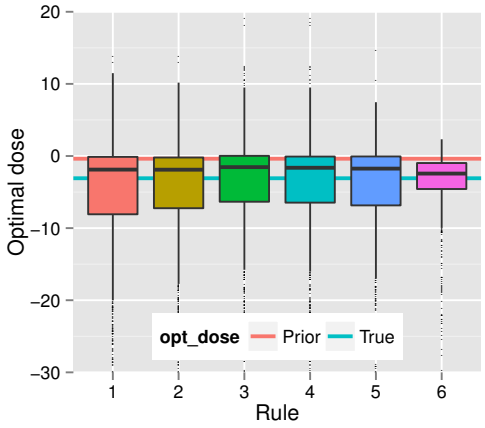

(c) Optimal dose estimates

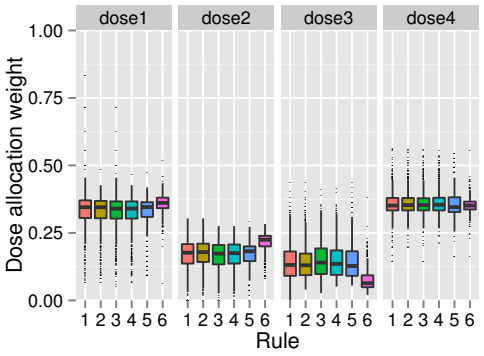

(d) Dose allocation

Figure 4.2: Operating characteristics with parameter ratio stopping rules in Table 4.1.

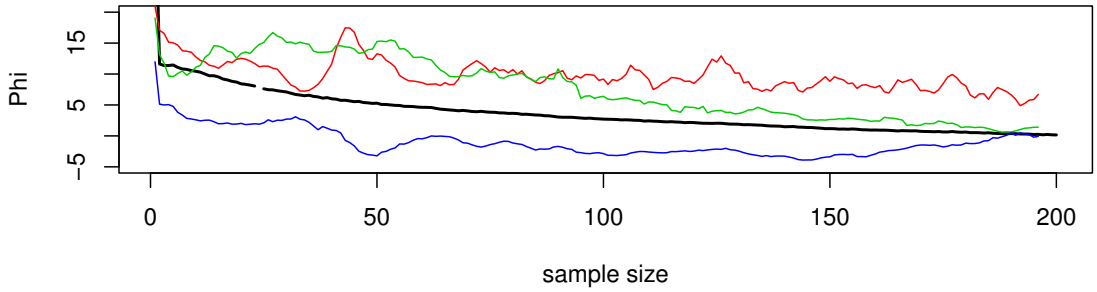

(a) History of scaled Φ_0

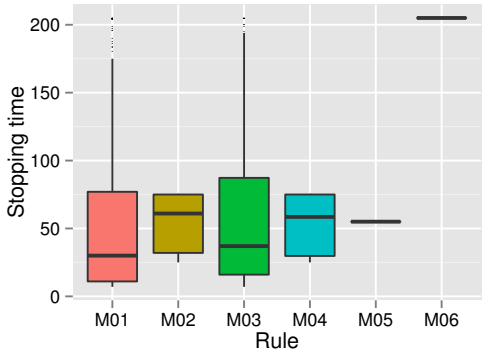

(b) Stopping time

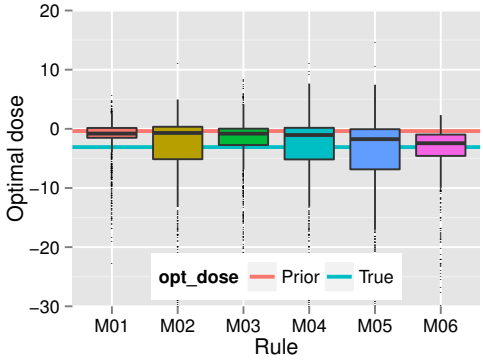

(c) Optimal dose estimates

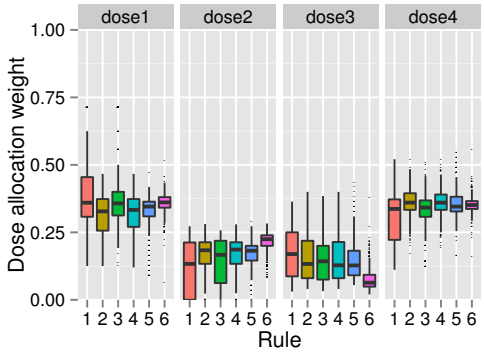

(d) Dose allocation

Figure 4.3: Operating characteristics with information criterion stopping rules in Table 4.2.

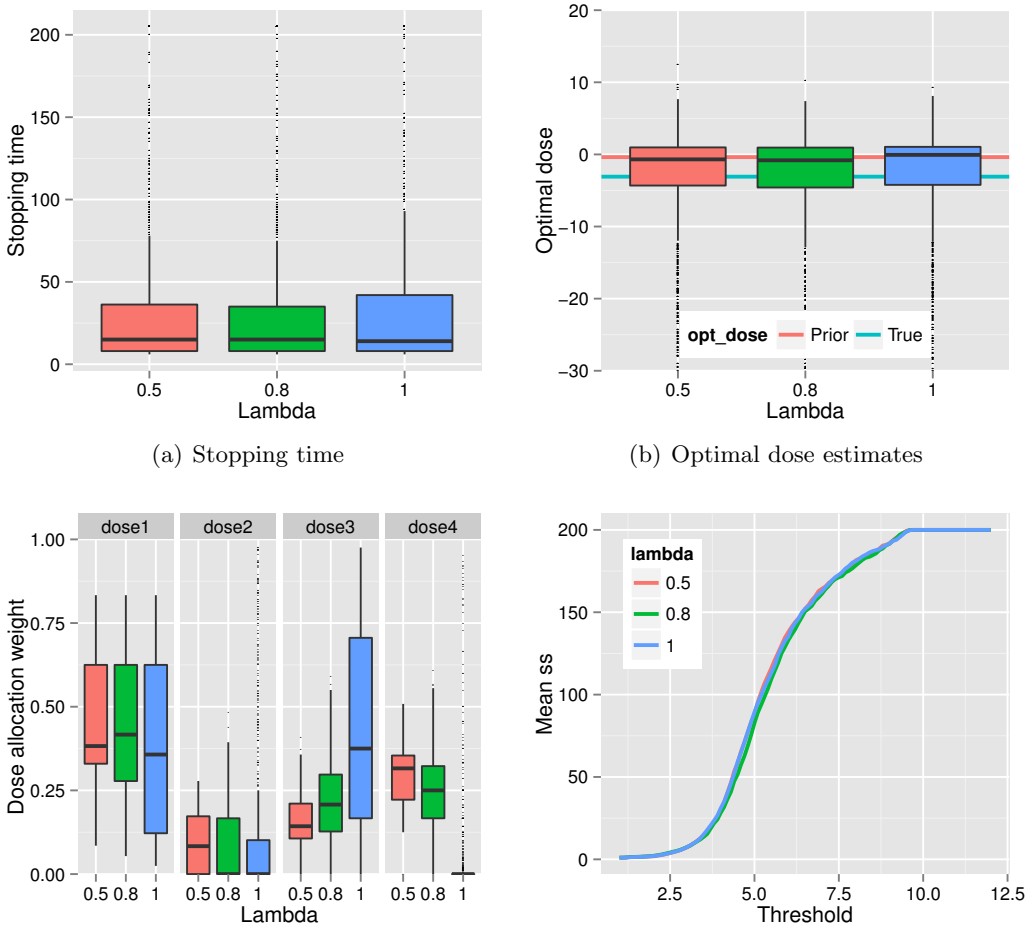

(a) Stopping time

(b) Optimal dose estimates

(c) Dose allocation

(d) Mean sample size vs threshold

Figure 4.4: Operating characteristics with ethical stopping rules based on MAP criterion in (4.3).

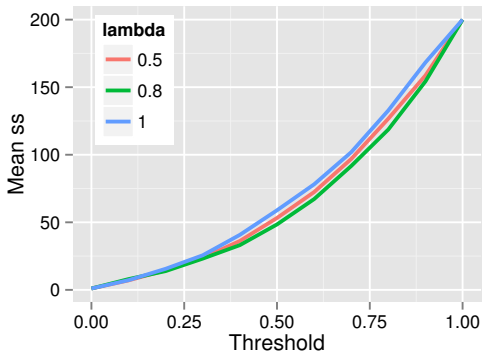

(a) Mean sample size vs threshold (U1)

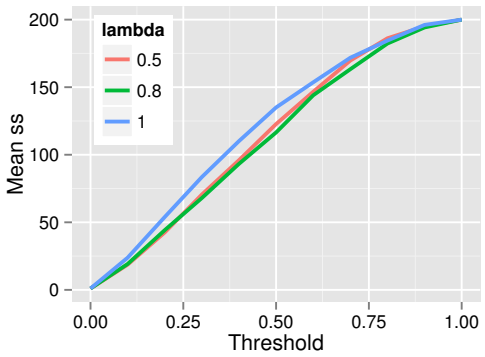

(b) Mean sample size vs threshold (U2)

Figure 4.5: Characteristic curve with ethical stopping rules based on posterior criteria in (4.4). U1 stands for the scenario with correctly specified correlation. U2 is the scenario with misspecified correlation.

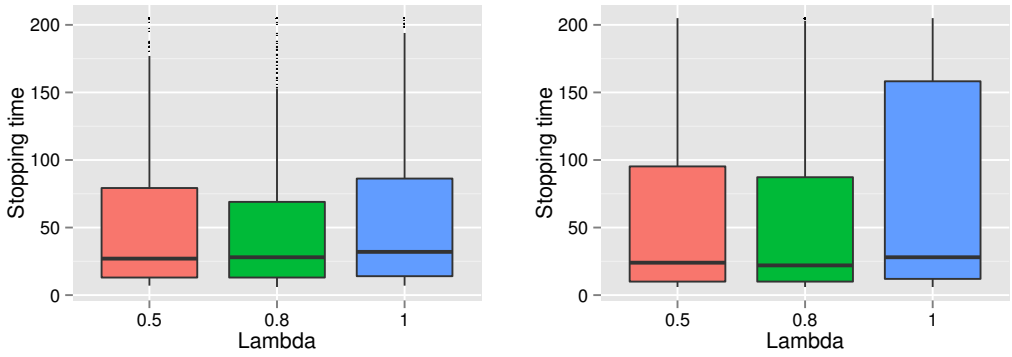

(a) Stopping time (U1)

(b) Stopping time (U2)

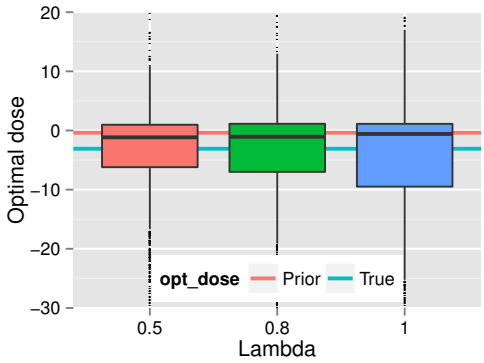

(c) Optimal dose estimates (U1)

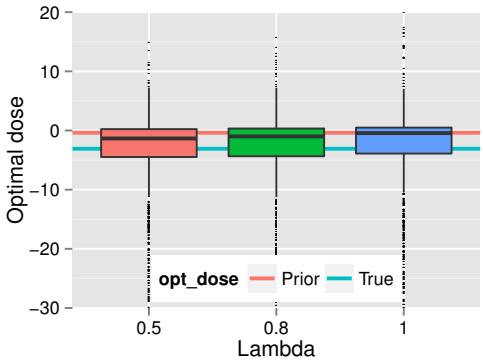

(d) Optimal dose estimates (U2)

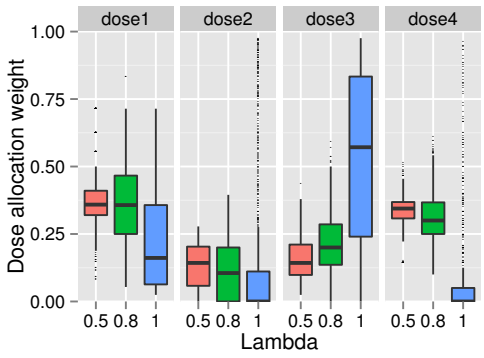

(e) Dose allocation (U1)

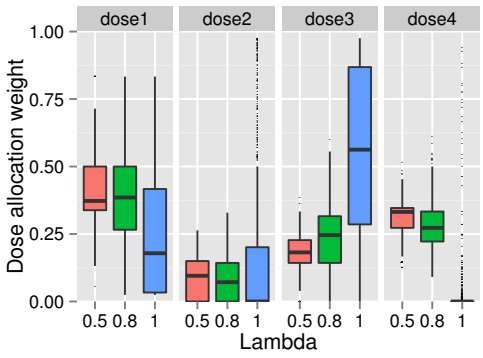

(f) Dose allocation (U2)

Figure 4.6: Operating characteristics with ethical stopping rules based on posterior criterion in (4.4). U1 stands for the scenario with correctly specified correlation. U2 is the scenario with misspecified correlation.

Chapter 5: Redesigns

In this chapter, we demonstrate the proposed Bayesian sequential design with penalty by redesigning two clinical trials in the literature. Specifically, we provide guidance on how to choose parameters in the design, how to select an appropriate design, and how to introduce stopping rules.

5.1 The CML trial

A kinase inhibitor Gleevec (STI571) was investigated for treating chronic myeloid leukemia (CML), in which case patients have higher than normal blood-cell counts. It was hypothesized that the BCR-ABL is a key factor to cause CML and an inhibitor of it can be an effective treatment (Druker et al. 2001). One of the efficacy responses in the study is complete hematologic response, defined as "reduction in the white-cell count to less than 10,000 per cubic millimeter and in the platelet count to less than 450,000 per cubic millimeter, maintained for at least four weeks" (Druker et al. 2001). We redesign a phase I trial of STI571 by using a pharmacokinetics and pharmacodynamics (PK/PD) study to learn prior information about the drug.

5.1.1 Redesign setting

In a PK/PD study (Peng et al. 2004) with 64 patients, the dose range is selected to be 25 mg to 1,000 mg. The researchers targeted steady-state of plasma concentrations of 1 $\mu mol/L^2$, which exceeds the 50% concentration to inhibit BCR-ABL in vitro. This information can be used to select the dose range. According to Table 5.1, the criterion is satisfied for at least 24 hours when doses are above 300 mg. With the dose 85 mg, the criterion is satisfied for the first 6.6 hours. Also, the white blood cell (WBC) counts drop under the normal

level when doses are at least 400 mg. This study indicates the dose-response relationship and provides information about the dose-range we use in the redesign. We therefore select four doses 85 mg, 200 mg, 400 mg and 800 mg to design a hypothetical phase I study.

The phase I study is considered by Druker et al. (2001) as an dose-escalating design for treating CML using STI571. A total of 83 patients are assigned to 14 doses varying from 25 mg/day to 1000 mg/day. Several toxicity responses are defined such as nausea, edema, etc. Among these, a thrombocytopenia toxicity occurs when the platelet count drops below the normal range. The toxicity responses are ordered from 1 to 4, where 1-2 denote mild and moderate toxicity and 3-4 denote severe and life-threatening toxicity. We choose complete hematologic response as our efficacy response and severe thrombocytopenia response as toxicity response.

We determine the prior information from the PK/PD study (Peng et al., 2004). CML patients have higher than normal WBC. In the PK/PD study, WBC at day 28 was recorded for each patient. It is an indication of efficacy when WBC drops below $11 \times 10^9/L$. However, it is an indication of toxicity when WBC is below $4 \times 10^9/L$. Peng et al. (2004) construct relationship between WBC and dose in an Emax model:

$$\log(WBC) = log(E_{\max}) + log\left(1 - \frac{x}{Ex_{50} + x}\right) + \epsilon, \tag{5.1}$$

where x is the log-transformed dose, and $E_{\max}$, Ex_{50} are PK parameters in Table 5.3. Then the probability of WBC below 11 and 4 are calculated for multiple doses in Figure 5.1. The 25th and 50th efficacy quantiles are 75 and 187, while the quantiles for toxicity are 281 and 609. The efficacy parameter should be close to the true parameter for they are based on the WBC counts. There is more uncertainty about the toxicity parameter. We hypothesize two additional scenarios where the quantiles are included in Table 5.4. The parameters are calculated for each scenario and the prior is chosen to be an interval that contains parameters in all scenarios.

We use the data in Table 5.2 to fit a bivariate response model. Specifically, we fit

the efficacy and toxicity data separately with logistic regression models using maximum likelihood estimates. We use logarithms of all raw doses. Because we have dose intervals and dose pairs in the summary data, we assume the patients are uniformly assigned to the specific dose ranges. As a result, for the dose interval $[d_1, d_2]$, we use the integration between d_1 and d_2 to estimate the event probability. For the dose pair "d_1 or d_2", we assume the patients are assigned to the two doses with equal probability and obtain the likelihood. The estimated toxicity parameters are $(6.98, 3.43)$ and the efficacy parameters are $(5.10, 3.21)$. Because there is no summary data for the combined events, we assume a mild correlation with $\rho = 0.2$.

To determine the efficacy-toxicity trade-offs, we assume the two equivalent probability pairs selected by the clinicians are (0.5, 0) and (0.9, 0.1). Therefore $C_T = 5.58$. With this penalty function, the optimal doses are 110, 193, 601 and 318 under scenario 1, 2, 3 and the true value. Then we assume the thresholds of p_{10} and p_T are 0.3 and 0.2 respectively. This corresponds to a penalty threshold of 11.58. This serves as a futility stopping rule.

5.1.2 Simulation results

We conduct 1000 simulations under the scenarios in Table 5.4. For each scenario, we use $\lambda = 0.5, 0.8$ and 1.0. All the operating characteristics are recorded and we want to select the best design to be passed to the redesign.

First, we want to select an appropriate λ for the redesign. This can be determined by multiple factors. One factor is the proportion of different events, which are included in Figure 5.2. We see that $\lambda = 0.5$ performs aggressively with the highest efficacy and toxicity proportions in scenarios 2 and 3. It raises a concern about safety and thus is not selected in the design. The designs with $\lambda = 1$ and 0.8 are competing in scenario 1-2. However, the latter performs better than the former for it has higher efficacy and no toxicity. Another factor is the percentage of patients assigned on the most desirable dose, which is summarized in Table 5.5. Keeping the penalty on each dose in mind, we see that the design with $\lambda = 0.5$ performs better by assigning more patients to the most desirable doses in scenario 1 and 3.

However, it has the tendency to assign patients to the highest dose level. The design with $\lambda = 1$ performs better in scenario 2 by assigning 95% patients to the most desirable dose. But it does not perform well in scenario 3 by assign 31% patients to the dose 2 which has much higher penalty. The design with $\lambda = 0.8$ performs equally well across scenarios. The third factor is how stable the dose allocation is. We plot the distribution of dose allocation in Figure 5.2. The design with $\lambda = 1$ performs extremely stable in scenarios 1 and 2 with almost fixed dose allocation. However, it becomes quite unstable in scenario 3, which has wide range of dose allocation. This means the design is more conservative than the other two information-based designs when the best dose is the highest dose. The two information-based design performs more stably in scenario 3. The last factor is how much information is learned from each design. This can be indicated in Figure 5.3. We see that the design with $\lambda = 0.5$ performs the best with close to truth median estimated optimal dose in all scenarios. It is followed by the design with $\lambda = 0.8$. However, the design with $\lambda = 1$ performs the worst by having a large interquartile ranges in scenarios 1 and 2 and having an optimal dose far from the true value. Its insufficient learning about the dose-response relationship is explained by poorer parameter estimates plotted in Figure 5.3. Based on the above observation, we choose $\lambda = 0.8$ as our design parameter.

We also determine the futility stopping rule based on the ethical posterior criteria proposed in the last chapter. Specifically, we set the minimal efficacy and no toxicity probability to be 0.3 and maximal toxicity probability to be 0.2. The corresponding threshold is 11.58. Then we use the following stopping rule:

$$Pr(\min_{1\leq i\leq k} \Phi_1(\delta_{x_i}; \theta^{(n)}) > 11.58 | \mathcal{D}_n) > p_1,$$

We need to determine p_1 based on the stopping characteristic curves for scenarios 1-3 in Figure 5.4. Intuitively, the test implies that the minimal ethical penalty is greater than 11.58. Then a 80% power test is roughly corresponds to a p_1 at which the mean stopping sample size is $80\% \times 83 = 66.4$. Thus p_1 is 0.375 in scenario 1, 0.25 in scenario 2 and 0.675

in scenario 3. Thus we choose p_1 to be 0.675.

In summary, we choose the design with $\lambda = 0.8$ because it gives rise to relatively high efficacy and no toxicity probability and low toxicity. Also, we use a stopping rule with $p_1 = 0.675$ so that we can stop the trial early for high toxicity or low efficacy.

5.1.3 The redesign

We implement a redesign with the selected design parameters. Here we conduct 1000 simulations under the true parameter settings. We summarize the operating characteristics in Figure 5.5. For comparison, we also include the design with $\lambda = 1$ and 0.5. None of the designs stop for futility. It can be seen that the selected design performs the best with the highest "efficacy and no toxicity" proportion. The true penalty is 9.2, 1.6, 1.3, 7.7 for the four doses. The selected design assigns the most patients to the most desirable dose. The overall estimated optimal doses scatter tightly around the true value, which is an indication of good learning about the dose-response relationship.

We summarize the first simulation as if we were conducting a real design in Table 5.6. The design begins quite conservatively by assigning most patients to the first two doses, even when there is no toxicity. When 20 patients are assigned, the design begins to explore higher doses by assigning most patients to high doses. Because there is no toxicity observed, the design keeps assigning patients to the third dose and assigns only a few to the highest dose. Finally, the design starts to explore more on the second dose and the last dose while maintain a relative high assignments on the third dose. The design ends up with 83% "efficacy and no toxicity" and 3% toxicity. The estimated optimal dose is 341 mg. Comparing with the design with known true parameter, the simulated design shows convergence to the true optimal design with near optimal dose allocation and estimated optimal dose.

5.2 The INFa trial

The dose-response relationship has been examined in an interferon-α-2a (INFa) trial for treating stage IV melanoma (Zohar et al., 2009). The aim of the study is to obtain the optimal dose with a limited sample size. We redesign this trial using the information-based method.

5.2.1 Redesign setting

Because IFNa had been prescribed to patients with other disease before the study was planned, there were knowledge about the dose range and possible adverse events. The four doses were selected as 3, 6, 9, 12 MIU/day. High doses with possibly high toxicity were selected. The efficacy was defined by "complete regression of all the detectable lesions confirmed by two consecutive evaluations performed at a 4-week interval" (Zohar et al., 2009). The toxicity was defined to be " nausea or vomiting (grade 3 and 4) for more than 24 hours" (Zohar et al., 2009). The design was planned with a maximum number of 25 patients.

The investigators decided four plausible scenarios for probabilities of "efficacy and no toxicity" and "toxicity", included in Table 5.7. Thall and Russell's (1998) proportional odds model was used in the original design. Therefore, there is a certain degree of information loss when we fit a Gumbel model to the original trinary outcomes. We fit a Gumbel model to each scenario using the MLE and obtain the estimates (3.00, 1.57, 3.04, 1.62, 0.21) for scenario 1, (3.01, 1.47, 3.02, 1.63, 1) for scenario 2, (14.78, 0.07, 2.56, 1.00, -1) for scenario 3, and (22.94, 0.13, 3.02, 0.46, -0.64) for scenario 4. Because of the information loss, we see that the estimates for the last two scenario do not line up with other scenarios. However, we still believe the estimates reflect the investigators' estimation about the scenarios. Therefore, we restrict the prior to be in the ranges of [2, 23], [0, 2], [2, 4], [0.2, 2] and [-1,1].

To obtain the true dose-response relationship, we use the data in Table 5.8. The MLE is (3.50, 0.64, 2.10, 2.67, 1). The strong positive correlation is due to a small sample size

and is rare in practice. Therefore, we assume the true parameter is (3.50, 0.64, 2.10, 2.67, 0.5).

When it comes to the efficacy-toxicity trade-offs, we assume the equivalent pairs are (0.3, 0) and (0.5, 0.2). This means C_T= 2.29. The optimal doses are 9.5, 8.2, 0.27, 0.16 and 2.8 in scenarios 1-4 and the true model, with penalties 9.9, 10.9, 4.3, 34 and 7.2 respectively. Also, toxicity and "efficacy and no toxicity" probability thresholds are chosen as 0.2 and 0.3, which corresponds to an ethical penalty of 5.56. Then all the scenarios should be stopped as early as possible due to greater than threshold penalty. We apply the same futility stopping rule as in the last trial and determine p_1 through simulation:

$$Pr(\min_{1\leq i\leq k} \Phi_1(\delta_{x_i}; \theta^{(n)}) > 5.56|\mathcal{D}_n) > p_1.$$

5.2.2 Simulation results

We conduct 1000 simulations under scenarios 1, 3 and 4. We determine the best design based on the factors we used in the last example. The first factor is the proportion of events shown in Figure 5.6. The design with $\lambda = 1$ performs uniformly better than the information-based designs in all scenarios by having the lowest toxicity. It also attains fairly high proportion of "efficacy and no toxicity" in scenario 3 and 4. The designs with $\lambda = 0.5$ and 0.8 gives rise to relatively high proportions of "efficacy and no toxicity" in scenario 1 at the cost of even higher toxicity proportions. The second factor is how many patients are assigned to the most desirable dose shown in Table 5.9. The design with $\lambda = 1$ performs conservatively by assign 80% patients to the first dose in scenario 1 where the third dose is the most desirable. However, it successfully identify the most desirable dose in scenarios 2 and 3 by assigning almost all patients to the dose. In contrast, the information-based designs fails to select the most desirable dose and a large number of patients are assigned to the highest dose. Next, we examine the stability of dose allocation in Figure 5.6. The design with $\lambda = 1$ is not stable in scenario 1. It is very stable in scenarios 3 and 4. Although it wastes resources by assigning patients to the least desirable dose in scenario 1, it does

not raise safety issues as we see for the information-based designs. Last, we examine the learning about the dose-response relationship in Figure 5.7. The information-based design performs better than the best intention design in scenario 1 with estimated dose closer to the optimum. This is probably due to a better estimates for α_E. The best intention design performs better in scenarios 3 and 4 with the smallest interquartile ranges for both the estimated optimal dose and the estimated α_E. Therefore, we see that for this over-dosed trial, the information-based designs do not have any edge over the best intention design, whose conservative behavior provides better protection. We select the best intention design for the redesign.

To select the stopping threshold, we consider the 80th quantile of the probability across simulations:

$$Pr(\min_{1\leq i\leq k} \Phi_1(\delta_{x_i}; \theta^{(n)}) > 5.56|\mathcal{D}_n).$$

The quantiles are included in Table 5.10. The quantiles start with a value at the beginning of each simulation due to prior knowledge. It drops after the first six assignments. Therefore, we set p_1 to be 0.8 and effective after 6 patients are assigned.

5.2.3 The redesign

We conduct 1000 simulations with the true parameter. The operating characteristics are included in Figure 5.8. We see median stopping time for the selected design with $\lambda = 1$ is 10, whereas the median stopping times for the other designs are both 7. The selected design also demonstrate the greatest protection of patients with the lowest toxicity proportions shown in Figure 5.8(a). This is due to almost all patients being assigned to the lowest dose level, whereas about 1/4 of the patients are assigned to the highest dose in the information-based design. All the design have good learning about the dose-response relationship by providing accurate estimated optimal dose in Figure 5.8(c). However, the MAP estimates from the best intention design tends to overestimate the optimal dose as oppose to its conservative

behavior. This is because we use a criterion that is based on the mean. The information-based designs have better MAP parameter estimates as the medians of estimates of α_E are closer to the true value. This example shows the discrepancy between the MAP estimates and the mean estimates in small samples.

We summarize the first simulation as if we were conducting a real design in Table 5.11. The trial starts by assigning all patients to the first dose even though the efficacy probability is high. The observed probability of "efficacy and no toxicity" is as high as 0.33 after 6 patients are assigned. Because the posterior probability of futility is 0.76, the trial proceeds by assigning more patients to the first dose. The probability of "efficacy and no toxicity" drops from 0.40 to 0.28 when 25 patients are assigned. The futility rule does not apply because of the low posterior probabilities around 0.4. Due to the low observed toxicity and relatively high "efficacy and no toxicity", a dose of 0.91 mg is recommended for further study.

5.3 Conclusion

In this chapter, we redesign two trials with the aid of simulation under hypothetical scenarios. The information-based design is more aggressive than the best intention design. This can be useful in a trial where the toxicity is expected to be low and the goal is to estimate the optimal dose to be passed onto further studies. The best intention design is relatively conservative by assigning patients to lower dose levels. This usually results in better protection of patients, especially when the toxicity is expected to be high. However, it does not provide sufficient learning about the dose-response relationship, which is important in model-based designs. Therefore, it is possible that decisions based on the best intention design can be inaccurate. In the INFa trial redesign, despite the small sample size, the fact that preselected information-based designs do not perform well is likely caused by the use of noninformative priors.

Table 5.1: Time above targeted concentration at steady-state. (Peng, et al., 2004)

Dose (mg)		85	200	300	400	600	800	1000
Time	Mean	6.6	10.2	32.3	49.3	43.6	62.6	69.3
	SD	6.2	10.9	17.3	17.1	15.1	11.4	22.4

Table 5.2: Efficacy and toxicity data from Druker et al. (2001).

Toxicity - severe to life-threatening thrombocytopenia						
dose	25-140 mg			200-300 mg	350-500 mg	600-1000 mg
Response/Pat. Number	0/14			0/23	1/18	7/28
Efficacy - complete hematologic response						
dose	25 or 50	85	140	200 or 250	300-1000	
Response/Pat. Number	0/6	1/4	1/3	9/16	53/54	

Table 5.3: PK parameters for the Emax model in (5.1) (Peng, et al., 2004).

Parameter	$\log(E_{\max})$	$\log(Ex_{50})$	ϵ
Mean	4.06	3.85	0
SD	0.505	0.643	0.693

Table 5.4: Hypothetical scenario for planning the CML trial.

Scenario	ED25	ED50	LD25	LD50	α_E	β_E	α_T	β_T
1	75	187	281	609	5.23	1.20	6.41	1.42
2	100	200	400	600	5.30	1.58	6.40	2.71
3	400	600	1000	2000	6.40	2.71	7.60	1.58
Prior					$[5, 7]$	$[1, 3]$	$[6, 8]$	$[1, 3]$
True	116	163	782	1075	5.10	3.20	6.98	3.43

Table 5.5: Dose allocation for scenarios 1-3.

		Mean weight of assigned patients			
Scenario	λ	dose 1	dose 2	dose 3	dose 4
1	1.0	0.08	0.92	0	0
1	0.8	0.31	0.65	0.04	0
1	0.5	0.34	0.47	0.18	0.01
Penalty		5.4	6.9	26	480
2	1.0	0.03	0.95	0.01	0
2	0.8	0.05	0.83	0.11	0
2	0.5	0.18	0.47	0.34	0.01
Penalty		5.0	2.8	9.0	2244
3	1.0	0.03	0.31	0.65	0
3	0.8	0.03	0.07	0.73	0.16
3	0.5	0.03	0.08	0.51	0.37
penalty		212	25	6.8	6.1

Table 5.6: An example for the redesign. Number of assigned patients, number of events and estimated optimal dose are included. The optimal dose allocation weights and event proportion given the true parameter is also included in the table.

n	n_1	n_2	n_3	n_4	n_E	n_T	n_{ENT}	Optimal dose	Raw dose
0 - 10	3	4	3	0	6	0	6	6.64	765
11 - 20	0	1	8	1	9	0	9	6.42	614
21- 45	0	2	21	2	22	1	21	5.94	382
46 - 83	0	11	23	4	33	2	31	5.82	336
Redesign (%)	4	22	66	8	84	4	81	5.82	336
True (%)	0	31	67	1	86	2	83	5.76	318

Table 5.7: Selected scenarios for planning the INFa trial. Source: personal correspondence with Dr. Sarah Zohar.

	P_T at dose				P_{10} at dose			
Scenario	3	6	9	12	3	6	9	12
1	0.05	0.10	0.20	0.30	0.05	0.10	0.20	0.20
2	0.05	0.10	0.20	0.30	0.05	0.10	0.20	0.15
3	0.20	0.30	0.40	0.50	0.25	0.23	0.22	0.20
4	0.30	0.35	0.40	0.45	0.05	0.05	0.05	0.05

Table 5.8: Data resulted in the INFa trial. Source: personal correspondence with Dr. Sarah Zohar.

Cohort	Patient No.	Dose	Toxicity	Efficacy	Efficacy and no Toxicity
1	3	1	0	1	1
2	3	2	2	0	0
3	3	2	1	1	0
4	3	2	0	0	0
5	4	3	2	2	0

Table 5.9: Dose allocation for scenarios 1, 3 and 4.

		Mean weight of assigned patients			
Scenario	λ	dose 1	dose 2	dose 3	dose 4
1	1.0	0.8	0.2	0	0
1	0.8	0.24	0.46	0.06	0.24
1	0.5	0.27	0.36	0.02	0.35
Penalty		24	11.8	10	10.4
3	1.0	1	0	0	0
3	0.8	0.68	0.21	0.04	0.07
3	0.5	0.48	0.30	0.00	0.22
Penalty		6.3	9.9	15	22
4	1.0	0.99	0.01	0	0
4	0.8	0.34	0.35	0.08	0.23
4	0.5	0.33	0.32	0.02	0.34
penalty		48	60	71	81

Table 5.10: The mean of the simulation posterior threshold probability.

Scenario	λ	$n = 4$	$n = 7$	$n = 10$	$n = 15$	$n = 20$	$n = 25$
1	1	0.95	0.93	0.93	0.93	0.93	0.93
3	1	0.78	0.73	0.72	0.71	0.72	0.72
4	1	0.95	0.96	0.97	0.99	0.99	1.00

Table 5.11: Example redesign for the INFa trial. Cumulative dose allocation weights, event proportion and estimated optimal dose. The optimal dose allocation and event proportion given the true parameter is also included in the table.

n	w_1	p_E	p_T	p_{10}	optimal dose	raw dose	$Pr(\min \Phi_1 > 5.56 \mid \mathcal{D}_n)$
6	1	0.50	0.167	0.33	-0.90	0.41	0.76
10	1	0.50	0.10	0.40	-2.2	0.11	0.37
15	1	0.47	0.13	0.40	-0.73	0.48	0.41
20	1	0.40	0.10	0.35	-0.79	0.45	0.41
25	1	0.32	0.08	0.28	-0.09	0.91	0.34
True	1	0.18	0.07	0.16	1	2.8	-

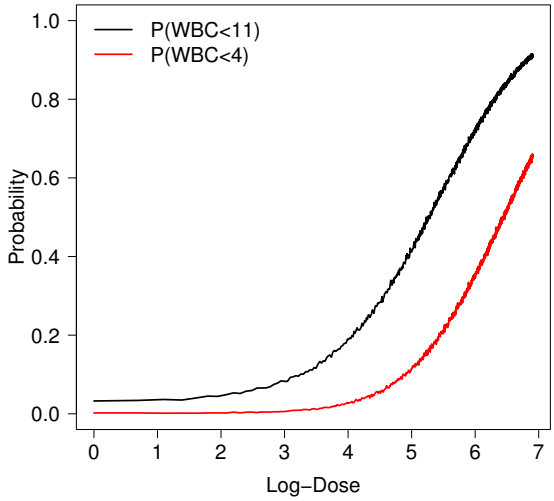

Figure 5.1: Probability of WBC below 11 and 4 based on simulation using the Emax model (5.1).

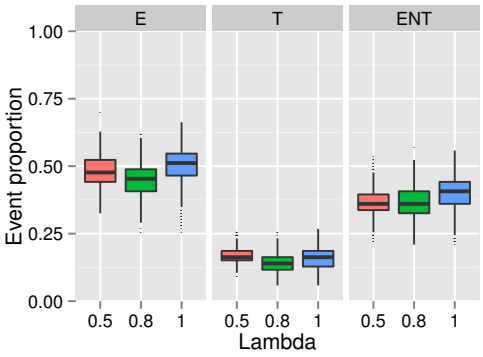

(a) Scenario 1, event proportion

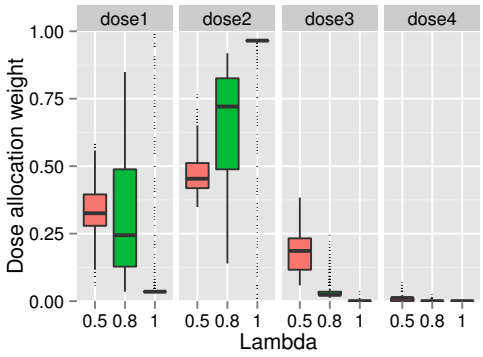

(b) Scenario 1, dose allocation

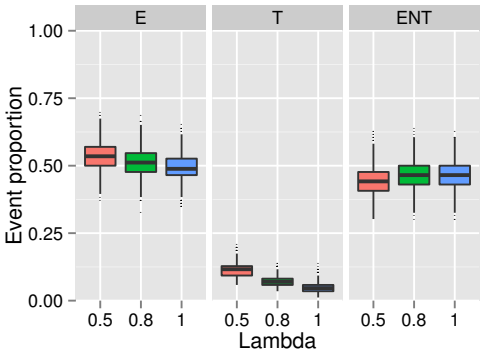

(c) Scenario 2, event proportion

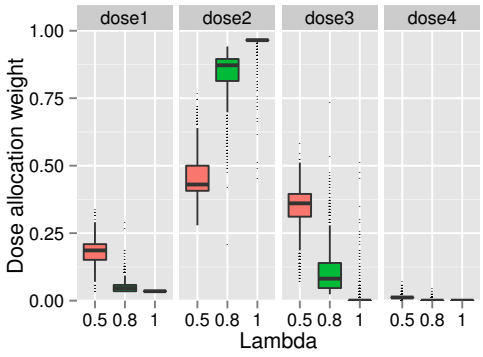

(d) Scenario 2, dose allocation

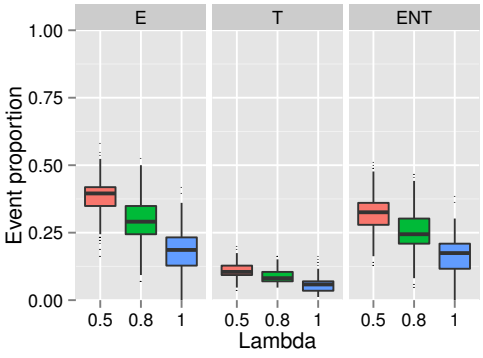

(e) Scenario 3, event proportion

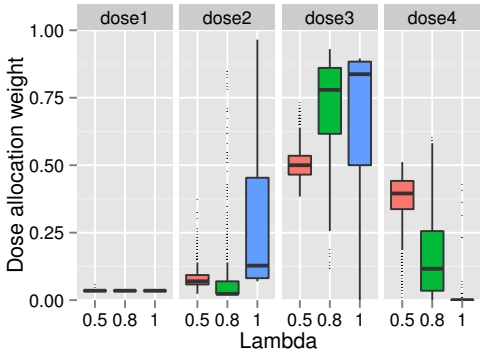

(f) Scenario 3, dose allocation

Figure 5.2: Proportions of events and dose allocation in scenarios 1-3. Efficacy is denoted by "E". Toxicity is denoted by "T". Efficacy and no toxicity is denoted by "Et".

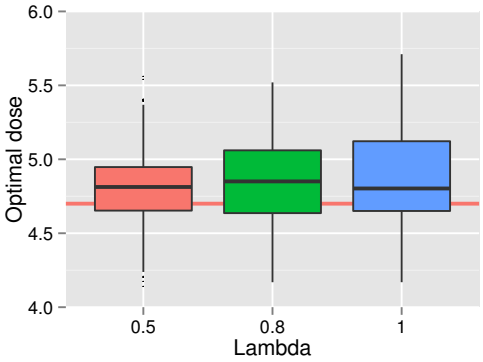

(a) Scenario 1, estimated optimal dose

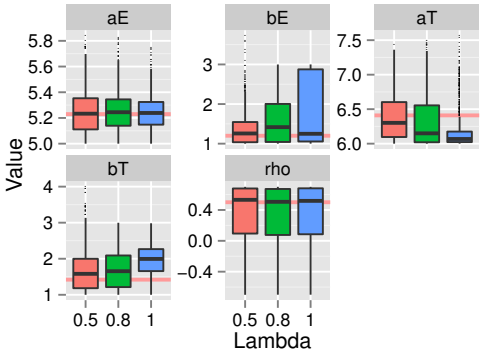

(b) Scenario 1, parameter estimates

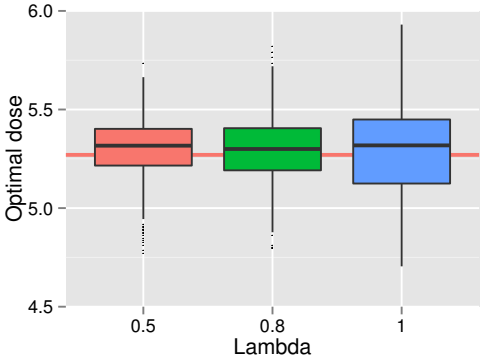

(c) Scenario 2, estimated optimal dose

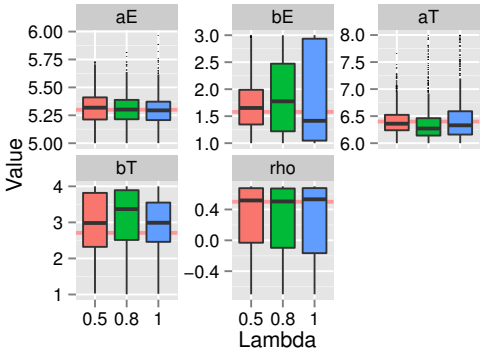

(d) Scenario 2, parameter estimates

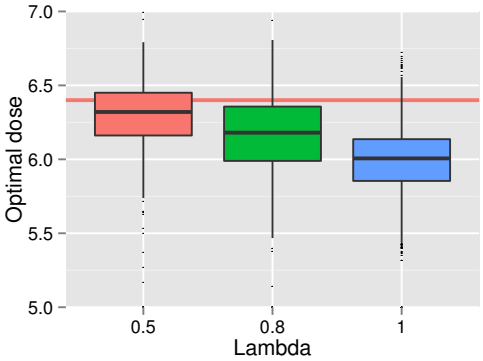

(e) Scenario 3, estimated optimal dose

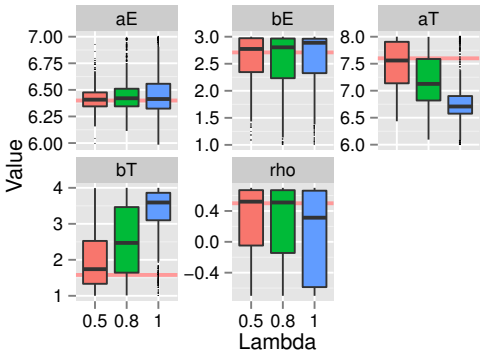

(f) Scenario 3, parameter estimates

Figure 5.3: Estimated optimal dose and parameter estimates in scenarios 1-3.

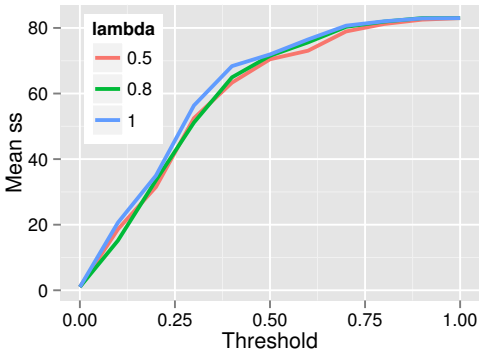

(a) Scenario 1, mean stopping sample size

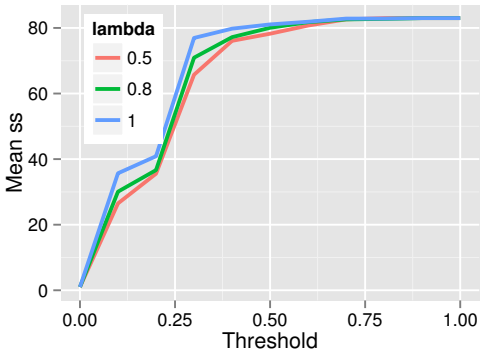

(b) Scenario 2, mean stopping sample size

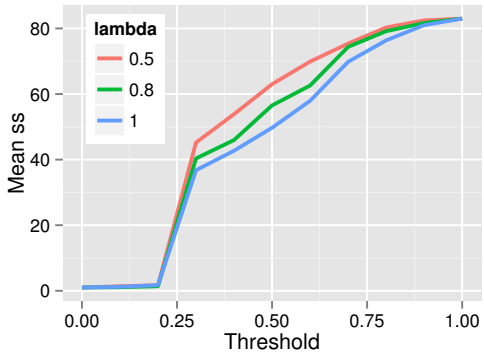

(c) Scenario 3, mean stopping sample size

Figure 5.4: Mean stopping sample size for scenarios 1-3.

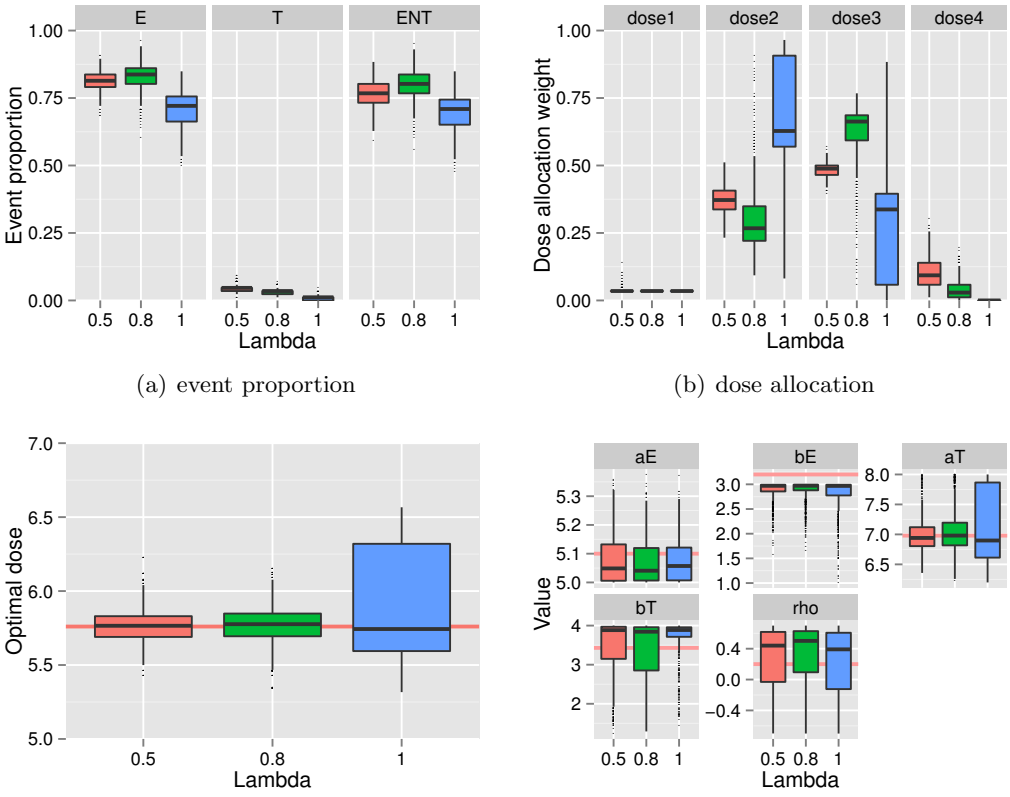

(a) event proportion

(b) dose allocation

(c) estimated optimal dose

(d) parameter estimates

Figure 5.5: Operating characteristics of the redesign.

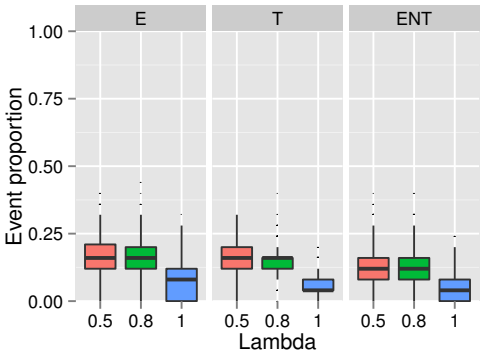

(a) Scenario 1, event proportion

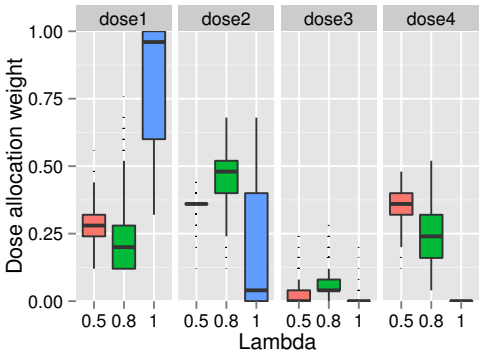

(b) Scenario 1, dose allocation

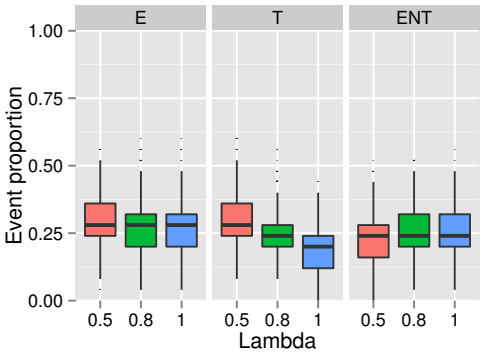

(c) Scenario 3, event proportion

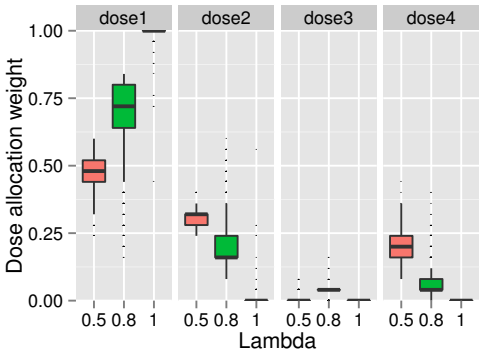

(d) Scenario 3, dose allocation

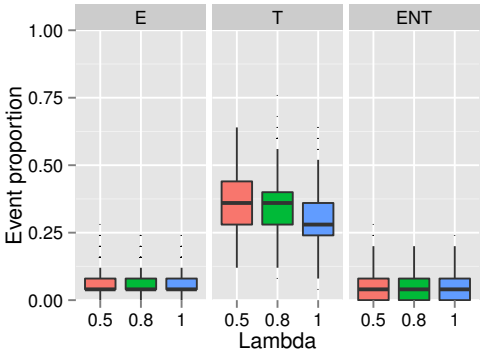

(e) Scenario 4, event proportion

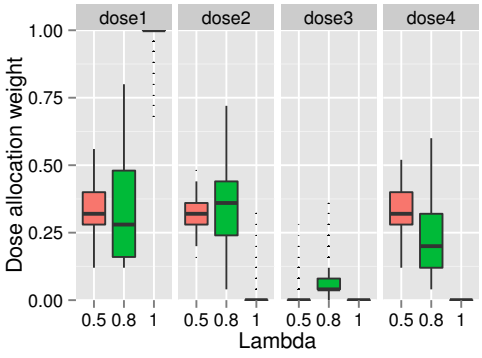

(f) Scenario 4, dose allocation

Figure 5.6: Proportions of events and dose allocation in scenarios 1, 3 and 4. Efficacy is denoted by "E". Toxicity is denoted by "T". Efficacy and no toxicity is denoted by "Et".

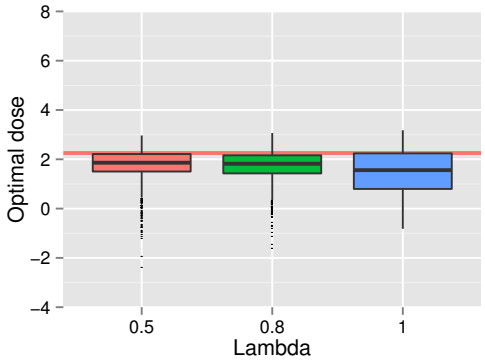

(a) Scenario 1, estimated optimal dose

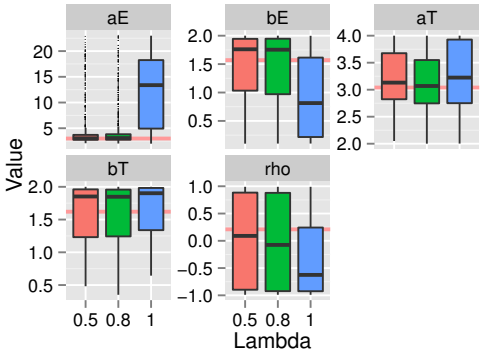

(b) Scenario 1, parameter estimates

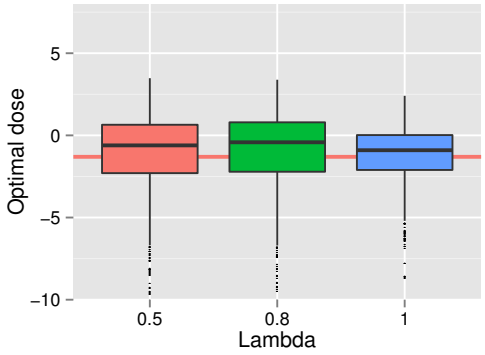

(c) Scenario 2, estimated optimal dose

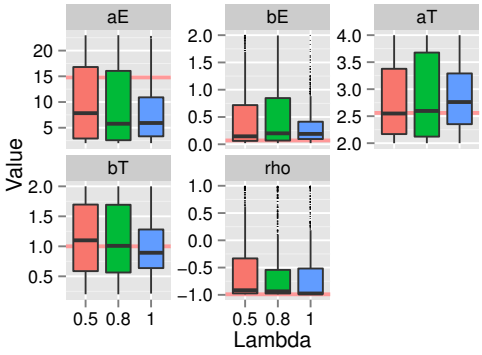

(d) Scenario 2, parameter estimates

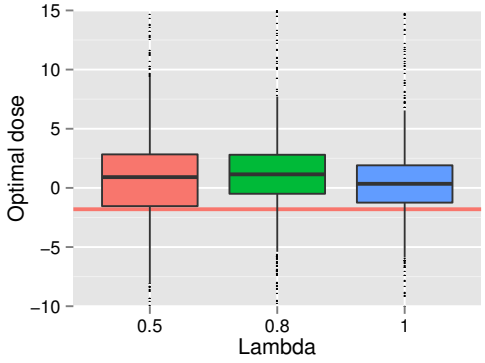

(e) Scenario 3, estimated optimal dose

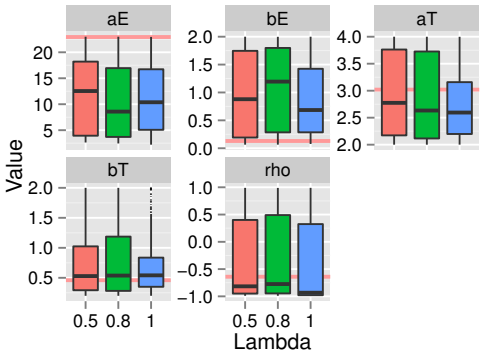

(f) Scenario 3, parameter estimates

Figure 5.7: Estimated optimal dose and parameter estimates in scenarios 1, 3 and 4.

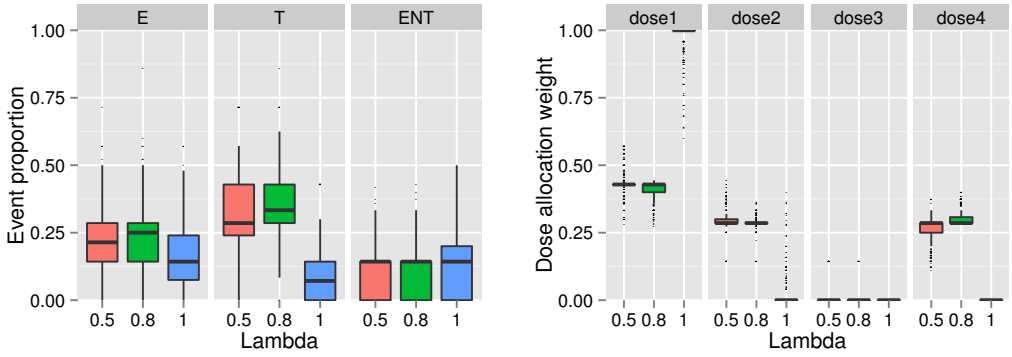

(a) event proportion

(b) dose allocation

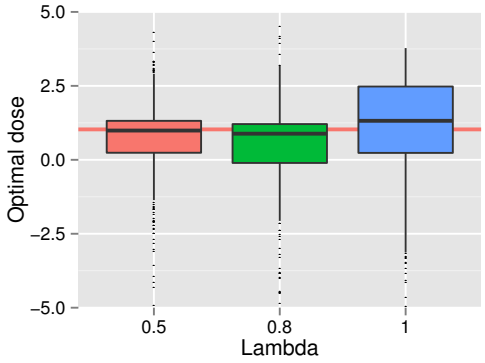

(c) estimated optimal dose

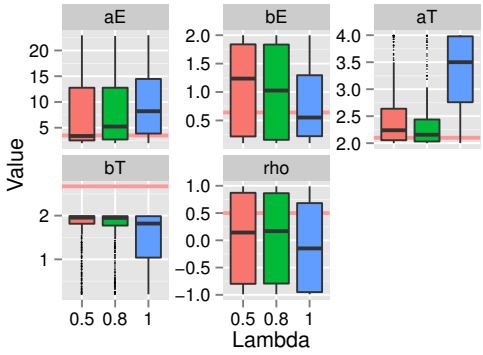

(d) parameter estimates

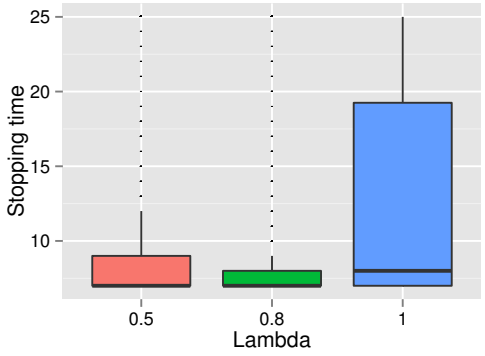

(e) stopping time

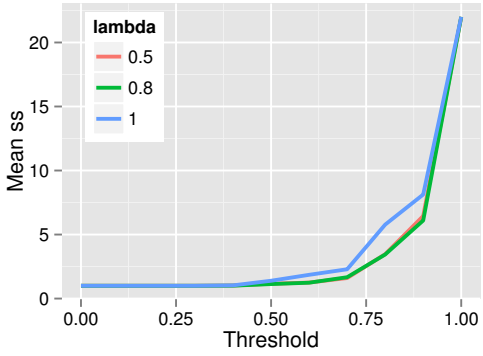

(f) stopping characteristic curve

Figure 5.8: Operating characteristics of the INFa trial redesign.

Chapter 6: Conclusions

In this dissertation, we examine Bayesian dose-finding procedures using simulation. One commonly used design philosophy is to find the optimal dose by assigning most patients to it. The philosophy is adapted by best intention designs. We identify the non-convergence issue with best intention designs. In the most extreme case, the best intention designs cease to gain new information about the dose-response relationship once they begin assigning patients to a single dose. The loss is the additional information for a better understanding of the drug and more powerful future studies. By proposing to use the information criterion to maximize the learning about the dose-response relationship, we introduce information-based designs. These designs are shown to have a better convergence property and better learning about the dose-response relationship, especially when priors are misspecified. The enhanced learning leads to better protection of patients due to earlier stopping. Despite its advantages, the information-based design tends to be aggressive by exploring high toxic doses. Therefore, a penalty coefficient is introduced to trade off information goal and ethical goal. Using simulation, we show that a comprise between the two goals can be achieved by tuning the penalty coefficient. In hypothetical redesigns, we provide guidance on how to select priors, trade off efficacy and toxicity, and select the best design.

From the philosophical perspective, the information and ethical goals appears to be competing objectives for Bayesian dose-finding studies. The former is achieved across multiple dose levels while the latter is achieved with a single dose. Behind this, there is the debate about selection between collective ethical goals and individual ethical goals. From the practical perspective, we demonstrate that the information goal is crucial, because without efficient parameter estimation, the ethical measure is not reliable and ethical gains cannot be maximized. In addition, stopping rules are used to introduce further protection. In model-based designs, stopping rules are usually based on parameter estimation. Therefore

the information goal is also important for these designs. To obtain a good design, multiple doses should be explored while patients are protected. And both goals can be weighted and achieved in a systematic manner using information-based designs.

Admittedly, there are other ways to achieve multiple-dose allocation. For example, one can conduct a two stage design with the first stage estimating the parameter and the second stage aiming for the ethically optimal dose. Also, one can assign participants to doses with probabilities that are proportional to the ethical utilities (Lei, Yuan and Yin, 2011). The efficiency of parameter estimation and performance stability of these designs need careful examination.

From the technical perspective, whether a design leads to a stable multiple-dose allocation is an important design criterion. Instead of listing only the mean operating characteristics, biostatisticians need to check stability under different scenarios using simulation and graphics. The worst case scenario should be addressed in simulation to check designs' robustness. Also, information-based design should be considered as a gold standard to assess the efficiency of parameter estimation. In this way, reliable designs can be obtained without producing unexpected outcomes.

A list of contributions include the following:

1. A class of adaptive designs with both the information goal and ethical goal is introduced and explored. The best intention design can be considered a special case of this class;

2. The convergence issue of best intention design is identified and solutions has been proposed; An example is given where the best intention design converges to the wrong dose, yet the information-based designs are effective;

3. Design robustness to prior misspecification and different priors are examined through simulation;

4. Different stopping rules are considered and used to compare different designs;

5. Practical guidance on prior elicitation, toxicity-efficacy trade-off specification and selection of design parameters is given in two redesign examples.

Several questions remain unsolved. Theoretical convergence properties of the proposed information-based design are to be explored. It is interesting to incorporate information criterion into different models, including the bivarite probit model, survival models and continual reassessment methods. Different ethical penalty functions may be used, for example the ratio of area under curve (Yuan and Yin, 2009). Different stopping criteria may be explored such as the sequential testing method (Zohar and O'Quigley, 2006).

The issue of appropriate sample sizes for using Bayesian sequential information-based designs is an open question. Under correctly specified priors, a smaller trial is more feasible, but this is unknown before the trial. One can simulate the trial under different scenarios and sample sizes to determine a minimum sample size that will result in reasonable estimates of the dose-response relationship.

Bibliography

Bibliography

[1] Atkinson, A. C., Donev, A. N., and Tobias, R. D. (2007) *Optimal Experimental Design, with SAS.* Oxford University Press: Oxford.

[2] Burkardt, J. (2006) NORMAL - Normal Random Number Generators. *http://people.sc.fsu.edu/∼jburkardt/c_src/normal/normal.html*

[3] Clayton, D. G. (1978) A model for association in bivariate life tables and its application in epidemiological studies of familial tendency in chronic disease incidence. *Biometika* **65** 141–151.

[4] Chaloner, K., Church, T., Louis, T. A., Matts, J. P. (1993) Graphical elicitation of a prior distribution for a clinical trial. *Journal of the Royal Statistical Society D* **42** 341–353.

[5] Chaloner, K., Larntz, K. (1989) Optimal Bayesian design applied to logistic regression experiments. *Journal of Statistical Planning and Inference* **21** 191–208.

[6] Chaudhuri, P., Mykland, P. A. (1993) Nonlinear experiments: optimal design and inference based on likelihood. *Journal of the American Statistical Association* **88** 538–546.

[7] Chaudhuri, P., Mykland, P. A. (1995) On efficient designing of nonlinear experiments. *Statistica Sinica* **5** 421–440.

[8] Chuang-Stein, C., D'Agostino, R. (2007) Optimal design of experiments in pharmaceutical applications. In *Pharmaceutical Statistics Using SAS, A Practical Guide* (Dmitrienko, A., Chuang-Stein, C., D'Agostino, R., eds.), SAS Institute: Cary, NC.

[9] de Lima, M., Champlin, R. E., Thall, P. F., Wang, X., Martin III, T. G., Cook, J. D., McCormick, G., Qazilbash, M., Kebriaei, P., Couriel, D., Shpall, E., J., Khouri, I., Anderlini, P., Hosing, C., Chan, K. W., Andersson, B. S., Patah, P. A., Caldera, Z., Jabbour, E., Giralt, S. (2008) Phase I/II study of gemtuzumab ozogamicin added to fludarabine, melphalan and allogeneic hematopoietic stem cell transplantation for high-risk CD33 positive myeloid leukemias and myelodysplastic syndrome. *Leukemia* **22** 258–264.

[10] Dette, H., Grigoriev, Y. (2000) A unified approach to second order optimality criteria in nonlinear design of experiments. *Annals of the Institute of Statistical Mathematics* **52** 574–597.

[11] Dragalin, V., Fedorov, V. (2006) Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* **136** 1800–1823.

[12] Dragalin, V., Fedorov, V., Wu, L. (2008) Adaptive designs for selecting drug combinations based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* **138** 352–373.

[13] Druker B. J., Talpaz M., Resta D. J., Peng B., Buchdunger E., Ford J. M., Lydon N. B., Kantarjian H., Capdeville R., Ohno-Jones S., Sawyers C. L. (2001) Efficacy and safety of a specific inhibitor of the BCR-ABL tyrosine kinase in chronic myeloid leukemia. *The New England Journal of Medicine***344** 1031–1037.

[14] Elfving, G. (1952) Optimum allocation in linear regression theory. *The Annals of Mathematical Statistics* **23** 255–262.

[15] Fedorov, V. V., Flournoy, N., Wu, Y., Zhang, R. (2011) Best intention designs in dose-finding studies. *http://www.newton.ac.uk/preprints/NI11065.pdf*

[16] Fedorov, V., Gagnon, R., Leonov S., Wu, Y. (2007) Optimal design of experiments in pharmaceutical applications. In *Pharmaceutical Statistics Using SAS, A Practical Guide* (Dmitrienko, A., Chuang-Stein, C., D'Agostino, R., eds.), SAS Institute: Cary, NC.

[17] Fraleigh, J. B., Beauregard, R. A. (1995) *Linear Algebra.* Addison-Wesley: Sydney.

[18] Gaydos, B., Krams, M., Perevozskaya, I., Bretz, F., Liu, Q., Gallo, P., Berry, D., Chuang-Stein, C., Pinheiro, J., Bedding, A. (2006) Adaptive dose-response studies. *Drug Information Journal* **40** 451–461.

[19] Gao, L., Rosenberger, W. F. (2013) Adaptive Bayesian design with penalty based on toxicity-efficacy responses. In *mODa 10- Advances in Model-Oriented Design and Analysis* (Ucinski, D., Atkinson, A. C., Patan, M., eds.), Springer: Swizerland, 91–98.

[20] Gharibvand, L., Liu, L. (2009) Analysis of survival data with clustered events. *Statistics and Data Analysis*. SAS Global Forum.

[21] Haines, L. M., Perevozskaya, I., Rosenberger, W. F. (2003) Bayesian optimal designs for phase I clinical trials. *Biometrics* **50** 591–600.

[22] Hu, I. (1998) On sequential designs in nonlinear problems. *Biometrika* **85** 496–503.

[23] Kadane, J. B., Wolfson, L. J. (1996) Priors for the design and analysis of clinical trials. In *Bayesian Biostatistics* (Berry, D. A., Stangl, D. K. eds.), Marcel Dekker: New York, 157–184.

[24] Lai, T. L. (1994) Asymptotic properties of nonlinear least squares estimates in stochastic regression models. *Annals of Staitics* **22** 1917–1930.

[25] Lehmann E. L., Casella, G. (1998) *Theory of Point Estimation.* Springer-Verlag: New York.

[26] Manukyan, Z. (2009) *Sequential designs for etimation toxicity and efficacy in a dose-response setting*. Dissertation. George Mason University: Fairfax, VA.

[27] Manukyan, Z., Rosenberger, W. F. (2010) D-optimal design for a five-parameter logistic model. In *mODa 9- Advances in Model-Oriented Design and Analysis* (Giovagnoli, A., Atkingson, A. C., Torsney, B., eds.), Physica-Verlag: Heidelberg, 113–120.

[28] Murphy, S. A. (1994) Consistency in a proportional hazards model incorporating a random effect. *The Annals of Statisitcs* **22** 712–731.

[29] Murphy, S. A. (1995) Asymptotic theory for the frailty model. *The Annals of Statistics* **23** 182–198.

[30] O'Quigley, J., Zohar, S. (2006) Experimental designs for phase I and phase I/II dose-finding studies. *British Journal of Cancer* **94** 609–613.

[31] Peng, B., Hayes, M., Resta, D., Racine-Poon, A., Druker, B. J., Talpaz, M., Sawyers, C. L., Rosamilia, M., Ford, J., Lloyd, P., Capdeville, R. (2004) Pharmacokinetics and pharmacodynamics of Imatinib in a phase I trial with chronic myeloid Leukemia patients. *Journal of Clinical Oncology* **22** 935–942.

[32] Perevozskaya, I., Rosenberger, W. F., Haines, L. M., (2003) Optimal design for the proportional odds model. *The Canadian Journal of Statistics* **31** 225-235.

[33] Plummer, M. (2003) JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling. *Proceedings of the 3rd International Workshop on Distributed Statistical Computing* Vienna, Austria.

[34] Pronzato, L. (2010a) One-step ahead adaptive D-optimal design on a finite design space is asymptotically optimal. *Metrika* **71** 219–238.

[35] Pronzato, L. (2010b) Penalized optimal designs for dose-finding. *Jounal of Statistical Planning and Inference* **140** 283–296.

[36] R Core Team (2012) R: A Language and Environment for Statistical Computing. *http://www.R-project.org/* R Foundation for Statistical Computing, Vienna, Austria.

[37] Rosenberger, W. F., Canfield, G. C., Perevozskaya, I., Haines, L. M., Hausner, P. (2005) Development of interactive software for Bayesian optimal phase I clinical trial design. *Drug Information Journal* **39** 89–98.

[38] Roy, A., Ghosal, S., Rosenberger, W. F. (2009). Convergence properties of sequential Bayesian *D*-optimal designs. *Journal of Statistical Planning and Inference* **139** 425–440.

[39] Thall, P. F., Cook, J. D. (2004) Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* **60** 684–693.

[40] Thall, P. F., Cook, J. D., Estey, E. H. (2006) Adaptive dose selection using efficacy-toxicity trade-offs: illustrations and practical considerations. *Journal of Biopharmaceutical Statistics* **16** 623–638.

[41] Thall, P. F., Russell, K. E. (1998) A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* **54** 251–264.

[42] Wang M. (2007) *An Adaptive Bayesian Approach to Jointly Modeling Response and Toxicity in Phase I Dose-finding Trials.* Dissertation. University of Pittsburger.

[43] Wei, L. J., Lin, D. Y., Weissfeld, L. (1989) Regression analysis of multivariate incomplete failure time data by modeling marginal distributions. *Journal of the American Statistical Association* **84** 1065–1073.

[44] Wickham, H. (2009) *ggplot2: An implementation of the Grammar of Graphics.* Springer: New York.

[45] Whittle, P. (1973) Some general points in the theory of optimal experimental design . *Journal of the Royal Statistical Society B* **35** 123–130.

[46] Wu, C. F. J. (1988) Optimal design for percentile estimation of a quantal response curve. In *Optimal Design and Analysis of Experiments* (Dodeg, Y., Fedorov, V. V., Wynn, H. P., eds.), North-Holland: Amsterdam, 213–223.

[47] Yuan, Y., Yin, G. (2009) Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society C* **58** 719–736.

[48] Zohar, S., O'Quigley, J. (2006) Identifying the most successful dose in dose-finding studies in cancer. *Pharmaceutical Statistics* **5** 187–199.

[49] Zohar, S., Medioni, J., Lebbe, C., Avril, MF, Kerob, D., Eftekhari, P., Driheme, A., Lebras, K. C., Bruzzoni-Giovanelli, H., Levy, V., Chevret, S., Calvo, F. (2009) Interest in an original methodology to define the optimal dosage of interferon-alpha-2a in meastatic melanoma patients. *Melanoma Research* **19** 379–384.

BIOGRAPHY

Lei Gao graduated from Peking University in 2005 with a B.S. degree in Mathematics. He received his M.S. degree from University of New Hampshire in 2010. He worked as a Teaching Assistant at the Department of Mathematics and Statistics, University of New Hampshire from 2008 to 2010 and then as a Graduate Research Assistant at the Department of Statistics, George Mason University from 2011 to 2014.