

# Early-phase Dose-finding Designs for Bivariate Outcomes

Christopher Michael Tait  
Nokesville, Virginia

B.S. Hampden-Sydney College, 2010  
M.S. University of Virginia, 2012

A Dissertation Presented to the Graduate Faculty of the University of Virginia in  
Candidacy for the Degree of Doctor of Philosophy

Department of Statistics

University of Virginia  
December 2014

## **Abstract**

The primary goal of a Phase I clinical trial is to estimate the maximum tolerated dose (MTD), defined as the highest dose which can be administered with a “tolerable” level of toxicity. The “tolerable” level of toxicity is based on the probability that a patient in the trial experiences a dose-limiting toxicity (DLT). Once a dose has been deemed safe in a Phase I trial, the next step is to conduct a Phase II trial where the goal is to see if the drug is effective, yet with a close eye remaining on the level of toxicity. Usually, Phase I and Phase II trials are performed independently, without formally sharing information across the separate phases. There has been a recent shift in the paradigm of drug development in oncology to integrate Phase I and Phase II trials so that the drug development process may be accelerated. In dose-finding trials of chemotherapeutic agents, the goal of identifying the MTD is usually determined by considering information on toxicity only, with the assumption that the highest safe dose also provides the most promising outlook for efficacy. A growing interest in targeted therapies challenges accepted dose-finding methods because minimal toxicity may arise over all doses under consideration and higher doses may not result in greater response. This dissertation proposes a collection of dose-finding methods for targeted agent trials involving bivariate outcomes. Various simulation studies are presented to evaluate the operating characteristics and to examine the robustness of the proposed method against varying design specifications. A discussion of the theoretical properties is presented. An extension of the proposed method into situations where there is delayed response is also explored. We close with some conclusions and some areas for future research.

# Acknowledgments

I would like to express my sincere gratitude to my advisors, Professor Nolan Wages and Professor Dan Spitzner. This dissertation could not have been written without their guidance. Prof. Wages has been a great mentor to me for many years beginning with him advising me on an Honor's Project my senior year of undergraduate studies at Hampden-Sydney College and leading all the way through my graduate studies before kindly agreeing to become one of my advisors. I would also like to extend my sincere thanks to Prof. Spitzner for helping with the theoretical aspect of my research.

I would also like to thank my other committee members, Professor Mark Conaway and Professor Jeff Holt, for reading my dissertation and offering useful comments. Both of them have been instrumental to my success throughout graduate school and my beginning career. Thanks must also go to the faculty and staff of the Department of Statistics, who have taught me a great deal in my pursuit of this degree.

I would like to thank my parents and siblings for their continued support and encouragement throughout the past few years. They have provided me with constant motivation to pursue this goal.

# Contents

<b>1</b>	<b>Early-phase Clinical Trials</b>	<b>1</b>
1.1	The Problem . . . . .	1
1.2	Background of Dose-finding . . . . .	1
1.3	Phase I Trials . . . . .	3
1.3.1	Trial Objective . . . . .	3
1.3.2	Algorithm-based Designs . . . . .	4
1.3.3	Model-based Designs . . . . .	6
1.4	New Paradigm of Dose-finding . . . . .	10
1.4.1	Locating the OBD . . . . .	11
1.4.2	Literature Review . . . . .	13
1.4.3	Areas of Improvement in Targeted Agents . . . . .	15
1.5	Outline of Dissertation . . . . .	16
<b>2</b>	<b>Models and Inference: Bivariate Binary Outcomes</b>	<b>18</b>
2.1	Modeling Toxicity . . . . .	19
2.2	Models and Inference for Efficacy . . . . .	19
2.3	Algorithm . . . . .	21
2.3.1	Adaptive Randomization Phase . . . . .	23
2.3.2	Maximization Phase . . . . .	23
2.3.3	Starting the Trial . . . . .	24

2.4	Method Based Solely on Efficacy . . . . .	24
2.5	Simulation Studies of Proposed Method . . . . .	26
2.5.1	Illustration . . . . .	26
2.5.2	Results . . . . .	30
2.6	Design Specifications . . . . .	37
2.6.1	Adaptive Randomization Sample Size . . . . .	37
2.6.2	Skeleton Choice . . . . .	38
2.6.3	Working Model Choice . . . . .	40
2.6.4	Utility Functions . . . . .	44
2.6.5	Model Selection Criteria . . . . .	46
2.7	Terminating the Trial . . . . .	48
<b>3</b>	<b>Theoretical Properties</b>	<b>51</b>
3.1	Convergence . . . . .	51
3.1.1	Model Assumptions in CRM . . . . .	51
3.1.2	Bivariate Extension of Model Assumptions . . . . .	54
3.2	Coherence of CRM . . . . .	68
3.2.1	Coherence of Bivariate Endpoints . . . . .	69
3.2.2	Coherence for Bivariate Binary Targeted Agents . . . . .	70
<b>4</b>	<b>Application to Delayed Response</b>	<b>72</b>
4.1	Lagged Response . . . . .	73
4.2	Review of TITE CRM . . . . .	74
4.3	Extension of TITE to Bivariate Outcomes . . . . .	76
4.3.1	TITE Model on Efficacy . . . . .	76
4.3.2	TITE Model on Toxicity and Efficacy . . . . .	82

<b>5</b>	<b>Conclusions and Future Research</b>	<b>86</b>
5.1	Introduction . . . . .	86
5.2	Future Research . . . . .	87
5.3	Concluding Remarks . . . . .	91
<b>A</b>	<b>Toxicity and Efficacy Scenarios</b>	<b>92</b>
<b>B</b>	<b>Simulation Results: Proposed Method</b>	<b>93</b>
<b>C</b>	<b>Simulation Results: Delayed Response</b>	<b>105</b>
	<b>Bibliography</b>	<b>109</b>

# Chapter 1

## Early-phase Clinical Trials

### 1.1 The Problem

This dissertation focuses on early-phase, dose-finding clinical trial designs for targeted agents in oncology. Targeted therapies are challenging accepted methods in this area, which were developed with cytotoxic agents in mind. For cytotoxic therapies, dose-finding studies are initial safety trials that evaluate patients based on a single toxicity endpoint, defined by the presence/absence of severe adverse events. In contrast, many targeted agents are considered intrinsically safe, often with minimal toxicity observed throughout the trial duration. Therefore, patients must be evaluated on more than one endpoint in order for the primary objective of the trial to be achieved. In this work, new dose-finding methods are proposed that account for the bivariate outcomes present in early-phase targeted therapy trials.

### 1.2 Background of Dose-finding

When a new drug is being investigated for use in humans, traditionally the first step is to carry out a Phase I clinical trial. If the drug is for a rare disease, such

as cancer, then the patients who are enrolled into the trial are ones in which other therapies have failed and are sometimes participating in the trial as a last resort. As a result, Phase I trials in oncology must take caution in order to not expose patients to an overly toxic dose and to not expose patients to doses which do not have enough therapeutic benefit. The trials must have a balance between maintaining an acceptable toxicity level and keeping the effectiveness of the doses high (O’Quigley, Pepe, and Fisher, 1990). Once a dose has been deemed safe in a Phase I trial, the next step is to conduct a Phase II trial where the goal is to see if the drug is effective, yet with a close eye remaining on the level of toxicity. Usually, Phase I and Phase II trials are performed independently, without formally sharing information across the separate phases. There has been a recent shift in the paradigm of drug development in oncology to integrate Phase I and Phase II trials so that the drug development process may be accelerated, while potentially reducing costs (Yin, 2012). To this end, several published Phase I/II methods have extended dose-finding methodology to allow for the modeling of both toxicity and efficacy (O’Quigley, Hughes, and Fenton, 2001; Braun, 2002; Thall and Cook, 2004; Yin, Li, and Ji, 2006; among others). Each of these designs, as well as the one we propose in this dissertation, is referred to by Federov and Leonov (2013) as a “best intention” (BI) design, meaning that it aims to allocate patients in the current trial to the “best” treatment based on the current knowledge available. We distinguish these designs from “optimal” designs in the area of dose-finding that seek to maximize global information about dose-response curves in order to optimize doses for future patients. In addition to the detailed discussion provided by Federov and Leonov (2013), we refer any reader interested in this type of design to Dragalin and Federov (2006) and Whitehead et al. (2006), among many others.



## 1.3 Phase I Trials

### 1.3.1 Trial Objective

The typical goal of a Phase I trial is to find a dose with an acceptable level of toxicity while maximizing the benefit that the patient receives. This goal is attained simultaneously because a standard assumption for Phase I trials is that both toxicity and efficacy increase with dose. Therefore, when we find the dose closest to some target rate we also find the dose with the maximum benefit and acceptable toxicity rate. The dose level with an estimable level of toxicity closest to the target is known as the maximum tolerated dose (MTD). Historically, a phase I trial typically aims to find the MTD from a discrete set of  $I$  doses,  $\{d_1, \dots, d_I\}$ , approved for experimentation. Therefore, the methods developed for Phase I trials were originally designed for cytotoxic agents. However with the transition to targeted agent therapies in recent years, the National Cancer Institute formed the Clinical Trials Design Taskforce. In summary of this taskforce, LoRusso, Boerner, and Seymour (2010) stated:

“Improperly or inefficiently designed early clinical trials results in the exposure of patients to ineffective and sometimes toxic agents, delay in the development times for potentially effective novel agents, and incur enormous costs to society. An appropriately designed phase I study is critical to inform investigators whether the drug is of sufficient interest to pursue further, and to determine the appropriate dose, schedule, and patient population for further study.”

The taskforce highlighted two major areas of early-phase clinical trials, targeted agents, which are agents that inhibits the growth and multiplication of cells by targeting a specific tumor pathway for example, and combinations of agents, in which there is a need to develop new designs for Phase I trials. Methods for early-phase clinical trials can be grouped into two categories: (1) algorithm based designs and (2) model-based designs.

### 1.3.2 Algorithm-based Designs

Algorithm-based designs have the advantage of being simple to implement in practice. Storer (1989) gives an extensive review of algorithm-based designs including the traditional “3 + 3” and other Up-and-Down designs. Another advantage of these designs is the fact that they do not make large jumps within the discrete set of doses since the dose given to the next patient is always equal to the current dose or an adjacent (the next higher or next lower) dose. Most Up-and-Down designs can be classified as random walk rules or, more formally, first-order Markov procedures. A drawback of these designs is they determine the dose for the next patient based only on the most recent dose and outcome. However, this simplicity allows for the direct application of random walk theory in order to obtain the exact limiting stationary distribution for the method. For example, the biased coin design (BCD) developed by Durham and Flournoy (1994) determines the  $\phi_T^{th}$  quantile,  $0 < \phi_T \leq 0.50$ , of a dose-toxicity curve using possible doses,  $d_1 < d_2 < \dots < d_I$ . The dose given to the  $j^{th}$  patient,  $x_j \in \{d_1, \dots, d_I\}$ , is determined through the use of a biased coin with the probability of heads of  $\phi_T/(1 - \phi_T)$ , the dose given to the previous patient,  $x_{j-1} = d_i$ , and the previous subject’s observed toxic response,  $y_{j-1}$ , through the following rule:

$$x_j = \begin{cases} d_{\max\{1, i-1\}} & \text{if } y_{j-1} = 1, \\ d_{\min\{i+1, I\}} & \text{if } y_{j-1} = 0 \text{ and the coin lands heads,} \\ d_i & \text{if } y_{j-1} = 0 \text{ and the coin lands tails.} \end{cases}$$

The authors showed that as long as  $F(d) = \Pr(y_j = 1 | x_j = d)$  is non-increasing the limiting distribution of the dose allocation for the BCD is unimodal around the  $\phi_T^{th}$  percentile of  $F(d)$  (Bartroff, Lai, and Shih, 2013).

Another commonly used method for early-phase clinical trials is the “3 + 3” design. This design involves the dose being administered group by group. The drug is not escalated within each group but between groups. Three patients are seen at a certain dose level and dose escalation for the next group occurs based on how many toxicities were seen in the previous group. The first group of three patients is seen on the lowest dose level and escalation occurs until there is an unacceptable rate of toxicity in the three patients. Escalation to the next highest dose occurs if none of the three patients experience a dose-limiting toxicity (DLT). However, if one of the patients in the group experience a DLT and only three patients have been seen at the current dose, then three more patients are enrolled at the current dose level. If two or more patients experience a DLT, then the trial is stopped and the next lower dose is determined to be the MTD. The algorithm also requires that at least six patients are treated at the MTD (Chevret, 2006).

The advantages of the “3 + 3” design are that it is easy to conduct and it does not need to assume a dose-response curve like model-based designs (described below). One disadvantage is that the algorithm-based method bases the recommended dose to the next group of patients using accumulated data only on the previous cohort(s) of patients instead of all the previously enrolled patients, resulting in poor statistical properties. The pitfalls of the “3 + 3” design are well-documented in the literature. Storer (1989) demonstrated that the “3 + 3” design is conservative and therefore treats too many patients at too low of a dose, which is not in accordance with the goals of a Phase I trial. Iasonos et al. (2008) performed an extensive simulation study which demonstrated the superiority of the Continual Reassessment Method (CRM) over the “3 + 3”, especially when the true dose is amongst the higher dose levels. Lastly, the method is not based on any underlying statistical properties and therefore the distribution of the doses recommended to be the MTD depends on the unknown dose-toxicity curve and the number of discrete dose levels (Cheung, 2011).

### 1.3.3 Model-based Designs

Another class of designs used in Phase I trials for cytotoxic agents are designs which assume some form of a dose-response curve. In these designs, a patient's toxicity response  $y$  to the drug at dose level  $d$  is commonly modeled by assuming  $y$  is a Bernoulli random variable where  $y = 1$  indicates that the patient experienced a DLT. The distribution of the response depends on the dose level  $d$  and an unknown parameter vector  $\beta$  through the function


$$F(d, \beta) = \Pr(y = 1 | \text{dose} = d).$$

In order to reflect the probability scale and the fact that the probability of toxicity is monotonically increasing, we assume that  $F(d, \beta)$  increases with  $d$ , approaching 0 as  $d \rightarrow -\infty$  and 1 as  $d \rightarrow \infty$ . The toxicity responses  $y_1, \dots, y_n$  are assumed to be independent in a sequential trial of  $n$  patients, except possibly through the dose levels  $x_1, \dots, x_n$  since the dose given to the  $j^{\text{th}}$  patient,  $x_j$ , will be chosen based on a function of the previous dose and response pairs  $(x_1, y_1), \dots, (x_{j-1}, y_{j-1})$ . The MTD is then the  $\phi_T^{\text{th}}$  percentile of the function  $F(d, \beta)$ . In particular, the  $\text{MTD} = F^{-1}(\phi_T, \beta)$  (Bartroff, Lai, and Shih, 2013).

#### **Continual Reassessment Method:**

The CRM is a model-based design that was first introduced by O'Quigley, Pepe, and Fisher (1990) as an alternative for the standard "3 + 3" algorithm based design which assumes a dose-toxicity curve in the form of  $F(d, \beta)$  where  $\dim(\beta) = 1$ . The basic idea of the CRM is that it uses information gathered from all previous patients to recommend a dose to the new patients. The patients for the CRM are entered sequentially and the dose-toxicity curve is refitted after the response of each patient is observed.

The original CRM assumes a dose-toxicity curve,  $F(d, \beta)$ , that is monotonic in



the dose,  $d$ , and the parameter,  $\beta$ . The curve can be defined by any number of functions in a “CRM class” such as an empiric model, logistic model, or hyperbolic tangent model. Figure 1.1 illustrates an example of one such dose-toxicity curve. The dose given to the  $j^{th}$  patient is a random variable,  $X_j = x_j$ , taking values  $x_j \in \{d_1, d_2, \dots, d_I\}$ . The true probability of toxicity at  $x_j$ ,  $\pi_T(x_j)$ , is modeled through the equation:

$$\pi_T(x_j) = \Pr(Y_j = 1 \mid X_j = x_j) = E(Y_j \mid x_j) \approx F(x_j, \beta) \quad (1.1)$$

for some one parameter model  $F(d, \beta)$  and  $\beta$  defined on the set  $\mathcal{B}$ , where  $Y_j$  is a Bernoulli random variable where 1 denotes a severe toxic response for the  $j^{th}$  patient and 0 otherwise. It is required that, for a fixed  $d$ ,  $F(d, \beta)$  must be strictly monotonic in  $\beta$ . Also, for a fixed  $\beta$ , it is required that, for a set of discrete dose levels  $d_i, i = 1, \dots, I$ ,  $F(d_i, \beta) > F(d_{i'}, \beta)$  when  $i > i'$ . Lastly, it is assumed that there exists some  $\beta_0$  from the parameter space for  $\beta$  such that we have  $F(d_\ell, \beta_0) = \phi_T$ , that is  $\beta_0$  reflects the underlying truth and the dose level  $d_\ell$  is the level required to achieve the desired target level of toxicity,  $\phi_T$ . Therefore,  $F(d_i, \beta_0) \leq \phi_T \forall i \leq \ell$

when the monotonicity assumption holds. A common example of the model in the literature is the empiric model  $F(d_i, \beta) \approx p_i^{\exp(\beta)}$ , where  $0 < p_1 < \dots < p_I < 1$  are the standardized units representing the discrete dose levels  $d_i$ . These values are referred to as the skeleton. Other models include a one-parameter logistic function,

$$F(d_i, \beta) = \frac{\exp(\alpha_0 + \beta \times d_i)}{1 + \exp(\alpha_0 + \beta \times d_i)}$$

where  $\alpha_0$  is a fixed constant, a two-parameter logistic function,

$$F(d_i, \alpha, \beta) = \frac{\exp(\alpha + \beta \times d_i)}{1 + \exp(\alpha + \beta \times d_i)},$$

and a hyperbolic tangent model,

$$F(d_i, \beta) = \left( \frac{\tanh d_i + 1}{2} \right)^{\exp(\beta)}.$$

The original CRM uses a Bayesian approach to model the probability of toxicity where the model parameter  $\beta$  is assumed random. Therefore, a prior distribution,  $g(\beta)$ , must be specified. In this dissertation, we will focus on a normal prior distribution

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

where  $\hat{\beta}_0$  and  $\sigma_\beta^2$  are the prior mean and variance and was introduced in the context of the CRM by O'Quigley and Shen (1996).

Once we have chosen a model, the CRM begins a trial by allocating the first patient to the prior MTD. Specifically,  $x_1 = d_i$  such that  $\arg \min_i \Delta(\hat{\pi}_{T_{i1}}, \phi_T)$  where  $\hat{\pi}_{T_{i1}} = F(d_i, \hat{\beta}_0)$ , and  $\Delta(\hat{\pi}_{T_{i1}}, \phi_T)$  is some loss function such as

$$\Delta(\hat{\pi}_{T_{i1}}, \phi_T) = |\hat{\pi}_{T_{i1}} - \phi_T|.$$

The remaining dose allocation decisions are determined based on data in the form of the set  $\Omega_j = \{(x_l, y_l) : l < j\}$ , which reflects the previous patients' doses and toxicity outcomes. At the enrollment of each new patient, the dose given is determined based on the current model-based MTD estimate, given  $\Omega_j$ . In particular, we determine the dose given to the  $j^{th}$  patient,  $x_j$ , to be

$$x_j = \arg \min_i \Delta(\hat{\pi}_{T_{ij}}, \phi_T) \quad (1.2)$$

where  $\hat{\pi}_{T_{ij}} = F(d_i, \hat{\beta}_j)$ ,

$$\hat{\beta}_j = \frac{\int_{\mathcal{B}} \beta L(\beta | \Omega_j) g(\beta) d\beta}{\int_{\mathcal{B}} L(\beta | \Omega_j) g(\beta) d\beta}$$

is the posterior mean of  $\beta$  given the set  $\Omega_j$ , and

$$L(\beta | \Omega_j) = \prod_{l=1}^{j-1} \{F(x_l, \beta)\}^{y_l} \{1 - F(x_l, \beta)\}^{(1-y_l)}$$

is the binomial likelihood. The remaining dose decisions are determined sequentially through Equation 1.2 until all patients have been treated at a dose and the response for each patient has been recorded. Then the recommended MTD will be the dose  $d_i$  such that  $\Delta(\hat{\pi}_{T_{i,n+1}}, \phi_T)$  is minimized, which means that it will be the dose given to the hypothetical  $(n+1)^{th}$  patient (Cheung, 2011).

In the years since the introduction of the CRM by O'Quigley, Pepe, and Fisher (1990) for cytotoxic agents there have been many adaptations on the original method to make the design more practical and extend the method into new areas. In contrast to the original Bayesian framework, O'Quigley and Shen (1996) extended the design to take into account a likelihood approach. Goodman, Zahurak, and Piantadosi (1995) extended the CRM to handle group accrual while, along with Korn et al. (1994) and Faries (1994), proposed a modification to the CRM to limit escalation to

no more than one level at a time. Various stopping rules due to too much toxicity for the CRM have been proposed but were first introduced by O’Quigley and Reiner (1998) and Heyd and Carlin (1998). An important extension of the CRM for this dissertation is that of Wages, Conaway, and O’Quigley (2011) who extended the CRM to combinations of agents by introducing the partial order CRM (POCRM). Yin and Yuan (2009) extended the CRM by combining it with Bayesian model averaging to eliminate some of the uncertainty in the choice of the skeleton. Iasonos et al. (2008) and Garrett-Mayer (2006) give recent overviews of the various other extensions of the CRM to more complicated clinical trials.

## 1.4 New Paradigm of Dose-finding

There has been a recent shift to more novel therapies, which include targeted agents. Significant discussion has taken place in the literature, Korn et al. (2001), Korn (2004) and Parulekar and Eisenhauer (2002), emphasizing the need for trial designs different from cytotoxic agents to be implemented for targeted agents. Targeted agent therapies offer potentially low toxicity drug profiles with potentially more efficacy. Since these agents are designed to target certain cells or pathways they may not bring about clinically significant toxicity. Using this fact Jain et al. (2010) demonstrated that in trials for targeted agents, patients who are treated at lower dose levels do not necessarily fare worse like in the cytotoxic agent setting. Therefore, more relevant endpoints need to be used in early-phase trials determining the dose to carry forward into a Phase II clinical trial. In these agents, the assumption of a higher dose leading to a higher efficacious response may not hold. Due to the lack of significant toxicity, the goal in these trials is no longer one of identifying a maximum tolerated dose but one of identifying some Optimal Biologic Dose (OBD), which is defined as a dose with the highest level of efficacious benefit to a patient



and a tolerable level of toxicity.

### 1.4.1 Locating the OBD

As noted in the previous section, the assumption of monotonicity with regards to dose-toxicity relationships is usually appropriate for targeted agents, yet this assumption for dose-efficacy curves may fail for targeted agents. For these types of therapies, dose-efficacy curves could follow a non-monotone pattern, such as a unimodal or plateau relationship. Let  $\pi_E(d_i) = \Pr(Z_j = 1|d_i)$  be the true probability of efficacy at dose  $d_i$  where  $Z_j$  equals 1 when the  $j^{th}$  subject experiences an efficacious response and is 0 if the  $j^{th}$  subject does not experience an efficacious response. In the presence of unimodal (including monotone increasing and monotone decreasing) or plateau dose-efficacy relationships, the primary objective of the trial is to find the OBD,  $d_\nu$ , defined such that

$$\pi_E(d_1) \leq \cdots \leq \pi_E(d_\nu) \geq \cdots \geq \pi_E(d_I). \quad (1.3)$$

In other words,  $\pi_E(d_\nu)$  corresponds to the dose  $d_\nu$  where the peak of the unimodal relationship occurs, or where the dose-efficacy curve begins to plateau. Using the terminology of Hwang and Peddada (1994),  $\pi_E(d_\nu)$  is said to be a *nodal* parameter (or a node) in that it is known that  $\pi_E(d_\nu) \geq \pi_E(d_i)$  for all  $i \neq \nu$ . Order restriction (1.3) is referred to as an *umbrella* ordering with the node of the umbrella occurring at  $d_\nu$ . For instance, if the node occurs at the highest dose, we have a monotone increasing dose-efficacy relationship. In actuality, we do not know where in the dose range this node occurs and we must account for this uncertainty. The node could occur at any of the  $I$  available dose levels, with non-decreasing efficacy probabilities for doses before the node and non-increasing efficacy probabilities for doses after the node. The strategy we propose is to use this information to formulate a set

Table 1.1: Possible dose-efficacy relationships for  $I = 4$  dose levels

Shape of relationship	Location of node	Possible Dose-efficacy relationship
Unimodal or plateau	$d_1$	$\pi_E(d_4) \leq \pi_E(d_3) \leq \pi_E(d_2) \leq \pi_E(d_1)$
Unimodal or plateau	$d_2$	$\pi_E(d_1) \leq \pi_E(d_3) \leq \pi_E(d_4) \leq \pi_E(d_2)$
Unimodal or plateau	$d_3$	$\pi_E(d_1) \leq \pi_E(d_2) \leq \pi_E(d_4) \leq \pi_E(d_3)$
Unimodal	$d_4$	$\pi_E(d_1) \leq \pi_E(d_2) \leq \pi_E(d_3) \leq \pi_E(d_4)$

of possible dose-efficacy relationships and make progress by appealing to Bayesian model choice, in order to account for the uncertainty surrounding the shape of the true dose-efficacy curve.

As an example, we describe the Phase I trial of Gerber et al. (2011) that assessed the safety and pharmacokinetics of a vascular targeting agent, bavituximab, in patients with advanced solid tumors. The trial investigated four dose levels (0.1, 0.3, 1, 3 mg/kg) and the goal of the study was to dose escalate to the predicted biologically effective dose, rather than the MTD, because for monoclonal antibodies, the MTD may not correspond to optimal efficacy. Any of the four dose levels could be the node of the unimodal or plateau dose-efficacy curve, corresponding to the optimal dose. These four possible dose-efficacy relationships are provided in Table 1.1. There are other possibilities, in addition to those in Table 1.1, that could be considered for inclusion into this subset of possible relationships. For example, for a unimodal curve with node at  $d_3$ , another possibility is given by  $\pi_E(d_1) \leq \pi_E(d_4) \leq \pi_E(d_2) \leq \pi_E(d_3)$ . This notion of formulating a subset of possible curves has been undertaken by other authors (Conaway, Dunbar, and Peddada, 2004; Wages, Conaway, and O’Quigley, 2011) within the context of dose-finding in drug combinations. These papers assessed the possible loss of information in only using a subset of relationships and they found that, in general, using a subset did not drastically reduce the performance of their methods. An important idea to this approach is that we do not need to identify the “correct” relationship at every dose level. We just need one which is

“close” to correct, in that it is correct in specifying the node (optimal dose) of the umbrella ordering.

### 1.4.2 Literature Review

There are two ways of thinking within targeted agents about how to define the possible outcomes for bivariate methods. One camp is to consider all four possible outcomes including (i) no toxicity and no efficacy, (ii) no toxicity and efficacy, (iii) toxicity and efficacy, and (iv) toxicity and no efficacy. On the other hand, one can think of a model which uses trinary outcomes: (1) no efficacy and no toxicity, (2) efficacy and no toxicity, and (3) toxicity. The thinking behind the trinary outcomes is that if a patient experiences a toxic response, then the patient will not experience an efficacious response as well.

Hoering et al. (2011, 2013) proposed methods for assessing bivariate binary endpoints for targeted agents that use a standard “3 + 3” method to identify the MTD in a Phase I portion and then allocates patients in the Phase II portion to the recommended Phase I MTD and dose level below (2011). Then the authors (2013) used a three arm approach, consisting of the recommended dose (RD), the dose level below the RD, denoted  $(RD^-)$ , and the dose level above the RD, denoted  $(RD^+)$ , for the Phase II portion of the study and randomizes patients to the different arms based upon the estimated toxicity and efficacy probabilities. Let  $\phi_T$  represent the toxicity upper bound and  $\phi_E$  represent the efficacy lower bound. If  $\hat{\pi}_T(RD^+) \leq \phi_T$ ,  $\hat{\pi}_E(RD^+) \geq \phi_E$  and has the highest estimated efficacy, then the patient is assigned to the  $RD^+$  arm. However, if the patient is not assigned to the  $RD^+$  arm and if  $\hat{\pi}_T(RD) \leq \phi_T$  and  $\hat{\pi}_E(RD) \geq \hat{\pi}_E(RD^-) \geq \phi_E$ , then the patient is assigned to the RD arm. Otherwise, when  $\hat{\pi}_T(RD^-) \leq \phi_T$  and  $\hat{\pi}_E(RD^-) \geq \phi_E$  then the patient is assigned  $RD^-$  arm.

Thall and Cook (2004) expanded the idea proposed by Thall and Russell (1998) when they proposed a more general method which uses an efficacy-toxicity trade-off contour to allocate patients to the best dose within a set of admissible doses. Their design can accommodate cases in which there are trinary outcomes and cases involving all four bivariate binary outcomes. In the bivariate binary case, Thall and Cook (2004) assume that the marginal model for toxicity is

$$\pi_T(d_i, \beta) = \text{logit}\{\eta_T(d_i, \beta)\}$$

where  $\eta_T(d_i, \beta) = \mu_T + d_i\beta$ . The marginal model for efficacy is assumed to be

$$\pi_E(d_i, \theta_1, \theta_2) = \text{logit}\{\eta_E(d_i, \theta_1, \theta_2)\}$$

where  $\eta_E(d_i, \theta_1, \theta_2) = \mu_E + d_i\theta_1 + d_i^2\theta_2$ . Assuming this form of  $\eta_E$  allows for the method to account for non-monotone dose-efficacy relationships as well as monotone dose-efficacy curves.

Hunsberger et al. (2005) introduced a method for when no significant toxicity is expected from the targeted agent. Their method provides a simple design to identify the OBD using only the efficacy response. The design uses an idea similar to the “3 + 3” method and enrolls patients in cohorts of three. The algorithm starts at the lowest dose enrolling three patients and escalates when less than or equal to one response is observed and expands the number of patients at the current dose level to six patients when greater than one patient experiences a response. The escalation continues if less than four patients experience a response and the dose recommended for further testing is the lowest dose with greater than three out of six responses. They also propose a design which will stop the dose escalation if the probability of efficacy remains constant across dose levels. They deem that the efficacy rate is

Table 1.2: Hypothetical toxicity and Efficacy Scenario used to demonstrate a pitfall of the methods developed by Hoering et al. (2011, 2013)

	$d_1$	$d_2$	$d_3$	$d_4$	$d_5$	$d_6$
Toxicity	0.03	0.07	0.12	0.18	0.25	0.31
Efficacy	0.8	0.6	0.5	0.4	0.3	0.2

constant when the calculated slope of the dose-response curve at the highest dose levels is less than or equal to zero.

Zang, Lee, and Yuan (2014) proposed three methods to identify the OBD for targeted agents. The first proposed design assumes the dose-efficacy curve follows a Bayesian logistic model with a quadratic term in order to properly account for the unimodal or plateau relationship of targeted agents. The second design they proposed is a nonparametric approach which uses double-sided isotonic regression. Lastly, the third method proposed in the paper is in the spirit of a “semi-parametric” approach since it assumes a logistic model for the dose-efficacy curve but only fits the model at the current dose level and two adjacent doses.

### 1.4.3 Areas of Improvement in Targeted Agents

The Hoering et al. methods (2011, 2013) are problematic because they use the traditional “3 + 3” method as the Phase I component of the trial, but in targeted agents there is low toxicity and therefore targeting a dose with a certain level of toxicity does not match with the objective of the trial. This misalignment in objectives can lead to the methods proposed by Hoering et al. (2011, 2013) to miss the true OBD. For example, consider the hypothetical scenario presented in Table 1.2. The “3 + 3” will likely identify either dose  $d_3$  or dose  $d_4$  as the MTD since the toxicity target is 0.33, and then using the method proposed by Hoering et al. (2011, 2013) will take that dose and the surrounding doses. However in this scenario dose  $d_1$  is the OBD, which will most likely be disregarded after the Phase I portion of the trial and not

carried forward into the Phase II portion of the trial.

Methods similar to that of Hunsberger et al. (2005) have the drawback that they do not take toxicity into account. Even though targeted agents are assumed to have no significant toxicity, safety of the patients participating in the trial is of upmost importance ethically. Therefore, the design may accidentally expose patients to doses that are too toxic. The methods presented in this dissertation will take toxicity and efficacy data into account simultaneously in order to recommend a dose to the next enrolled patient.

## 1.5 Outline of Dissertation

This dissertation will focus on the design of early-phase dose-finding studies that involve agents in which the dose-efficacy curve plateaus or peaks in the dose range and does not necessarily adhere to the monotonicity assumption. In this dissertation, we address aspects of current designs for targeted agents that can be inefficient and ineffective. The proposed methods will be developed within the original continual reassessment method framework as well as lean upon recent advances in its methodology including the notion of partial ordering. It is believed that the methods presented will provide accurate estimates of the OBD.

Chapter 2 presents the specifics of the proposed method. It details the models, inference, and algorithm, used to extend the continual reassessment method to bivariate binary outcomes. Various simulation studies are presented in order to demonstrate the operating characteristics of the method and compare it to existing designs for targeted agents.

Chapter 3 provides reassurance in the use of the proposed method by proving the theoretical properties of the proposed design. These properties will provide a complement to the extensive simulations presented in Chapter 2. A discussion of

properties important in the clinical setting concludes this chapter.

Chapter 4 extends the proposed methodology of Chapter 2 to a case in which the outcomes are not observed in a reasonable time frame. The chapter discusses how the proposed method can accommodate time-to-event outcomes for both toxicity and efficacy. Lastly, some concluding remarks and areas of future research are presented in Chapter 5.

## Chapter 2

# Models and Inference: Bivariate Binary Outcomes

The method introduced is one which handles the case when there is monotonicity for toxicity and non-monotonicity for efficacy. This non-monotonicity is due to the fact that each dose can potentially be the most efficacious dose since efficacy may not increase with increased dose in targeted agents. Therefore, we must incorporate this uncertainty into our model. Another key difference between the proposed method and existing CRM methodology is that a set of acceptable doses is used instead of identifying a dose with a certain level of toxicity. We use the acceptable dose approach because we want to control toxicity so as not to expose patients to too toxic a dose while acknowledging the fact that the most efficacious dose can be any of the identified doses in the acceptable set.

Consider a trial investigating  $I$  doses,  $d_1, \dots, d_I$ . Toxicity and efficacy are binary random variables, where for each subject  $j$ , we observe:

$$Y_j = \begin{cases} 0 & \text{if no toxicity,} \\ 1 & \text{if toxicity} \end{cases} \quad Z_j = \begin{cases} 0 & \text{if no response,} \\ 1 & \text{if response.} \end{cases}$$



The dose given to the  $j^{th}$  entered patient,  $X_j$ ,  $j = 1, \dots, n$ , can be thought of as a random variable taking values  $x_j \in \{d_1, \dots, d_I\}$ .

## 2.1 Modeling Toxicity

The method starts out by assuming some dose-toxicity curve,  $F(d, \beta)$ , from a CRM class. Toxicity estimates are obtained through the model presented in Section 1.3.3 as developed by O’Quigley, Pepe, and Fisher (1990).

## 2.2 Models and Inference for Efficacy

As noted above, we are uncertain where the node of the dose-efficacy curve is amongst the discrete doses. For a given model,  $k$ ,  $k = 1, \dots, 2I - 1$  ( $I$  models to represent the peak of the dose-response curve at each of the  $I$  dose levels and  $I - 1$  to account for the plateau occurring at dose levels  $1, \dots, I - 1$ ; there are only  $I - 1$  plateau orderings since a plateau at the highest dose level corresponds to the peak being at the highest dose), we model  $\pi_E(d_i)$ , the true probability of efficacy at dose  $d_i$ , by

$$\pi_E(d_i) \approx G_k(d_i, \theta) = q_{ik}^{\exp(\theta)} \quad (2.1)$$

for a class of working dose-efficacy models,  $G_k(d_i, \theta)$  and  $\theta \in \Theta$ . Similar to that of toxicity  $0 < q_{1k} < \dots < q_{Ik} < 1$  and the  $q_{ik}$ ’s are the skeleton values of the  $k^{th}$  model for efficacy. We can incorporate more uncertainty in the model by taking into account any prior information available concerning the plausibility of each model. This prior information is incorporated by introducing prior weights,  $\tau(k) = \{\tau(1), \dots, \tau(2I - 1)\}$ , where each  $\tau(k) \geq 0$  and  $\sum_k \tau(k) = 1$ . If we had no reason to believe that any one model was more plausible than another, then the weights would be discrete uniform. At the enrollment of the  $j^{th}$  patient into the

trial, we have efficacy data in the form of the set  $\mathcal{D}_j = \{(x_l, z_l) : l < j\}$ . Then for each model we obtain an estimate of  $\hat{\theta}_{jk}$ . Once the prior probability distribution,  $g(\theta)$ , for the parameter  $\theta$  of each model has been specified, we obtain the likelihood for the parameter  $\theta$ . At the enrollment of the  $j^{th}$  patient, the likelihood for model  $k$  is given by

$$L_k(\theta|\mathcal{D}_j) = \prod_{l=1}^{j-1} \{G_k(x_l, \theta)\}^{z_l} \{1 - G_k(x_l, \theta)\}^{(1-z_l)}$$

which, for each model, can be used to obtain a summary value,  $\hat{\theta}_{jk}$ , for  $\theta$ . Then the posterior density for  $\theta$  given the set  $\mathcal{D}_j$  and the likelihood is given through the equation:

$$f_k(\theta|\mathcal{D}_j) = \frac{L_k(\theta|\mathcal{D}_j)g(\theta)}{\int_{\Theta} L_k(\theta|\mathcal{D}_j)g(\theta)d\theta}.$$

Using this information, we can establish the posterior probabilities of the models given the data as

$$\Pr(M_k|\mathcal{D}_j) = \frac{\tau(k) \int_{\Theta} L_k(\theta|\mathcal{D}_j)g(\theta)d\theta}{\sum_{k=1}^{2I-1} \tau(k) \int_{\Theta} L_k(\theta|\mathcal{D}_j)g(\theta)d\theta}.$$

The prior model probabilities,  $\tau(k)$ , are updated by the efficacy response data  $\mathcal{D}_j$ . It is expected that the more the data support model  $k$ , the greater its posterior probability will be. Thus, we appeal to sequential Bayesian model choice in order to guide allocation decisions. Each time a new patient is enrolled, the single model,  $k^*$ , with the largest posterior probability is chosen such that

$$k^* = \arg \max_k \Pr(M_k|\mathcal{D}_j), \quad k = 1, \dots, 2I - 1.$$

We then take the working model,  $G_{k^*}(d_i, \theta)$  associated with the model and obtain estimates of the efficacy probabilities at each dose through

$$\hat{\pi}_E(d_i) = G_{k^*}(d_i, \hat{\theta}_{jk^*})$$

where

$$\hat{\theta}_{jk^*} = \int_{\Theta} f_{k^*}(\theta | \mathcal{D}_j) \theta d\theta.$$

Dose assignments can then be made according to these estimates and  $\phi_T$ , which is redefined in the targeted agent case as the toxicity tolerance instead of the target toxicity rate defined in the original CRM.

An important feature of the design of the proposed method is its flexibility to work under different assumptions for the dose-efficacy curve. It is possible to incorporate various degrees of uncertainty about the shape of the dose-efficacy curve. If there is strong reason to believe that the dose-efficacy curve will not plateau in the dose range, then the method will only take  $I$  models into account, one for each location of the peak of the curve. On the other hand, if there is strong knowledge that the efficacy curve will plateau in the range of doses, then only the plateau models can be considered as possible models by the method.

## 2.3 Algorithm

At each interim dose decision, the method proposed here will allocate the next patient to the dose that is deemed the most efficacious from the doses that have an acceptable level of toxicity. Once the estimates for probability of toxicity and efficacy at each dose are obtained, an acceptable set is defined as those doses having an estimated level of toxicity less than  $\phi_T$ . At the enrollment of the  $j^{th}$  patient, we

define the acceptable dose set as

$$\mathcal{A}_j = \{d_i : \hat{\pi}_T(d_i) \leq \phi_T \ i = 1, \dots, I\}$$

where  $\hat{\pi}_T(d_i) = \hat{\pi}_{T_{ij}}$ . The next patient is then allocated to a dose based on estimated efficacy probabilities,  $\hat{\pi}_E(d_i)$ . Using the acceptable set notion allows us to exclude any dose that the model deems too toxic based on the toxicity upper bound. Since the acceptable set is redefined after the inclusion of each patient based upon the current toxicity estimates it is possible, once more data has been observed, for  $\mathcal{A}_j$  to include a dose that was previously excluded when a limited amount of data existed. This stage of the dose-finding algorithm depends on how much data has been observed in the trial. Note that the proposed design is “greedy” in nature since at each interim dose decision the dose given to the next patient is the current estimated OBD. However, greedy designs have a practical problem that, in some cases, little or no information is gained by choosing the current estimated OBD. This occurs when only a few doses have been experimented with and the design sticks on a suboptimal dose. Thall and Nguyen (2012) give merit to employing the use of this randomization when they note,

“The general phenomenon of a greedy sequential decision procedure becoming stuck at a suboptimal treatment is well-known. A common solution for this problem is to randomly assign some patients to suboptimal treatments. This distributes patients more evenly among treatments and consequently more is learned about the design space, often with a resulting improvement in the method’s reliability”

Therefore, the algorithm is broken up into an adaptive randomization (AR) phase at the beginning and then a maximization phase for the remaining portion of the trial.

### 2.3.1 Adaptive Randomization Phase

As noted above, the AR phase is designed to prevent the algorithm from locking-in on a dose early in a trial simply because the method doesn't have enough information about other doses. Therefore, we need experimentation at other levels in order to gather information on the toxicity and efficacy profiles. Due to this fact, we do not rely entirely on the maximization of the estimated efficacy probabilities to guide the dose allocation decisions but implement a randomized allocation algorithm to obtain information over many doses that are deemed acceptable. Based on the estimated efficacy probabilities,  $\hat{\pi}_E(d_i)$ , for doses in the acceptable set,  $\mathcal{A}_j$ , we calculate a randomization probability  $R_i$ ,

$$R_i = \frac{\hat{\pi}_E(d_i)}{\sum_{d_i \in \mathcal{A}_j} \hat{\pi}_E(d_i)},$$

and randomize the next patient to dose  $d_i$  with a probability of  $R_i$ . We only want to randomize for a subset of patients, just enough to allow information to accumulate on untried safe doses. In the simulations presented in this dissertation, the AR phase is set to a third of the total maximum sample size. For example, if we are running a trial with a sample size of 36 patients, then the AR phase will end after 12 patients. Other AR sample sizes will be explored. We denote the AR phase sample size as  $n_R$ .

### 2.3.2 Maximization Phase

The remaining patients' dose allocation will be determined through a maximization phase. After the adaptive randomization phase has ended after  $j$  patients then we define the acceptable set  $\mathcal{A}_j$  and allocate the next patient to the dose with the highest estimated efficacy within this set. Formally, the dose given to the  $j^{th}$  patient

is  $x_j = \arg \max_{d_i \in \mathcal{A}_j} \hat{\pi}_E(d_i)$ . At the end of the maximization phase, the OBD is defined as the dose that is given to the hypothetical  $(n + 1)^{th}$  patient as determined by the model.

### 2.3.3 Starting the Trial

In order to get the trial underway, we will choose the efficacy skeleton with the largest prior probability,  $\tau(k)$ , among the orders being considered. If several, or all, of the models have the same maximum prior probability, then the algorithm will choose at random from these models. Given the starting skeleton,  $k^*$ , for efficacy, the starting dose  $x_1$  is then chosen. Specifically, based on the toxicity skeleton,  $p_i$ , we define the acceptable set, a priori, to be

$$\mathcal{A}_1 = \{d_i : p_i \leq \phi_T; i = 1, \dots, I\}.$$

Based on the efficacy skeleton  $q_{ik^*}$ , for doses in  $\mathcal{A}_j$ , the algorithm calculates the randomization probability  $R_i$ ,

$$R_i = \frac{q_{ik^*}}{\sum_{d_i \in \mathcal{A}_1} q_{ik^*}}$$

and randomizes the first patient to dose  $x_1 = d_i$  with probability  $R_i$ .

## 2.4 Method Based Solely on Efficacy

Since there may not be clinically significant toxicity for targeted agents, it was suggested by Hunsberger et al. (2005) to determine dose decisions based solely on efficacy. Therefore, we present a method which only determines interim dose decisions based on the efficacy model presented in Section 2.2. The algorithm for this method will proceed in much the same way as presented in Section 2.3 with the

exception that there will be no acceptable set since the dose-toxicity curve is not taken into account. Specifically, we allocate the next patient to dose  $d_i$

$$R_i = \frac{\hat{\pi}_E(d_i)}{\sum_{d_i \in 1, \dots, I} \hat{\pi}_E(d_i)}$$

in the AR phase of the algorithm and allocate the  $j^{th}$  patient to

$$x_j = \arg \max_{d_i \in 1, \dots, I} \hat{\pi}_E(d_i)$$

in the maximization phase.

A simulation study comparing the slope-sign method of Hunsberger et al. (2005), which stops the dose escalation when the probability of efficacy remains constant across dose levels, and the local logistic (L-logistic) and Isotonic regression methods of Zang, Lee, and Yuan (2014) to the proposed method, denoted WTE, is presented in Table B.1 to demonstrate the ability of the proposed method to model efficacy on its own. In all the scenarios except Scenario 6, the proposed method outperforms the slope-sign method and even in this scenario the proposed method has a PCS of 29.4% as compared to the slope-sign PCS of 31.3%. However, the improvement of identifying the true OBD for the proposed method over the slope-sign method can be seen in Scenario 8 in particular where the PCS for each method is 74.6% and 45.7%, respectively. The proposed method is competitive with the L-logistic design from Zang, Lee, and Yuan (2014) and outperforms the method in Scenarios 7 and 8. It is interesting to note that the proposed method significantly outperforms the nonparametric isotonic regression method in most of the presented scenarios, particularly Scenario 3 where the OBD is at the highest dose level. This fact demonstrates the improvement gained by assuming some parametric model for the dose-efficacy curve. Even in Scenarios 6 and 8, where the isotonic regression slightly outperforms

the proposed method, the difference is not very large, about 5% worse in the proposed method, and the proposed method still identifies the OBD in a large portion of the simulated trials. Scenario 6 is an interesting situation where the dose-efficacy curve plateaus at dose level  $d_2$ . Since we are assuming there is no clinically significant toxicity there is no way clinical way to break the tie between the last four doses. However,  $d_2$  is considered the OBD due to economic reasons since increased dosage means increased costs to the drug companies.

## 2.5 Simulation Studies of Proposed Method

### 2.5.1 Illustration


In this section, we present the results from a single trial. Consider a scenario as presented in Thall and Cook (2004), where there are four possible dose levels,  $\{d_1, d_2, d_3, d_4\}$ . The true toxicity probabilities are  $\pi_T(d_1) = 0.05$ ,  $\pi_T(d_2) = 0.15$ ,  $\pi_T(d_3) = 0.42$ , and  $\pi_T(d_4) = 0.65$  and the true efficacy probabilities are  $\pi_E(d_1) = 0.25$ ,  $\pi_E(d_2) = 0.65$ ,  $\pi_E(d_3) = 0.50$ , and  $\pi_E(d_4) = 0.05$ . The toxicity tolerance is set to  $\phi_T = 0.40$  and therefore the underlying true OBD is  $d_2$  since the true acceptable set is  $\mathcal{A} = \{d_i : \pi_T(d_i) \leq 0.40 \ i = 1, \dots, 4\} = \{d_1, d_2\}$  and  $\pi_E(d_2) > \pi_E(d_1)$ . The trial will treat a total of  $n = 36$  patients so the adaptive randomization size,  $n_R$ , is set to 12 patients. In order to get the trial underway, skeletons need to be specified for both the toxicity probabilities and the efficacy probabilities for each ordering. Let  $p_1 = 0.05$ ,  $p_2 = 0.20$ ,  $p_3 = 0.35$ , and  $p_4 = 0.45$  and we consider  $2 \times 4 - 1 = 7$  possible models for efficacy as presented in Table 2.1. Since we believe each of the seven models is equally likely, we will assume that, a priori,  $\tau(1) = \dots = \tau(7) = \frac{1}{7}$ . The results of the trial are presented in Table 2.2 along with an illustration in Figure 2.1, where an open circle represents subject  $j$  did



Table 2.1: Models for efficacy with 4 dose levels

Model (k)	Type	Dose Level			
		1	2	3	4
1	Peak	0.1	0.3	0.5	0.7
2	Peak	0.3	0.5	0.7	0.5
3	Peak	0.5	0.7	0.5	0.3
4	Peak	0.7	0.5	0.3	0.1
5	Plateau	0.3	0.5	0.7	0.7
6	Plateau	0.5	0.7	0.7	0.7
7	Plateau	0.7	0.7	0.7	0.7

Figure 2.1: Sequential trial of 36 patients



not experience a toxicity, a red-colored circle represents the subject experienced a toxicity, a dot within the circle represents that the subject experienced an efficacious response, and no dot within the circle represents no efficacious response.

When determining which dose to give to the first patient, we first determine which doses are deemed safe through the definition of the acceptable set  $\mathcal{A}_1 = \{d_i : p_i \leq 0.40\} = \{d_1, d_2, d_3\}$ . Since we assumed all of the models are equally likely, we randomly choose a single model, which in this trial is  $M_6$ . Therefore,  $\hat{\pi}_E(d_1) = 0.5$  and  $\hat{\pi}_E(d_2) = \hat{\pi}_E(d_3) = 0.7$ . Now that we have the estimated efficacy

Table 2.2: Sequential trial of 36 patients

$j$	$x_j$	$y_j$	$\hat{\beta}$	$z_j$	$\hat{\theta}_k$	$k$	$j$	$x_j$	$y_j$	$\hat{\beta}$	$z_j$	$\hat{\theta}_k$	$k$
1	3	1	0	1	0	6	20	3	1	-0.068	0	-0.128	1
2	1	0	-0.852	0	-0.465	2	21	2	0	-0.151	0	0.854	3
3	1	0	-0.553	0	-0.003	5	22	2	1	-0.113	0	0.907	3
4	1	0	-0.367	0	-0.14	1	23	2	0	-0.222	0	1.085	6
5	2	1	-0.239	1	-0.001	1	24	2	0	-0.186	1	0.054	1
6	1	0	-0.62	0	-0.384	1	25	2	0	-0.152	1	1.05	6
7	1	0	-0.51	0	-0.266	1	26	2	0	-0.121	1	0.87	3
8	2	0	-0.423	0	-0.175	1	27	2	1	-0.092	1	0.812	3
9	2	1	-0.308	1	-0.052	1	28	2	0	-0.181	1	0.759	3
10	2	0	-0.52	1	-0.265	1	29	2	0	-0.153	0	0.709	3
11	2	0	-0.421	0	-0.423	1	30	2	0	-0.126	1	0.752	3
12	1	0	-0.337	0	-0.318	1	31	2	0	-0.101	0	0.706	3
13	2	0	-0.295	1	-0.261	1	32	2	0	-0.077	1	0.747	3
14	2	0	-0.228	0	-0.384	1	33	2	0	-0.054	0	0.703	3
15	2	0	-0.169	1	-0.303	1	34	2	0	-0.033	1	0.742	3
16	3	0	-0.116	0	-0.407	1	35	2	0	-0.012	0	0.696	3
17	3	1	-0.052	0	-0.319	1	36	2	1	0.007	1	0.738	3
18	2	0	-0.152	0	0.811	6	37	2	0	-0.066	1	0.699	3
19	2	0	-0.108	0	0.877	6							

probabilities at each of the three dose levels in the acceptable set, we determine the randomization probabilities. Note that  $R_1 = \frac{0.5}{0.5 + 0.7 + 0.7} \approx 0.26$  and  $R_2 = R_3 = \frac{0.7}{0.5 + 0.7 + 0.7} \approx 0.37$ . Then, we randomly determine that the dose given to the first patient is  $d_3$ .

Given that the first patient now has data for toxicity and efficacy, we can begin modeling the dose-toxicity and dose-response curves. Following Section 1.3.3 to obtain an estimate of the parameter,  $\beta$ , we estimate the probabilities of toxicity through the equation  $\hat{\pi}_T(d_i) = p_i^{\exp(\hat{\beta})}$  at each dose to be  $\hat{\pi}_T(d_1) = 0.279$ ,  $\hat{\pi}_T(d_2) = 0.503$ ,  $\hat{\pi}_T(d_3) = 0.639$ , and  $\hat{\pi}_T(d_4) = 0.711$ . Therefore, the acceptable set is estimated to be  $\mathcal{A}_2 = \{d_1\}$  since  $d_1$  is the only dose that has an estimated probability of toxicity to be less than 0.4. Note that at this juncture, we are able to determine that the dose to be given to the third enrolled patient should be  $d_1$  and

thus no efficacy modeling is required since there is only one dose in the acceptable set.

At the enrollment of the third patient, the Bayesian posterior mean estimate for the  $\beta$  parameter in the model for toxicity was calculated to be  $-0.553$ . Therefore, the estimated probabilities of toxicity were  $\hat{\pi}_T(d_1) = 0.178$ ,  $\hat{\pi}_T(d_2) = 0.396$ ,  $\hat{\pi}_T(d_3) = 0.547$ , and  $\hat{\pi}_T(d_4) = 0.632$ , which leads to an acceptable set of  $\mathcal{A}_3 = \{d_1, d_2\}$ . The posterior model probabilities for efficacy were estimated to be  $\Pr(M_1|\mathcal{D}_2) = 0.191$ ,  $\Pr(M_2|\mathcal{D}_2) = 0.210$ ,  $\Pr(M_3|\mathcal{D}_2) = 0.100$ ,  $\Pr(M_4|\mathcal{D}_2) = 0.035$ ,  $\Pr(M_5|\mathcal{D}_2) = 0.210$ ,  $\Pr(M_6|\mathcal{D}_2) = 0.155$ , and  $\Pr(M_7|\mathcal{D}_2) = 0.098$ . Since  $M_2$  and  $M_5$  have the same posterior model probability, we randomly select one model, which in this case is  $M_5$ . Therefore, we obtain  $\hat{\theta}_k = -0.465$  which leads to estimated probabilities of efficacy of  $\hat{\pi}_E(d_1) = 0.301$ ,  $\hat{\pi}_E(d_2) = 0.501$ ,  $\hat{\pi}_E(d_3) = 0.701$ , and  $\hat{\pi}_E(d_4) = 0.701$  through the equation  $\pi_E(d_i) = q_{ik}^{\exp(\hat{\theta}_k)}$ . Now, the dose administered to the fourth entered patient was determined to be  $d_1$  based on randomization with the estimated probabilities of  $R_1 = 0.375$  and  $R_2 = 0.625$ .

The method continues in this manner through the end of the AR phase. Once the AR phase comes to a conclusion after patient 12, we enter the maximization phase. Toxicity and efficacy probabilities are modeled in the same way, except now we simply choose the dose level with the highest estimated efficacy within the acceptable set. At the enrollment of the 13<sup>th</sup> patient, the estimated toxicity probabilities were  $\hat{\pi}_T(d_1) = 0.100$ ,  $\hat{\pi}_T(d_2) = 0.289$ ,  $\hat{\pi}_T(d_3) = 0.445$ , and  $\hat{\pi}_T(d_4) = 0.541$ , which leads to an acceptable set of  $\mathcal{A}_{13} = \{d_1, d_2\}$ . Note that  $M_1$  is chosen and the estimated efficacy probabilities at doses 1 and 2 are  $\hat{\pi}_E(d_1) = 0.396$  and  $\hat{\pi}_E(d_2) = 0.586$ . Therefore,  $d_2$  is chosen as the dose administered to the 13<sup>th</sup> patient in the trial. The proposed method continues in this way until all 36 patients have been entered into the trial and the dose which would have been allocated to the hypothetical 37<sup>th</sup> patient is  $d_2$ . Therefore, the estimated OBD from this trial is determined to

be  $d_2$  with an estimated probability of toxicity of  $\hat{\pi}_T(d_2) = 0.222$  and an estimated efficacy probability of  $\hat{\pi}_E(d_2) = 0.488$ .

### 2.5.2 Results

A simulation study is shown below to demonstrate the performance of the proposed method as compared to the method developed by Thall and Cook (2004). In each set of 5000 simulations, in order to define the acceptable set  $\mathcal{A}$ , the toxicity tolerance was specified to be  $\phi_T = 0.40$ . Only the peak models from Table 2.1 were taken into consideration as possible candidate models by the method. The value of  $\phi_T$  and the set of candidate models were chosen in order to provide a justifiable comparison to Thall and Cook (2004). We ran simulations using outcomes generated with the five fixed values of  $\psi$  used in Thall and Cook (2004) where  $\psi$  represents the association parameter and  $\psi = 0$  corresponds with  $\pi_T$  and  $\pi_E$  being independent;  $\psi = \{-2.049; -0.814; 0; 0.814; 2.049\}$ . This analysis was performed to assess the sensitivity of the proposed method to differing values of the correlation between toxicity and efficacy. The sample size for each simulated trial was set to  $n = 36$ .

The performance of the proposed method, which is denoted WT, and the design of Thall and Cook (2004), denoted TC, is summarized in Table B.2. Table B.2 reports the percentage of trials in which each dose was selected as the OBD at the end of each trial for each of the Scenarios 1-5. The true (toxicity, efficacy) probabilities for each dose are given next to the number of the scenario. In Scenario 1, dose level 4 is the only overly toxic dose and dose 3 is defined as the OBD having the largest efficacy among those with tolerable toxicity. The performance of the proposed method and TC are comparable in this scenario, with the proposed method recommending the OBD in a slightly higher percentage of trials than TC (max 86.0 v. 80.3). However, TC does a better job of not recommending the

overly toxic dose 4, doing so approximately 5% less than the proposed method. In Scenarios 2 and 3, the performance of the two methods is nearly identical across the dose range in terms of recommendation percentages. In Scenario 4, TC does a little better in not recommending the dose level above the OBD (4.8) compared to the proposed method (max 7.5). For this scenario, this is not a major concern for the proposed method since it still recommends the OBD in a large majority of trials and the second dose and none of the other doses have true toxicity probabilities larger than  $\phi_T$ . Therefore, the method proposed did not expose patients that were not treated at the OBD to overly toxic levels.


Although across nearly all of the scenarios considered, the performance of WT and TC (for  $\psi = 0$ ) is quite similar, some differences need to be noted. In Scenario 5, the performance of TC is superior. The true optimal dose is at level 1 with true (toxicity, efficacy) probabilities (0.05, 0.45). TC recommends this dose in 72.9% of trials, while WT does so in a maximum of 46.5% of trials. WT selects dose level 2, with true probabilities (0.45, 0.50) in 53.3% of trials. This dose has a true toxicity probability just above the toxicity tolerance value  $\phi_T = 0.40$  used to define the acceptable set, and it appears as though the toxicity probability estimates include dose level 2 often enough to where it is recommended as the optimal dose in a high percentage of trials. We reran Scenario 5 with  $\phi_T = 1/3$  and the previous recommendation percentages (46.5, 53.3) for the first two dose levels became (64.2, 35.7), improving the performance in this particular scenario from when  $\phi_T = 0.40$ . Despite this improvement, it appears as though TC may have an advantage when a dose has a true toxicity probability close to, but above,  $\phi_T$ .

The simulation results indicate that an advantage of the proposed method lies in its robustness to the specification of the association between toxicity and efficacy. As is evident from Table B.2, the performance of the proposed method is practically unchanged for all values of  $\psi$ . In Thall and Cook (2004), the authors report

recommendation percentages for the various values of  $\psi$  for Scenario 2 only. When  $\psi = 0$ , the recommendation percentages for their method for doses 3 and 4 were (13.6, 86.0). For  $\psi = -2.049$ , these two percentages become (64.0, 35.6); for  $\psi = -0.814$ , (69.4, 30.4); for  $\psi = 0.814$ , (40.4, 41.1); and for  $\psi = 2.049$ , (82.0, 15.8). With dose 4 being the OBD, it is desirable that it be recommended in a high percentage of trials. The recommendation percentage for dose 4 drops from 86.0% when  $\psi = 0$  to, at best, 41.1% for a non-zero value of  $\psi$ . Thall and Cook (2004) point out that a strong positive or negative association between toxicity and efficacy will diminish the “likelihood of acceptability of all doses.” The dependence of Thall and Cook (2004) on the association parameter can be contributed to the fact that they consider  $\psi$  a parameter in their model and therefore assumed a prior of  $N(0, 1)$  on  $\psi$  even when  $\psi = -2.049$ . However, the fact that the proposed method does not depend so heavily on the association parameter makes it an attractive alternative to the TC method. Overall, Table B.2 indicates that the proposed method is competitive with that of Thall and Cook (2004). Another advantage of the proposed method over that proposed by Thall and Cook (2004) is its relative simplicity compared to the TC method. Thall and Cook (2004) noted their method “is somewhat more structured than most dose-finding methods, and thus requires the statistician to work harder.” The growing familiarity of the CRM in early-phase clinical trials contributes to the fact that the proposed method is well understood by clinicians and review boards.

Another important comparison that demonstrates further merit of the proposed method is a comparison to the method proposed by Hoering et al. (2013) for identifying the OBD in a trial investigating  $I = 6$  dose levels. The comparison assesses the performance of each method over the twelve true toxicity/efficacy scenarios contained in their paper, which are provided in Figure A.1. True toxicity probabilities are denoted T1 through T4, while true efficacy probabilities are denoted R1 through

Figure 2.2: Simulation results comparing the proposed method to the method introduced by Hoering et al. (2013).



R3. In each set of 1000 simulated trials, the maximum sample size was set to  $n = 64$ . The Phase II portion of Hoering et al. (2013) randomized  $n_2 = 48$  patients to three doses levels (arms); (1) the recommended dose (RD) from Phase I, (2) the dose level immediately above the RD (RD+), and (3) the dose level immediately below the RD (RD-). For response, the authors test competing hypotheses  $H_0 : p = 0.05$  vs.  $H_1 : p = 0 : 30$ . The Phase I sample sizes for the Hoering et al. (2013) method

are reported in Table 1 of Hoering et al. (2011), with a minimum sample size of  $n_1 = 16$ . Therefore, we thought a justifiable total sample size for comparison was  $n = 16 + 48 = 64$ . For the proposed method, the size of the AR phase was set equal to one quarter the total sample size so that  $n_R = 16$ . In a subsequent section, we investigated the impact of various sizes,  $n_R$ , of the AR phase on operating characteristics.

The proposed design embodies characteristics of the CRM so we can utilize these features in specifying design parameters. The toxicity probabilities were modeled via the empiric model given in Section 1.3.3 with skeleton values  $p_i = \{0.01, 0.08, 0.15, 0.22, 0.29, 0.36\}$ . It has been shown (O’Quigley and Zohar, 2010) that CRM designs are robust and efficient with the implementation of “reasonable” skeletons. Simply defined, a reasonable skeleton is one in which there is adequate spacing between adjacent values. This skeleton also reflects the idea that, for targeted agents, it is often assumed that there is minimal toxicity over the dose range. For efficacy, probabilities were modeled via the class of empiric models in Section 2.2, using  $2I - 1 = 11$  skeletons that correspond to the possible dose-efficacy relationships. The six sets of values used for the unimodal relationships were:

1.  $q_{i1} = \{\mathbf{0.60}, 0.50, 0.40, 0.30, 0.20, 0.10\}$
2.  $q_{i2} = \{0.50, \mathbf{0.60}, 0.50, 0.40, 0.30, 0.20\}$
3.  $q_{i3} = \{0.40, 0.50, \mathbf{0.60}, 0.50, 0.40, 0.30\}$
4.  $q_{i4} = \{0.30, 0.40, 0.50, \mathbf{0.60}, 0.50, 0.40\}$
5.  $q_{i5} = \{0.20, 0.30, 0.40, 0.50, \mathbf{0.60}, 0.50\}$
6.  $q_{i6} = \{0.10, 0.20, 0.30, 0.40, 0.50, \mathbf{0.60}\}$

The five sets of values for the plateau relationships were:



7.  $q_{i7} = \{0.20, 0.30, 0.40, 0.50, \mathbf{0.60}, \mathbf{0.60}\}$
8.  $q_{i8} = \{0.30, 0.40, 0.50, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}\}$
9.  $q_{i9} = \{0.40, 0.50, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}\}$
10.  $q_{i10} = \{0.50, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}\}$
11.  $q_{i11} = \{\mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}\}$

We assumed, a priori, that each of these eleven models was equally likely and set  $\tau(k) = 1/11$ .

The results of the comparison are summarized in Table B.3 and Figure 2.2, with each bar in the figure representing the proportion of simulated trials that each method selected a good dose [Figure 2.2(a)] and the best dose [Figure 2.2(b)]. For each scenario, using the definitions of Hoering et al. (2013), the best dose is defined as the dose level(s) that maximizes efficacy while assuring safety, and a good dose as a level where efficacy is above some pre-specified efficacy threshold while also not being too toxic. In Figure A.1, a blue circle represents a good dose, while an orange circle denotes a best dose. When there is only an orange circle, then the good and best doses coincide with each other. The proposed method is in green while the method developed by Hoering et al. (2013) is in black. Overall, the results indicate a superior performance for the proposed method over that of Hoering et al. (2013). The left column of Figure A.1 and the first four rows of Table B.3 represent cases in which the dose-efficacy curves have monotone increasing probabilities. In the first three of these scenarios (R1T1, R1T2, and R1T3),  $d_4$  is the only good and best dose. WT outperforms Hoering et al. (2013) in these scenarios by selection percentages of 39%, 32% and 9%, respectively. In R1T4, the toxicity probabilities plateau at a value of 0.05 across the dose range, so the highest dose,  $d_6$ , is considered the best dose. The Hoering et al. (2013) method selects the best dose as the OBD

in 63% of simulated trials, compared to 54% by the proposed method. The next four rows of Table B.3 and the middle column of Figure A.1 illustrate situations in which the dose-efficacy curves increase until dose  $d_4$  and remain constant after  $d_4$  (plateau). Again, in terms of selecting both good and best doses, the proposed method improves upon the selection percentages of Hoering et al. (2013) in most cases (R2T1, R2T2, R2T4), by as much as 36% in R2T1. Hoering et al. (2013) performs better in selecting a good dose in R2T3, doing so in 71% of trials compared to 59% by the proposed method. Finally, in the last four rows of Table B.3 and the right column of Figure A.1, most of the doses are considered to be best/good doses, so, as expected, both methods produce high selection percentages in these cases. The proposed method does perform slightly better in these last four scenarios, yielding selection percentages that range from 3% to 13% higher than those produced by Hoering et al. (2013). We feel that this improved performance by the proposed method across the twelve scenarios considered here can mostly be attributed to Hoering et al. (2013)'s use of a Phase I portion to identify an MTD. This may not be appropriate for targeted agents due to the uncertainty surrounding the shape of the dose-efficacy curve, especially in scenarios similar to that presented in Table 1.2. If the MTD-based Phase I portion does not identify doses around the OBD, then the Phase II portion could possibly be randomizing to incorrect doses. In contrast, throughout the duration of the trial, our method continuously monitors toxicity and efficacy data to identify a set of safe doses, within which we allocate to a dose with high efficacy. This allows for the adaptive updating of the doses being considered after each patient inclusion, rather than restricting attention to two or three doses after a Phase I component designed to find the MTD.

The proposed method is superior in almost all cases when compared to Hoering et al. (2013). It is also competitive with Thall and Cook (2004) except in cases where a dose lies close to, but above, the toxicity tolerance  $\phi_T$ . However, the use of

the two-parameter logistic model will be explored in order to handle such scenarios. Overall, the strong showing of our method against published work in the area in extensive simulation studies makes us feel confident in recommending it.

## 2.6 Design Specifications

There are aspects of the design which require further exploration in order to provide merit to the way the algorithm is currently performed. Additionally, this section will demonstrate the robustness of the method in certain cases. Studies will include:

1. Adaptive Randomization Sample Size
2. Skeleton Choice
3. Working Model Choice
4. Utility Functions
5. Model Selection Criteria.

### 2.6.1 Adaptive Randomization Sample Size

It was stated above that  $n_R$  will be set to one-third of the total sample size but  $n_R$  could be set to any number between 0 and  $n$ . The simulation study below demonstrates that as long as some mixture of adaptive randomization and maximization phases are utilized, then one can contract or expand the amount of patients in the adaptive randomization phase and the operating characteristics of the method will be relatively unaffected. In the study presented,  $n_R$  takes values in  $\{0, 12, 18, 24, 36\}$ .

Four scenarios are presented in Table B.4. All design specifications, other than  $n_R$  are held constant throughout the study in order to allow for a direct comparison of the results. The results show the percentage of simulated trials which identify each of

the doses out of 1000 simulations. The adaptive randomization sample size is varied from all maximization ( $n_R = 0$ ) to all randomization ( $n_R = 36$ ). In the first scenario, the percentage correct selection (PCS), which measures how often the true OBD (indicated in bold) is chosen, ranges from 48.1% to 52.9% which demonstrates the robustness of the choice of  $n_R$ . As expected, the results demonstrate that using all randomization hinders the method’s ability to correctly identify the OBD, especially in the third scenario where the PCS is only 25.7% when  $n_R = 36$  which is about 9% less than the worst mixture of randomization and maximization. Based on simulation studies it is recommended to use a randomization sample size between one-fourth and one-half of the total sample size. However, the size of the AR phase can be expanded or contracted according to the desire of the clinician / statistician team and operating characteristics will be relatively unaffected.

### 2.6.2 Skeleton Choice

An important question to study is how robust the proposed method is to the choice of the skeleton, both for toxicity and efficacy. This topic has been vigorously studied in the context of the CRM by O’Quigley and Zohar (2010), Lee and Cheung (2009), and Daimon, Zohar, and O’Quigley (2011). They found that as long as a “reasonable” skeleton is used, then the CRM is robust and efficient. A “reasonable” skeleton is one in which there is adequate spacing between adjacent dose levels. The point of the spacing between adjacent dose levels is to allow the CRM to be able to identify each dose level as the correct dose dependent on the underlying truth.

Four skeleton choices for toxicity are presented in the simulation study presented in Table B.5. Additionally, two different efficacy skeleton choices for the models are studied. The toxicity skeletons explored were:

1.  $\alpha_1 = \{0.05, 0.20, 0.35, 0.45\}$

2.  $\alpha_2 = \{0.20, 0.30, 0.45, 0.60\}$
3.  $\alpha_3 = \{0.10, 0.25, 0.40, 0.55\}$
4.  $\alpha_4 = \{0.01, 0.08, 0.15, 0.22\}$ .

The first toxicity skeleton reflects the notion that, a priori, we believe that three out of the four doses have an acceptable level of toxicity. Skeletons 2, 3, 4 reflect the notion that we believe that two, three, and four doses, respectively, have an acceptable level of toxicity, a priori. One of the efficacy skeletons chosen for the models are the skeletons presented in Table 2.1, denoted  $Q_1$  in Table B.5. The other is a variation, denoted  $Q_2$ , with the following models:

1.  $q_{i1} = \{0.10, 0.20, 0.30, 0.50\}$
2.  $q_{i2} = \{0.20, 0.30, 0.50, 0.30\}$
3.  $q_{i3} = \{0.30, 0.50, 0.30, 0.20\}$
4.  $q_{i4} = \{0.50, 0.30, 0.20, 0.10\}$
5.  $q_{i5} = \{0.10, 0.20, 0.50, 0.50\}$
6.  $q_{i6} = \{0.30, 0.50, 0.50, 0.50\}$
7.  $q_{i7} = \{0.50, 0.50, 0.50, 0.50\}$

The toxicity skeletons all perform fairly similarly across the four scenarios. In Scenario 4, the four toxicity skeletons perform almost identically with the maximum PCS of 92.1% occurring for  $\alpha_2$  and the minimum PCS of 90.7% occurring for  $\alpha_1$  and  $\alpha_3$ . The biggest discrepancy in the performance occurs in Scenario 2 when the OBD is at the highest dose  $d_4$ . In this situation, skeleton  $\alpha_2$  has the hardest time of identifying the true OBD since  $p_4 = 0.60$ . This large value is harder to overcome

than in  $\alpha_1$  and  $\alpha_4$  where  $p_4$  takes the value of 0.45 and 0.22, respectively. Therefore, the method tends to treat patients at less toxic doses for this skeleton specification. It is not that surprising that  $\alpha_4$  outperforms the other toxicity skeletons since it assumes, a priori, that all doses have an acceptable level of toxicity. Even when there is a discrepancy, the proposed method still recommends the OBD in a large percentage of the trials so there is robustness to the choice of the toxicity skeleton, as long as reasonable skeletons are selected.

From the results in Table B.5, we see that the method is robust against the choice of the efficacy skeleton values as well. The method performs almost identically under efficacy skeletons  $Q_1$  and  $Q_2$  in Scenarios 1 and 2. There is a minor difference in Scenarios 3 and 4 where  $Q_2$  does a better job of not treating patients at doses higher than the OBD. In particular,  $Q_2$  only recommends dose  $d_3$  5.5% of the time as compared to  $Q_1$  recommending the dose 10.6% of the time. Overall, the method is fairly robust against the choice of the efficacy skeleton used to obtain the model-based OBD.

### 2.6.3 Working Model Choice

As mentioned in Section 1.3.3, there exists a variety of choices for the working model in order to model toxicity and efficacy. Therefore, a simulation study analyzing the operating characteristics of the proposed method under different choices of the working model from within the CRM class including

1. Empiric Model
2. One-parameter Logistic
3. Two-Parameter Logistic

is performed. We again used the scenarios from Thall and Cook (2004) to evaluate the operating characteristics of the proposed method under these different working models. The toxicity tolerance rate is set to  $\phi_T = 0.40$  and the PCS of correctly identifying the OBD is presented on the basis of 1000 simulated trials.

The skeleton values were held consistent to previous simulations and were set to  $p_1 = 0.05$ ,  $p_2 = 0.20$ ,  $p_3 = 0.35$ , and  $p_4 = 0.45$  for the empiric model. In the two logistic models, the dose scale is different from the empiric model so the skeleton values were rescaled by computing

$$d_i = \log \frac{p_i}{1 - p_i} - \alpha_0$$

where  $\alpha_0$  is a fixed constant for the one-parameter model and

$$d_i = \log \frac{p_i}{1 - p_i} - \bar{\alpha}$$

where  $\bar{\alpha}$  is the posterior mean for the intercept parameter  $\alpha$  for the two-parameter model. Similar rescaling was also performed on the efficacy skeletons  $q_{ik}$  in order to model the probability of efficacy. We obtained the parameter estimates in the two-parameter setting using the function `bayesglm` in the R package `arm` (R Core Team). In order to account for the different degrees of the monotonicity assumptions for toxicity and efficacy, a different prior is used on the intercept parameter in both logistic models. We need to specify a vague normal prior with a mean of 0 and a variance of 1.34 for efficacy in order for the model to be flexible enough to account for the uncertainty in the shape of the dose-efficacy curve. However, when we model toxicity, we want to exclude models where the dose-toxicity curve is not monotonically increasing. Therefore, we need to specify the prior on  $\beta$  as normal

with a mean of 1 and set the variance parameter,  $\sigma^2$ , such that

$$0.975 = \int_{-\infty}^0 \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-1)^2}{2\sigma^2}} dx.$$

Specifically, we set  $\sigma^2 = 1.3237754$  and there will be a small, a priori, probability that the slope parameter is negative.

Table B.6 presents a simulation study of these three working models. The empiric model and one-parameter logistic functions behave similarly for Scenarios 1-4 which demonstrates some robustness against the choice of the working model in the proposed method. However, the main difference occurs in Scenario 5 where the true optimal dose is at level 1 with true (toxicity, efficacy) probabilities (0.05, 0.45) and dose level 2 has true probabilities (0.45, 0.50). In this scenario, the one-parameter logistic function, recommending the true OBD in 66.5% of the trials, significantly outperforms the empiric model which only recommended dose 1 as the OBD in only 46.5% of the trials. Therefore, if there is a strong prior belief that the dose above the OBD is just above the toxicity tolerance, then we recommend using the logistic working model in order to capture this effect. The two-parameter logistic function is of particular interest because its use may be suited for scenarios where the toxicity of a dose lies close to, but just above, the toxicity tolerance. This is due to the fact that a two-parameter logistic function is more capable of capturing the entire shape of the curve, whereas the one-parameter function is best suited for identifying a target dose. However, the results in Table B.6 provide caution from using the two-parameter logistic function in practice and reinforce the notion that the under-parameterized working models outperform the multiple parameter working models in the CRM context in most situations. The performance in Scenario 1 is particularly discouraging since it only recommends level 3 as the OBD in 39.0% of the trials as compared with a minimum of 71.1% of the trials for the two one-parameter



working models. On the other hand, the use of the two-parameter logistic function is particularly encouraging in Scenario 5 where it recommends level 1 as the OBD in 92.9% of the trials.

Due to the different natures of the toxicity and efficacy responses, the possibility of using different working model choices for each response was explored. In particular, the combined use of the two-parameter logistic function to model one response and the one-parameter logistic function to model the other response was explored using three cases: (1) model toxicity and efficacy using the two-parameter and one-parameter working models, respectively, (2) model toxicity and efficacy using the one-parameter and two-parameter working models, respectively, and (3) model toxicity and efficacy both using the two-parameter logistic working model. The motivation for this simulation is the goal for modeling toxicity is to identify a set of doses that satisfy a toxicity tolerance can be thought of as targeting a certain threshold and then uses all doses below the identified dose since it is assumed that the toxicity probabilities are monotonically increasing or can be framed as wishing to model the entire curve. Also, the goal when modeling efficacy can be framed as choosing the dose with the highest efficacy from within the tolerable doses (modeling the entire curve) or viewed as the targeting the efficacy rate of 1. Table B.7 presents the results analyzing the effect of the use of the two-parameter logistic function for cases (1), (2), and (3). In general, case (1) is the most viable alternative across all scenarios to the use of the one-parameter logistic working model for both toxicity and efficacy. Case (1) outperforms cases (2) and (3) in Scenarios 2 and 4 and is competitive in scenarios 3 and 5. However, caution must still be exercised as demonstrated by the performance of all three cases which recommends the true OBD in only a maximum of 59.0% of trials in Scenario 1 where the OBD is level 3 and the one-parameter logistic correctly identifies level 3 as the OBD in 71.1% of the trials. Therefore, it is recommended that any use of the two-parameter logistic

working model in the proposed method is met with strong prior belief that there exists a dose with a probability of toxicity close to, but above, the toxicity tolerance.

### 2.6.4 Utility Functions

The utility based approach is one which chooses the dose level with the maximum utility value as the dose to be given to the next patient. Let  $G(d_i, \theta) = \pi_E(d_i, \theta)$  be the model for efficacy and  $F(d_i, \beta) = \pi_T(d_i, \beta)$  be the model for toxicity. In Section 2.4, the dose decisions were determined according the following utility function:

$$\mathcal{U}_{\mathcal{A}} = \pi_E(d_{i*}, \theta) \text{ for } i* \in \mathcal{A}.$$

However, alternative utility functions could have been implemented to guide the interim dose decisions. Other utility functions explored include the ones mentioned in Cheung (2011) including the probability ratio, denoted

$$\mathcal{U} = \frac{\pi_E(d_i, \theta)}{\pi_T(d_i, \beta)};$$

odds ratio, denoted

$$\mathcal{U} = \frac{\pi_E(d_i, \theta)\{1 - \pi_T(d_i, \beta)\}}{\pi_T(d_i, \beta)\{1 - \pi_E(d_i, \theta)\}};$$

a success probability which assumes independence between toxicity and efficacy, denoted

$$\mathcal{U} = \pi_E(d_i, \theta)\{1 - \pi_T(d_i, \beta)\};$$

and an acceptable set success probability, denoted

$$\mathcal{U}_{\mathcal{A}} = \pi_E(d_{i*}, \theta)\{1 - \pi_T(d_{i*}, \beta)\} \text{ for } i* \in \mathcal{A}.$$

A practical issue that arises when using the utility-based approach is which utility function to use because different utilities can lead to different dose decisions. This discrepancy is due to the fact that low doses are preferred based on the ratio utilities since a low estimated toxicity probability inflates the ratio but higher doses are preferred in the success probability approach. This favorability of high doses by the success probability is the reason for introducing the acceptable set success probability utility. Therefore, much care must be taken to choose an appropriate utility function for a given trial (Cheung, 2011).

In order to explore the feasibility of the utility function, a simulation study is presented for a design which uses the success probability to allocate patients to doses as compared to the proposed method in Section 2.3 for various true toxicity and efficacy scenarios and design specifications. Using success probabilities to guide the interim dose decisions is done the same way as the proposed method but replacing  $\hat{\pi}_E$  with the estimated success probability throughout the design. Therefore, the randomization probability for dose  $i$  in the adaptive randomization phase becomes

$$R_i = \frac{\hat{\pi}_E(d_i) \times (1 - \hat{\pi}_T(d_i))}{\sum_{d_i} \hat{\pi}_E(d_i) \times (1 - \hat{\pi}_T(d_i))}.$$

The study was also carried out for the probability ratio and odds ratio but the results are excluded from the table because they perform poorly overall. The ratios perform poorly because the methods lock-in on a dose early-on in the trial and become “stuck” at the dose no matter how much of the sample size is used as the exploration sample size. This drawback is due to the fact that the ratios do not have the restriction of being between zero and one that is apparent in the success probability and the proposed approach. In the ratio setting, a dose level could have a ratio estimate of around ten and the rest be below one because a

drastically low estimated toxicity causes inflation of the ratio estimate which makes the randomization probability for that dose to be close to one, which in turn, causes the procedure to lock-in on a single dose that may be the incorrect dose.

A subset of the results are shown in Table B.8. Results are presented for both the success probability and the acceptable set success probability utilities. When no acceptable set is used in the success probability utility function to control toxicity, the method often recommends a too toxic dose as the OBD. This notion is especially apparent in Scenario 1 where the method recommends  $d_4$  as the OBD in 50.6% of the simulated trials which is particularly troubling since  $\pi_T(d_4) = 0.80$ . This aggressiveness of the success probability utility function is noted by Cheung (2011). Therefore, it is our recommendation that if the success probability function is applied to determine interim dose decisions, then the acceptable set success probability utility is used to control toxicity and prevent treating patients at potentially harmful doses. Notice that using the acceptable set success probability utility function is a much better option than when no acceptable set is exploited. When using the acceptable set success probability utility to guide interim dose decisions, the results are almost identical to the proposed method across the scenarios presented in Table B.8.

### 2.6.5 Model Selection Criteria

Instead of taking into account all the  $2I - 1$  models into consideration and possibly including a model which is far worse than all the other models, a form of Occam's window criterion can be applied to cause the method to potentially converge to the correct model quicker. In a crude approximation, simulations were performed where the lowest half of the models in terms of the posterior model probability were adaptively eliminated at two different stages: at  $n_R$  and after three-fourths of the

total patients have been enrolled into the trial. Another crude approximation which was studied was to adaptively eliminate a quarter of the models at  $n_R$  and another quarter of the models from consideration at three-fourths of the way through the trial. The simulation results are presented in Table B.9. Notice that the results diminish quite substantially from considering all the models throughout the entire duration of the trial. Out of the three crude adaptive elimination methods, eliminating half of the models at three-fourths of the way through the trial seems to work best. However, based on the simulation results, eliminating the orders in this way does not seem to be a suitable way of conducting a trial.

In order to take into account more uncertainty into the model, the idea of using Bayesian Model Averaging (BMA) was introduced in the context of the CRM by Yin and Yuan (2009). The main purpose of using BMA is to incorporate the uncertainty surrounding dose-finding by taking multiple skeletons into account at the same time. The estimate for the efficacy probability at each dose level using BMA is given by

$$\bar{\pi}_E(d_i) = \sum_{k=1}^{2I-1} \hat{\pi}_E(d_{ik}) \Pr(M_k | \mathcal{D}_j)$$

where  $\hat{\pi}_E(d_{ik})$  is the estimate of the efficacy probability at dose  $d_i$  under model  $M_k$ . The notion of adaptively eliminating models from consideration fits in well when using BMA in order to estimate the efficacy probability at each dose. Therefore, using Occam's window criterion, we will include model  $M_k$  for consideration at the enrollment of the  $j^{th}$  patient only if

$$\frac{\Pr(M_k | \mathcal{D}_j)}{\max_{i \in 1, \dots, 2I-1} \Pr(M_k | \mathcal{D}_j)} > \delta.$$

Since we make use of this criteria at the enrollment of each patient it is possible that a model that was previously excluded could be included. In the simulation study

presented in Table B.10, BMA using Occam’s window, denoted BMAO, used  $\delta$  equal to 0.6, as suggested in Yuan and Yin (2009) or 0.7 to exclude models deemed far worse than the best fitting model. The BMAO as a model selection criteria fared better than using simple BMA since it is excluding models far from the best model. This notion is due to the nature of the possible efficacy models. Since the models take into account where the peak or plateau occurs, including a model where the peak is at dose level 4 is detrimental to the BMA average when the true peak of the dose-efficacy curve occurs at dose level 1. Therefore, when BMAO excludes this model from consideration, the estimates of the probability of efficacy at each dose are improved. The downside of BMA in this context is seen in the simulation results in Table B.10 where in Scenarios 2 and 4, when the OBD is on the boundary of the dose levels, the PCS is only 73.6% and 76.5%, respectively, as compared to the PCS of 74.2% for Scenario 2 and 90.7% for Scenario 4 for the proposed method. On the other hand, BMAO performs much closer to the proposed method, especially in Scenario 4 with a PCS of 90.2% for  $\delta = 0.6$  and 89.7 % for  $\delta = 0.7$ . A key feature of the BMAO method is that the simulation results show it to be fairly robust against the choice of  $\delta$ , which makes sense given  $\delta$  being either 0.6 or 0.7. However, note that the closer  $\delta$  is set to 0, the closer the results for BMAO would be to the results of BMA. If  $\delta$  is set to 1, then the method will only utilize the model with the highest posterior model probability to obtain the estimates for the efficacy probabilities and will therefore be the same as the proposed method.

## 2.7 Terminating the Trial

In practice, investigators will want some measure by which to stop the trial in the presence of undesirable toxicity in order to protect the patients safety. The notion of terminating the trial early is especially important in oncology trials where a lengthy

Table 2.3: An example of the proposed method using the Bayesian stopping rule as presented in Yuan and Yin (2011) for a scenario presented in Thall and Cook (2004).

	True (toxicity, efficacy) probability				None	$\bar{n}$
	(0.50, 0.05)	(0.75, 0.25)	(0.85, 0.50)	(0.87, 0.70)		
$\phi_T = 0.40$	38.8	0.0	0.0	0.0	62.2	25.37
$\phi_T = 0.30$	8.8	0.0	0.0	0.0	91.2	17.11

Phase II study of the response may be required to determine any potential efficacy of the experimental treatment and the concern is to quickly move to the Phase II trial (O’Quigley 2002). Even though we do not expect undesirable toxicity in targeted agents, it is still of utmost importance ethically to protect the patients enrolled in the trial. Therefore, we still must ensure patients are not treated with potentially harmful doses.

Yuan and Yin (2011) proposed a Bayesian stopping rule that terminates the trial when

$$\Pr(\pi_T(d_1) > \phi_T | \Omega_j) > 0.9.$$

In other words, the trial is halted if the posterior probability that the lowest dose’s probability of toxicity is greater than  $\phi_T$  is larger than 0.9. The stopping rule proposed for this design will only be assessed in the maximization phase, not in the exploration phase. This allows the trial to gain enough information to more accurately determine that the trial should be stopped, preventing it from stopping too early. Consider the scenario from Thall and Cook (2004) presented in Table 2.3 where the lowest dose has a probability of toxicity of 0.50 which is larger than the toxicity tolerance  $\phi_T = 0.40$ . Similarly to Section 2.5.1, let  $p_1 = 0.05$ ,  $p_2 = 0.20$ ,  $p_3 = 0.35$ , and  $p_4 = 0.45$  as well as consider the  $2 \times 4 - 1 = 7$  possible models for efficacy as presented in Table 2.1. In Thall and Cook (2004),  $\phi_T = 0.40$  therefore we expect that the trial would terminate fairly often before reaching the exhaustive

sample size. Table 2.3 demonstrates that the proposed method with the stopping rule stops 62.2% of the time with an average sample size of 25.37 as compared to the total sample size of  $n = 36$ . Therefore, an average of roughly 11 fewer patients were exposed to the toxic treatment by imploring the use of the stopping rule proposed by Yuan and Yin (2009). Note that if the value that the posterior probability has to be larger than decreases say to 0.8 from 0.9, then the method would stop more often. Given how close  $\phi_T$  is to  $\pi_T(d_1)$ , the method performs well at stopping the trial early. However, if the toxicity tolerance is decreased to  $\phi_T = 0.30$ , then the proposed method stops 91.2% of the time with an average sample size of 17.11. Therefore, the stopping rule proposed by Yuan and Yin (2011) has good operating characteristics when paired with the method proposed in this dissertation.



# Chapter 3

## Theoretical Properties

In the previous chapter, a large simulation study was presented to demonstrate the behavior of the CRM for bivariate binary endpoints in identifying the OBD. The method is recommending the true OBD or a treatment close to the true OBD in a large number of simulated trials. Here, we provide a theoretical argument to prove the large sample properties of the proposed method.

### 3.1 Convergence

#### 3.1.1 Model Assumptions in CRM

The theoretical properties of the CRM are well documented in the literature. Shen and O’Quigley (1996) showed that under certain conditions detailed below the original CRM is consistent under model misspecification.

**Condition 1.** *The parameter  $\beta$  belongs to a finite interval  $[A, B]$ .*

This condition eliminates degenerate cases (when the probability of toxicity is equal to zero or one) when the value of the toxicity parameter,  $\beta$ , lies on the boundary of the parameter space.

The following are restrictions on the working model  $F(d, \beta)$  as defined in Section 1.3.3 and thus are verifiable.

**Condition 2.**  $F(d, \beta)$  is strictly increasing in  $d$  for all  $\beta$ .

**Condition 3.**  $F(d, \beta)$  is continuous and is strictly monotone in  $\beta$  in the same direction for all  $d$ .

Let  $g_{ij}(\beta) = \{1 - F(d_i, \beta)\} / \{1 - F(d_j, \beta)\}$ .

**Condition 4.** The following inequality holds  $|F'(d_i, \beta)| \geq |F'(d_j, \beta)|g_{ij}(\beta)$  for all  $i > j$ .

This last condition puts a lower bound on  $|F'(d_i, \beta)|$  in terms of  $|F'(d_j, \beta)|$ . The next two conditions ensure that the model is flexible enough to approximate various dose-toxicity scenarios with all possible locations of the maximum of the acceptable set  $\ell$  and any given toxicity tolerance  $\phi_T$ .

**Condition 5.** For any given  $\phi_T \in (0, 1)$  and a given  $\ell$ , there exists  $\beta$  in the parameter space such that  $F(d_\ell, \beta) = \phi_T$ .

**Condition 6.** Suppose  $F(d_i, \beta)$  is decreasing in  $\beta$ . There exist constants  $\underline{b}$  and  $\bar{b}$  such that  $F(d_1, \underline{b}) > \phi_T$  and  $F(d_I, \bar{b}) < \phi_T$  and  $[\underline{b}, \bar{b}]$  is a compact subset of the parameter space of  $\beta$ .

**Condition 7.**  $F(d_i, \beta)$  is bounded away from 0 and 1 on  $[\underline{b}, \bar{b}]$  for all  $i$ ; and  $F'(d_i, \beta)$  is uniformly bounded in  $\beta$ .

**Condition 8.** For any given  $0 < p < 1$  and each  $i$ , the function

$$p \frac{F'(d_i, \beta)}{F(d_i, \beta)} + (1 - p) \frac{-F'(d_i, \beta)}{1 - F(d_i, \beta)}$$

is continuous and strictly monotone in  $\beta$ .

The following restriction:

**Condition 9.** For  $\beta \in [\underline{b}, \bar{b}]$ ,

$$F(d_i, \beta)F''(d_i, \beta) - \{F'(d_i, \beta)\}^2 \leq 0$$

and

$$-F''(d_i, \beta) + [F(d_i, \beta)F''(d_i, \beta) - \{F'(d_i, \beta)\}^2] \leq 0$$

with at least one inequality being strict for all  $i$ .

is required since we employ the use of the Bayesian posterior mean as the estimator over the maximum likelihood estimator. This is a reasonable assumption as illustrated by the following example. Assume that the dose-toxicity function has the following form

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

Note that

$$F(d_i, \beta)F''(d_i, \beta) - \{F'(d_i, \beta)\}^2 = \exp^\beta d_i^{\exp(2\beta)} \log d_i < 0 \quad (3.1)$$

where

$$F'(d_i, \beta) = \exp^\beta d_i^{\exp(\beta)} \log d_i$$

and

$$F''(d_i, \beta) = \exp^\beta d_i^{\exp(\beta)} \log d_i [\exp^\beta \log d_i + 1].$$

Furthermore, note that

$$\begin{aligned} -F''(d_i, \beta) + [F(d_i, \beta)F''(d_i, \beta) - \{F'(d_i, \beta)\}^2] = \\ \exp^\beta d_i^{\exp(\beta)} \log d_i \left[ d_i^{\exp(\beta)} - \exp^\beta \log d_i - 1 \right] < 0 \end{aligned} \quad (3.2)$$

The inequalities in Equations 3.1 and 3.2 are strict due to Condition 7.

Cheung and Chappell (2002) introduced a simple way to evaluate the sensitivity of the model in CRM. The authors introduced the notion of the home set at each dose level defined as

$$H_i = \{\beta \in \mathcal{B} : |\pi_T(d_i) - \phi_T| < |\pi_T(d_k) - \phi_T| \text{ for } k \neq i\}.$$

Then, they defined the least false parameter,  $\mu_i$ , at each dose level to be  $\mu_i = F^{-1}(\pi_T(d_i))$ . Then, the CRM is consistent if the following condition holds:

**Condition 10.**

$$\mu_\ell \in H_\ell$$

$$\mu_i \in \cup_{j=i+1}^I H_j \text{ for } i = 1, \dots, \ell - 1$$

$$\mu_i \in \cup_{j=1}^{i-1} H_j \text{ for } i = \ell + 1, \dots, I$$

where  $\ell$  is the location of the true maximum of the acceptable set.

Note that the above condition is a less-restrictive condition required for consistency than the original condition proposed by Shen and O'Quigley (1996) which required

$$\mu_i \in H_\ell.$$

### 3.1.2 Bivariate Extension of Model Assumptions

An extension to Condition 5 can be derived in the bivariate case to include both toxicity and efficacy into the following flexibility condition:

**Condition 11.** *For any given  $\phi_T, \phi_E \in (0, 1)$  and a given  $x$ , there exists a  $\beta$  and  $\theta$  in the parameter space such that  $\pi_E(x, \theta) = \phi_E$  and  $\pi_T(x, \beta) = \phi_T$ .*

Given the standard conditions of the CRM on the dose-toxicity function discussed in the previous section and an extension of the richness condition to the one from Cheung (2011) (Condition 5), then Theorem 1 will hold. The conditions on  $G_k(x, \theta_k)$  which are required in order to extend the theory developed in O’Quigley (2006) and outlined in Wages, Conaway, and O’Quigley (2011) are as follows.

**Condition 12.** *The parameter  $\theta_k$  belongs to a finite interval  $[A, B]$ .*

This condition protects against degenerate cases when the probability of efficacy equals 0 or 1 when  $\theta_k$  lies on the boundary of the parameter space.

**Condition 13.** *For a fixed  $\theta_k$ ,  $G_k(x, \theta_k)$  is continuous and strictly increasing in  $x$ .*

**Condition 14.** *For a fixed  $x$ ,  $G_k(x, \theta_k)$  is continuous and strictly decreasing in  $\theta_k$ .*

**Condition 15.** *There exist constants  $\theta_{k1}, \dots, \theta_{kI} \in [A, B]$  such that for  $1 \leq i \leq I$*

$$(i) \ G_k(d_i, \theta_{ki}) = \pi_E(d_i)$$

$$(ii) \ G_k(d_i, B) < \theta < G_k(d_i, A)$$

$$(iii) \ \text{For a unique } \theta_{k0} \in (\theta_{k1}, \dots, \theta_{kI}), \ G_k(d_0, \theta_{k0}) = \pi_E(d_0) = \theta_0.$$

These conditions restrict our choices of the working model  $G_k(x, \theta_k)$ . It will also be required that

**Condition 16.** *For each  $0 < t < 1$  and each  $x$  the function*

$$s(t, x, q) := t \frac{G'_k}{G_k}(x, \theta_k) + (1 - t) \frac{-G'_k}{1 - G_k}(x, \theta_k)$$

*is continuous and is strictly monotone in  $\theta_k$ .*

This is a standard requirement in order for the estimating equations to have unique solutions.

Since we are using the Bayesian posterior mean as the estimate of  $\theta$ , we require the additional condition that

**Condition 17.** For  $\theta_k \in [\underline{b}, \bar{b}]$ ,

$$G_k(d_i, \theta)G_k''(d_i, \theta) - \{G_k'(d_i, \theta)\}^2 \leq 0$$

and

$$-G_k''(d_i, \theta) + [G_k(d_i, \theta)G_k''(d_i, \theta) - \{G_k'(d_i, \theta)\}^2] \leq 0$$

with at least one inequality being strict for all  $i$ .

The proofs presented in this dissertation rely at first on maximum likelihood theory. Therefore, heterogeneity must exist in the toxic responses. Heterogeneity is defined as having at least one toxic and at least one non-toxic response. Once the maximum likelihood results are shown, the results are extended to the model's Bayesian framework. Since maximum likelihood theory is utilized, we must first prove that the maximum likelihood estimate exists after a certain point.

**Lemma 1.** *If the true probabilities of toxicity are monotone increasing,  $0 < \pi_T(d_1) < \dots < \pi_T(d_I) < 1$ , then the  $\Pr(n_0 < \infty) = 1$ , where  $n_0 = \inf\{k : 0 < \sum_{j=1}^k y_j < k\}$  and is the point at which heterogeneity exists.*

**Proof:** Let  $T_k = \sum_{j=1}^k [y_j - \pi_T(x_j)]$ . The continual reassessment method's definition implies that  $T_k$  is a martingale since each term in the summation is bounded,  $-1 \leq y_j - \pi_T(x_j) \leq 1$ . Therefore, we can apply the limit theorem for martingales presented in Shiryaev (1984) §7.3, Corollary 2 to show that  $k^{-1}T_k \xrightarrow{a.s.} 0$ . The monotonicity assumption leads to

$$0 < \pi_T(d_1) \leq \frac{1}{k} \sum_{j=1}^k \pi_T(x_j) \leq \pi_T(d_I) < 1$$

which demonstrates that  $1 \leq \sum_{j=1}^k y_j \leq k - 1$  as  $k$  tends to infinity. Therefore,  $\Pr(n_0 < \infty) = 1$ . □

We require a slight tweak on the definition of the home set to take into account the different goal of the trial we are performing. Therefore, the home set becomes

$$H_i = \{\beta \in \mathcal{B} : \Delta(F(d_i, \beta), \phi_T) < \Delta(F(d_k, \beta), \phi_T) \text{ for } k \neq i\}$$

where

$$\Delta(F(d_i, \beta), \phi_T) = \begin{cases} \delta [\phi_T - F(d_i, \beta)] & \text{if } F(d_i, \beta) \leq \phi_T \\ (1 - \delta) [F(d_i, \beta) - \phi_T] & \text{if } F(d_i, \beta) > \phi_T. \end{cases}$$

**Lemma 2.** *If Conditions 2, 3, 8, the monotonicity assumption from Lemma 1, and the target dose level is  $x_\ell$ , meaning that  $\pi_T(x_\ell) \leq \phi_T$  and  $\pi_T(x_{\ell+1}) > \phi_T$ , then  $H_\ell$  is an open and convex set.*

**Proof:** The continuity of  $F(d, \beta)$  guarantees the openness of the set  $H_i$ . In order to establish convexity, let  $\beta', \beta'' \in H_i$  such that  $\beta' < \beta_0 < \beta''$ . Without loss of generality, we can focus on a particular dose level  $x_i < x_\ell$ , Then by the definition of  $H_i$ :

$$\Delta(F(d_\ell, \beta'), \phi_T) < \Delta(F(d_i, \beta'), \phi_T) \quad (3.3)$$

$$\Delta(F(d_\ell, \beta''), \phi_T) < \Delta(F(d_i, \beta''), \phi_T) \quad (3.4)$$

It follows that  $\delta[\phi_T - F(d_\ell, \beta')] < \delta[\phi_T - F(d_i, \beta')]$  since  $F(d, \beta')$  is increasing due to Condition 2. Note that (3.3) implies that  $\delta[\phi_T - F(d_i, \beta')] > 0$ . Then we have

$$-\delta[\phi_T - F(d_i, \beta')] < |\delta[\phi_T - F(d_\ell, \beta')]| < \delta[\phi_T - F(d_i, \beta')].$$

Likewise, (3.4) leads to

$$-\delta[\phi_T - F(d_i, \beta'')] < |\delta[\phi_T - F(d_\ell, \beta'')]| < \delta[\phi_T - F(d_i, \beta'')]. \quad (3.5)$$

Now suppose that  $\beta \in (\beta_0, \beta'')$ , then Condition 3 implies that  $\delta[\phi_T - F(d_i, \beta)] > 0$  because

$$\delta[\phi_T - F(d_i, \beta_0)] > \delta[\phi_T - F(d_\ell, \beta_0)] = 0, \quad \delta[\phi_T - F(d_i, \beta'')] > 0.$$

Note that  $F(d_\ell, \beta) - \phi_T < 0$ , then  $F(d_\ell, \beta)$  must be decreasing in  $\beta$  because we know that  $F(d_\ell, \beta_0) - \phi_T = 0$ . Likewise,  $F(d_i, \beta)$  is decreasing by Condition 3. Therefore,

$$0 < \delta[\phi_T - F(d_\ell, \beta)] < \delta[\phi_T - F(d_\ell, \beta'')] < \delta[\phi_T - F(d_i, \beta'')] < \delta[\phi_T - F(d_i, \beta)]$$

Therefore,

$$\Delta(F(d_\ell, \beta), \phi_T) < \Delta(F(d_i, \beta), \phi_T) \quad (3.6)$$

and thus convexity is established. The argument for  $\beta \in (\beta', \beta_0)$  is similar.  $\square$

**Theorem 1.** *If Conditions 1, 2, 3, 5, 8, 9, 10, and the monotonicity assumption, hold, then as  $n$  tends to infinity  $\Pr(d_{n+1} = d_\ell) \rightarrow 1$  and almost surely  $\hat{\beta}_n^{Bayes} \rightarrow \beta$ , where  $\ell = \arg \min_i \Delta(\pi_T(d_i), \phi_T)$  and*

$$\Delta(\hat{\pi}(d_i), \phi_T) = \begin{cases} \delta [\phi_T - \hat{\pi}(d_i)] & \text{if } \hat{\pi}(d_i) \leq \phi_T \\ (1 - \delta) [\hat{\pi}(d_i) - \phi_T] & \text{if } \hat{\pi}(d_i) > \phi_T \end{cases}$$

with  $\delta \rightarrow 0$ .

Before the theorem for consistency is proven, we first note that the loss function in Theorem 1 is continuous since each part of the piecewise function is itself continuous on its given domain and  $\Delta(\phi_T, \phi_T) = \delta[\phi_T - \phi_T] = 0$ , so  $\Delta(\phi_T, \phi_T)$  exists,  $\lim_{\hat{\pi}(d_i) \rightarrow \phi_T^+} \Delta(\phi_T, \phi_T) = \lim_{\hat{\pi}(d_i) \rightarrow \phi_T^+} (1 - \delta)[\hat{\pi}(d_i) - \phi_T] = 0$ , and  $\lim_{\hat{\pi}(d_i) \rightarrow \phi_T^-} \Delta(\phi_T, \phi_T) = \lim_{\hat{\pi}(d_i) \rightarrow \phi_T^-} \delta[\phi_T - \hat{\pi}(d_i)] = 0$ . Therefore,  $\Delta(\hat{\pi}(d_i), \phi_T) = \Delta(\hat{\pi}(d_i), \phi_T)$  is continuous at



$\phi_T$  and thus continuous on  $[0, 1]$  since  $\delta[\phi_T - \hat{\pi}(d_i)]$  is continuous on  $[0, \phi_T)$  and  $(1 - \delta)[\hat{\pi}(d_i) - \phi_T]$  is continuous on  $(\phi_T, 1]$ .

**Proof:** The proof follows a similar idea as presented by Shen and O'Quigley (1996). Suppose the first  $n$  patients were treated at  $x_1, x_2, \dots, x_n$  with responses  $y_1, y_2, \dots, y_n$ . Define

$$I_n(\beta) = \frac{1}{n} \sum_{j=1}^n \left[ y_j \frac{F'(x_j, \beta)}{F(x_j, \beta)} + (1 - y_j) \frac{-F'(x_j, \beta)}{1 - F(x_j, \beta)} \right]$$

and

$$\tilde{I}_n(\beta) = \frac{1}{n} \sum_{j=1}^n \left[ \pi_T(x_j) \frac{F'(x_j, \beta)}{F(x_j, \beta)} + (1 - \pi_T(x_j)) \frac{-F'(x_j, \beta)}{1 - F(x_j, \beta)} \right].$$

Note that for each dose level  $d_i$ ,  $\frac{F'(x_j, \beta)}{F(x_j, \beta)}$  and  $\frac{F'(x_j, \beta)}{1 - F(x_j, \beta)}$  are uniformly continuous in  $\beta$  over the finite interval  $[A, B]$ . Then for any  $\varepsilon > 0$  and each  $d_i$ , there must exist a partition  $A = t_0 < t_1 < \dots < t_k = B$  such that  $\forall \beta \in [A, B]$ , there exists a  $t_{k_0}$  such that

$$\left| \frac{F'(x_j, \beta)}{F(x_j, \beta)} - \frac{F'(x_j, t_{k_0})}{F(x_j, t_{k_0})} \right| < \varepsilon \quad (*)$$

$$\left| \frac{F'(x_j, \beta)}{1 - F(x_j, \beta)} - \frac{F'(x_j, t_{k_0})}{1 - F(x_j, t_{k_0})} \right| < \varepsilon \quad (**)$$

Since there are only  $I$  possible dose levels, the partition may be chosen so that  $(*)$  and  $(**)$  are valid for all  $d_i$ .

Let  $I_n(\beta) - \tilde{I}_n(\beta) = I_{n1}(\beta) + I_{n2}(\beta) + I_{n3}(\beta)$  where

$$I_{n1}(\beta) = I_n(\beta) - \frac{1}{n} \sum_{j=1}^n \left[ y_j \frac{F'(x_j, t_{k_0})}{F(x_j, t_{k_0})} + (1 - y_j) \frac{-F'(x_j, t_{k_0})}{1 - F(x_j, t_{k_0})} \right]$$

$$I_{n2}(\beta) = \frac{1}{n} \sum_{j=1}^n \left[ (y_j - \pi_T(x_j)) \frac{F'(x_j, t_{k_0})}{F(x_j, t_{k_0})} + \frac{F'(x_j, t_{k_0})}{1 - F(x_j, t_{k_0})} \right]$$

and

$$I_{n3}(\beta) = \frac{1}{n} \sum_{j=1}^n \left( \pi_T(x_j) \left[ \frac{F'(x_j, t_{k_0})}{F(x_j, t_{k_0})} - \frac{F'(x_j, \beta)}{F(x_j, \beta)} \right] + \right. \\ \left. (1 - \pi_T(x_j)) \left[ \frac{-F'(x_j, t_{k_0})}{1 - F(x_j, t_{k_0})} - \frac{-F'(x_j, \beta)}{1 - F(x_j, \beta)} \right] \right).$$

It follows from (\*) and (\*\*) that

$$\sup_{\beta \in [A, B]} |I_{n1}(\beta)| \leq \frac{1}{n} \sum_{j=1}^n \left[ y_j \left( \frac{F'(x_j, \beta)}{F(x_j, \beta)} - \frac{F'(x_j, t_{k_0})}{F(x_j, t_{k_0})} \right) \right. \\ \left. + (1 - y_j) \left( \frac{F'(x_j, \beta)}{F(x_j, \beta)} - \frac{F'(x_j, t_{k_0})}{F(x_j, t_{k_0})} \right) \right] \\ = \frac{1}{n} \sum_{j=1}^n [y_j \varepsilon + (1 - y_j) \varepsilon] = \varepsilon$$

Similarly,

$$\sup_{\beta \in [A, B]} |I_{n3}(\beta)| \leq \frac{1}{n} \sum_{j=1}^n [\pi_T(x_j) \varepsilon + (1 - \pi_T(x_j)) \varepsilon] = \varepsilon$$

Now, it just remains to establish  $I_{n2} \xrightarrow{a.s.} 0$  for each  $t_{k_0}$ .

For a fixed  $t_{k_0}$ ,  $\{nI_{n2} : n \geq 1\}$  forms a martingale and each term in the summation in  $I_{n2}$  is bounded. Therefore,  $I_{n2} \xrightarrow{a.s.} 0$  by the Martingale Limit Theorem. Since there is a finite number of points in the partition

$$\sup_{\beta \in [A, B]} |I_n(\beta) - \tilde{I}_n(\beta)| \rightarrow 0.$$

Let  $S$  denote the set  $[\mu_{(1)}, \mu_{(I)}]$  where  $\mu_{(1)} = \min\{\mu_1, \dots, \mu_I\}$  and  $\mu_{(I)} = \max\{\mu_1, \dots, \mu_I\}$ , where  $\mu_i$  is the least false parameter defined in Condition 10. The monotonicity assumption and Lemma 2 imply that  $S \subset H_\ell$ . Then we rewrite

$$\tilde{I}_n(\beta) = \sum_{i=1}^I \hat{p}_i \left[ \pi_T(d_i) \frac{F'(d_i, \beta)}{F(d_i, \beta)} + (1 - \pi_T(d_i)) \frac{-F'(d_i, \beta)}{1 - F(d_i, \beta)} \right]$$

where  $\hat{p}_i$  is the frequency that level  $d_i$  has been administered to a patient. Suppose  $\tilde{\beta}_n$  is the solution to the equation

$$\sum_{i=1}^I \hat{p}_i \left[ \pi_T(d_i) \frac{F'(d_i, \beta)}{F(d_i, \beta)} + (1 - \pi_T(d_i)) \frac{-F(d_i, \beta)}{1 - F(d_i, \beta)} \right] = 0.$$

For each  $1 \leq i \leq I$ , the definition of  $\mu_i$  and Condition 8 imply  $\mu_i$  is the unique solution of  $\left[ \pi_T(d_i) \frac{F'(d_i, \beta)}{F(d_i, \beta)} + (1 - \pi_T(d_i)) \frac{-F(d_i, \beta)}{1 - F(d_i, \beta)} \right] = 0$ . Condition 3 implies that  $\tilde{\beta}$  must fall into the interval  $S$ . Since  $\hat{\beta}^{MLE}$  solves  $I_n(\beta) = 0$ ,

$$\sup_{\beta \in [A, B]} |I_n(\beta) - \tilde{I}_n(\beta)| \rightarrow 0$$

ensures that  $\hat{\beta}^{MLE} \in H_\ell$  a.s. eventually. Therefore,  $\hat{\beta}_n^{MLE}$  satisfies

$$\Delta(\hat{\pi}_T(d_\ell, \hat{\beta}_n^{MLE}), \phi_T) < \Delta(\hat{\pi}_T(d_i, \hat{\beta}_n^{MLE}), \phi_T)$$

Thus, as  $n$  tends to infinity  $d_{n+1} = d_\ell$ .

In order to obtain the consistency of  $\hat{\beta}_n^{MLE}$ , we observe that as  $n$  tends to infinity all the  $\hat{p}_i$ 's in  $\tilde{I}_n(\beta)$  become negligible, except the one corresponding to the level  $x_\ell$ , which tends to 1. Therefore,  $\tilde{\beta}_n$  is close to the solution of

$$\pi_T(d_\ell) \frac{F'(d_\ell, \beta)}{F(d_\ell, \beta)} + (1 - \pi_T(d_\ell)) \frac{-F'(d_\ell, \beta)}{1 - F(d_\ell, \beta)} = 0.$$

The solution to the above equation is inherently  $\beta$ . Applying the fact that

$$\sup_{\beta \in [A, B]} |I_n(\beta) - \tilde{I}_n(\beta)| \rightarrow 0,$$

shows that  $\hat{\beta}_n^{MLE} \xrightarrow{a.s.} \beta$ .

Now that we have proven the MLE,  $\hat{\beta}_n^{MLE}$ , is consistent, we now need to show

it holds for the posterior mean,  $\hat{\beta}_n^{Bayes}$ . Let the  $l_n(\beta) = \log L_n(\beta)$  with first and second derivatives:

$$l'_n(\beta) = \sum_{i=1}^I c_i \frac{F'(d_i, \beta)}{F(d_i, \beta)} + (n_i - c_i) \frac{-F'(d_i, \beta)}{1 - F(d_i, \beta)}$$

and

$$l''_n(\beta) = \sum_{i=1}^I \left[ c_i \frac{F(d_i, \beta)F''(d_i, \beta) - [F'(d_i, \beta)]^2}{[F(d_i, \beta)]^2} \right] + \\ (n_i - c_i) \left[ \frac{-F''(d_i, \beta) + F(d_i, \beta)F''(d_i, \beta) - [F'(d_i, \beta)]^2}{[1 - F(d_i, \beta)]^2} \right]$$

where  $n_i$  and  $c_i$  are the number of patients treated and the number of toxic responses at level  $i$ . Taking a Taylor's series expansion, leads to

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

for some  $\tilde{\beta}_n^*$  between  $\beta$  and  $\hat{\beta}_n^{MLE}$  because  $l'_n(\hat{\beta}_n^{MLE}) = 0$ .

Applying Condition 9, we have  $\bar{l}''_n \equiv \sup_{\beta \in [\underline{b}, \bar{b}]} l''_n(\beta) \rightarrow -\infty$  as  $n \rightarrow \infty$  almost surely.

Note, the posterior mean can be expressed as a function of the MLE through

$$\hat{\beta}_n^{Bayes} = \frac{\int_{-\infty}^{\infty} \beta L_n(\beta) g(\beta) d\beta}{\int_{-\infty}^{\infty} L_n(\beta) g(\beta) d\beta} = \hat{\beta}_n^{MLE} + \frac{\int_{-\infty}^{\infty} (\beta - \hat{\beta}_n^{MLE}) L_n(\beta) g(\beta) d\beta}{\int_{-\infty}^{\infty} L_n(\beta) g(\beta) d\beta}.$$

Notice,

$$\left| \int_{-\infty}^{\infty} (\beta - \hat{\beta}_n^{MLE}) L_n(\beta) g(\beta) d\beta \right| = \\ \left| \int_{-\infty}^{\infty} (\beta - \hat{\beta}_n^{MLE}) \exp \left[ l_n(\hat{\beta}_n^{MLE}) + \frac{1}{2}(\beta - \hat{\beta}_n^{MLE})^2 l''_n(\tilde{\beta}_n^*) \right] g(\beta) d\beta \right| \\ \leq L_n(\hat{\beta}_n^{MLE}) \int_{-\infty}^{\infty} \left| \beta - \hat{\beta}_n^{MLE} \right| \exp \left[ \frac{1}{2}(\beta - \hat{\beta}_n^{MLE})^2 \bar{l}''_n \right] g(\beta) d\beta$$

We can break the integral into two pieces:

$$\int_{\hat{\beta}_n^{MLE-\varepsilon}}^{\hat{\beta}_n^{MLE+\varepsilon}} \left| \beta - \hat{\beta}_n^{MLE} \right| \exp \left[ \frac{1}{2} (\beta - \hat{\beta}_n^{MLE})^2 \bar{l}_n'' \right] g(\beta) d\beta \leq \varepsilon$$

and

$$\left( \int_{-\infty}^{\hat{\beta}_n^{MLE-\varepsilon}} + \int_{\hat{\beta}_n^{MLE+\varepsilon}}^{\infty} \right) \left| \beta - \hat{\beta}_n^{MLE} \right| \exp \left[ \frac{1}{2} (\beta - \hat{\beta}_n^{MLE})^2 \bar{l}_n'' \right] g(\beta) d\beta \rightarrow 0$$

as  $n \rightarrow \infty$ . Since  $\varepsilon$  can be chosen to be arbitrarily small, we have shown that

$$\left| \hat{\beta}_n^{Bayes} - \hat{\beta}_n^{MLE} \right| \xrightarrow{a.s.} 0.$$

We now know that  $\hat{\beta}_n^{MLE} \in H_\ell$  eventually with probability one and since  $H_\ell$  is an open set we can make the same conclusion about  $\hat{\beta}_n^{Bayes}$ . We have thus demonstrated that  $\hat{\beta}_n^{Bayes} \xrightarrow{a.s.} \beta_0$ .  $\square$

As discussed in Bickel and Yahav (1966) and Doob (1949), the asymptotic distribution of  $\hat{\beta}_n^{Bayes}$  is equivalent to the asymptotic distribution of  $\hat{\beta}_n^{MLE}$  since the prior is diminished as  $n$  approaches infinity. Since the likelihood for the proposed method is identical to the likelihood in the original CRM, Theorem 2 from Shen and O'Quigley (1996) directly applies as long as the same conditions of Theorem 1 presented in this dissertation are satisfied. Therefore, the asymptotic distribution of  $\sqrt{n}(\hat{\beta}_n^{Bayes} - \beta_0)$  is  $N(0, \{F'(d_\ell, \beta_0)\}^{-2} \phi_T(1 - \phi_T))$ .

**Remark 1.** *The proof of Theorem 1 does not depend on the assumption that  $\pi_T(d_\ell) = \phi_T$  but the method converges to  $d_\ell$  as long as  $\pi_T(d_\ell)$  is the closest to  $\phi_T$  out of all the possible dose levels according to the function  $\Delta(\pi_T(d_i), \phi_T)$ .*

**Remark 2.** *The previous theorem provides reassurance in the use of the proposed method in scenarios such as Scenario 5 in Table B.6 where the proposed method does*

not recommend the true OBD in a large percentage of the trials in the short run. Therefore, it suggests that if  $n$  is increased from the small sample size of just 36 patients, then we would expect the proposed method to converge to the correct OBD,  $d_1$ . This fact is apparent when we increase the sample size to 150 patients and the PCS for  $d_1$  increases to from 46.5% to 70.5% of trials.

For the following theorem, we restrict the possible models to be unimodal or plateau. Note that with a single targeted agent multiple modal dose-response curves such as  $d_1 \leq d_2 \geq d_3 \leq d_4$  are not considered possible. Therefore, this assumption does not hinder the proposed method in any way.

**Theorem 2.** *If the prior is proper and not degenerate, the prior model probabilities are all equal, and Condition 17 holds, then as  $n \rightarrow \infty$ ,*

$$\lim_{n \rightarrow \infty} \Pr(BF_{mk^*} \leq BF_{mk}) = 1 \quad \forall k \neq k^*$$

where

$$BF_{mk^*} = \frac{\int_{\boldsymbol{\theta}} L_m(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}}{\int_{\boldsymbol{\theta}} L_{k^*}(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}}$$

and  $k^*$  is the model which minimizes the Kullback-Leibler distance.

**Proof:** Note  $\mathcal{D}_n = \{(x_1, z_1), \dots, (x_{n-1}, z_{n-1})\}$  and the prior  $g(\boldsymbol{\theta}) \sim N(\mathbf{0}, \mathbf{1.34})$  is proper and non-degenerate. Model selection is carried out using the notion of maximum posterior model probability, i.e. the method chooses the model,  $k^*$ , with the largest

$$\Pr(M_{k^*}|\mathcal{D}_n) = \frac{\tau(k^*) \int_{\boldsymbol{\theta}} L_{k^*}(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}}{\sum_{k=1}^{2I-1} \tau(k) \int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}}$$

Guo and Speckman (2009) noted that

$$\Pr(M_{k^*}|\mathcal{D}_n) \xrightarrow{p} 1 \iff \text{BF}_{mk^*} \xrightarrow{p} 0$$

since in order for the posterior model probability to converge to 1 in probability, the likelihood for model  $k^*$  must dominate the likelihoods of all other models being considered. This fact means that the denominator in  $\text{BF}_{mk^*}$  will diverge to infinity causing the Bayes Factor to converge in probability to 0.

Wasserman (2000) noted that two practical problems exist when using the Bayes Factor for model selection: the integration and the choice of the prior. However, Kass and Wasserman (1995) demonstrated that the issues diminish asymptotically, since through the Laplace transformation

$$\log \left( \int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} \right) \approx \log(L_k(\hat{\boldsymbol{\theta}}|\mathcal{D}_n)) - \frac{K \times \log n}{2} \quad (3.7)$$

where  $\hat{\boldsymbol{\theta}}$  is the maximum likelihood estimate and  $K$  is the number of parameters in  $\boldsymbol{\theta}$ , either 1 or 2. This is attained through defining

$$h(\boldsymbol{\theta}) = \log(L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta}))$$

so that

$$\int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} = \int e^{h(\boldsymbol{\theta})}d\boldsymbol{\theta}.$$

Then we apply a Taylor series expansion of  $h(\boldsymbol{\theta})$  around the posterior mode,  $\tilde{\boldsymbol{\theta}}$

$$h(\boldsymbol{\theta}) \approx h(\tilde{\boldsymbol{\theta}}) + \frac{1}{2}(\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}})'H(\tilde{\boldsymbol{\theta}})(\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}) \approx h(\tilde{\boldsymbol{\theta}}) - \frac{1}{2}(\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}})'(I_{\boldsymbol{\theta}})^{-1}(\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}})$$

where  $H = \frac{\partial^2 H}{\partial \theta_i \partial \theta_j}$  and  $I_{\boldsymbol{\theta}}$  is Fisher's Information matrix. When we combine the

Taylor series expansion and the Laplace approximation we obtain:

$$\begin{aligned} \int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} &\approx \exp(h(\tilde{\boldsymbol{\theta}})) \int_{\boldsymbol{\theta}} \exp\left[\frac{-1}{2}(\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}})'(I_{\boldsymbol{\theta}})^{-1}(\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}})d\boldsymbol{\theta}\right] \\ &= \exp(h(\tilde{\boldsymbol{\theta}}))(2\pi)^{K/2}|I_{\boldsymbol{\theta}}|^{1/2} \end{aligned}$$

Therefore,

$$\begin{aligned} \log\left(\int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}\right) &= \log(L_k(\tilde{\boldsymbol{\theta}}|\mathcal{D}_n)) + \log(g(\tilde{\boldsymbol{\theta}})) + \frac{1}{2}\log|I_{\boldsymbol{\theta}}| + O(1) \\ &= \log(L_k(\hat{\boldsymbol{\theta}}|\mathcal{D}_n)) + \log(g(\hat{\boldsymbol{\theta}})) - \frac{K}{2}\log n + O(1) \end{aligned}$$

$$\begin{aligned} \Pr(\text{BF}_{mk^*} - \text{BF}_{mk} \leq 0) &= \Pr\left(\frac{\int_{\boldsymbol{\theta}} L_m(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}}{\int_{\boldsymbol{\theta}} L_{k^*}(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}} - \frac{\int_{\boldsymbol{\theta}} L_m(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}}{\int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}} \leq 0\right) \\ &= \Pr\left(\frac{\int_{\boldsymbol{\theta}} L_m(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} \left[\int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} - \int_{\boldsymbol{\theta}} L_{k^*}(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}\right]}{\int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} \int_{\boldsymbol{\theta}} L_{k^*}(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta}} \leq 0\right) \end{aligned}$$

Now, we note that by definition of  $k^*$ :

$$\int_{\boldsymbol{\theta}} L_{k^*}(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} > \int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} \quad \forall k \neq k^*$$

is always true. Therefore,

$$\Pr(\text{BF}_{mk^*} - \text{BF}_{mk} \leq 0) \xrightarrow{P} 1$$

It now just remains to prove that  $k^*$  is the model that is closest to the underlying truth. Mathematically, this means that the model which is closest to the true



underlying model is model  $k^*$  which has the smallest Kullback-Leibler (KL) distance, defined as

$$I(T, k) = E_T \left[ \log \left( \int L_T(\boldsymbol{\theta} | \mathcal{D}_n) g(\boldsymbol{\theta}) d\boldsymbol{\theta} \right) \right] - E_T \left[ \log \left( \int L_k(\boldsymbol{\theta} | \mathcal{D}_n) g(\boldsymbol{\theta}) d\boldsymbol{\theta} \right) \right].$$

Since the term  $E_T \left[ \log \left( \int L_T(\boldsymbol{\theta} | \mathcal{D}_n) g(\boldsymbol{\theta}) d\boldsymbol{\theta} \right) \right]$  only depends on the truth it can be thought as some arbitrary constant  $C$ . Therefore, we cannot estimate how close model  $k$  is to the truth,  $T$ , based on  $I(T, k)$ . However, more importantly, we can identify which model out of the  $2I - 1$  candidate models is closest to the true underlying model by choosing the model such that

$$\arg \max_k E_T \left[ \log \left( \int L_k(\boldsymbol{\theta} | \mathcal{D}_n) g(\boldsymbol{\theta}) d\boldsymbol{\theta} \right) \right].$$

Note, that this is the *relative* KL distance as described by Bozdogan (1987), Kapur and Kesavan (1992), and Burnham and Anderson (2002) in relation to the absolute KL distance defined as

$$\hat{I}(T, k) - C = -\hat{E}_T \left[ \log \left( \int L_k(\boldsymbol{\theta} | \mathcal{D}_n) g(\boldsymbol{\theta}) d\boldsymbol{\theta} \right) \right].$$

Since we have shown in Equation 3.7 that the individual components of the Bayes Factor do not depend on the integration or the choice of the prior, we can proceed with using Akaike's (1973) result where he demonstrated that

$$-\log \left( L_k(\hat{\boldsymbol{\theta}} | \mathcal{D}_n) \right) + K = \hat{E}_{\hat{\boldsymbol{\theta}}} [\hat{I}(T, k)] - C$$

which demonstrates that the expected KL distance is approximately equal to the negative log likelihood plus the number of parameters in the model, where  $K = 1$  or  $2$ .

It has already been noted that

$$\int_{\boldsymbol{\theta}} L_{k^*}(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} > \int_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}|\mathcal{D}_n)g(\boldsymbol{\theta})d\boldsymbol{\theta} \quad \forall k \neq k^*$$

Therefore, model  $k^*$  minimizes the expected KL distance.  $\square$

**Corollary 1.** *The method consistently identifies the OBD.*

**Remark 3.** *Since Theorem 2 proves that the method consistently chooses the model which is closest to the truth in terms of the KL distance, the proposed method chooses the dose with the highest estimated efficacy within the acceptable set, and the dose-response curve is assumed to be unimodal or plateau, we know that the proposed method will correctly identify the OBD as  $n \rightarrow \infty$ .*

**Corollary 2.** *The consistency for the parameter  $\theta$  is directly tied in with consistently choosing the correct model. The consistency of Bayesian estimates is well documented in the literature going back to Doob (1949).*

## 3.2 Coherence of CRM

Coherence is defined as the important ethical principle where dose escalation only occurs when there is no toxicity in the previous patient. This limits exposing patients to overly toxic doses unnecessarily. Similarly, de-escalating in the face of a nontoxic outcome for the previous patient is counterintuitive and is also an incoherent move. Therefore, a dose-finding design  $\mathcal{D}$  is coherent if with probability one

$$\Pr_{\mathcal{D}}(x_j - x_{j-1} > 0 | Y_{j-1} = 1) = 0 \text{ for all } i$$

Table 3.1: Incoherent moves for Cytotoxic Agents with Bivariate Binary Outcomes

	Toxic Outcome	Nontoxic Outcome
Response	Incoherent escalation: toxicity	Incoherent de-escalation: toxicity
Nonresponse	Incoherent escalation: toxicity Incoherent de-escalation: efficacy	Incoherent de-escalation: toxicity, efficacy

and if with probability one

$$\Pr_{\mathcal{D}}(x_j - x_{j-1} < 0 | Y_{j-1} = 0) = 0 \text{ for all } i.$$

Therefore, a dose-finding design is one which does not involve any incoherent moves (Cheung, 2011).

### 3.2.1 Coherence of Bivariate Endpoints

The notion of extending the idea of coherence to the bivariate binary case is an important one. As noted in the previous section, de-escalating in the face of a nontoxic response is an incoherent move; as is escalating when faced with toxicity in the previous patient. Table 3.1 shows the possible incoherent moves for a bivariate binary design for drugs where monotonicity is assumed for both toxicity and efficacy. A design would also be incoherent if it de-escalates when the previous patient does not experience an efficacious response. An interesting situation from the table is that if the previous patient experiences no efficacious benefit and a toxic outcome then the design should stay at the current level (Cheung, 2011).

Mathematically, design  $\mathcal{D}$  which incorporates bivariate binary outcomes is said to be coherent if it satisfies conditions (a)-(e) with probability one.

- (a)  $\Pr_{\mathcal{D}}(x_j - x_{j-1} > 0 | Y_{j-1} = 1 \ \& \ Z_{j-1} = 1) = 0$  for all  $j$
- (b)  $\Pr_{\mathcal{D}}(x_j - x_{j-1} < 0 | Y_{j-1} = 0 \ \& \ Z_{j-1} = 1) = 0$  for all  $j$
- (c)  $\Pr_{\mathcal{D}}(x_j - x_{j-1} > 0 | Y_{j-1} = 1 \ \& \ Z_{j-1} = 0) = 0$  for all  $j$
- (d)  $\Pr_{\mathcal{D}}(x_j - x_{j-1} < 0 | Y_{j-1} = 1 \ \& \ Z_{j-1} = 0) = 0$  for all  $j$
- (e)  $\Pr_{\mathcal{D}}(x_j - x_{j-1} < 0 | Y_{j-1} = 0 \ \& \ Z_{j-1} = 0) = 0$  for all  $j$

If design  $\mathcal{D}$  satisfies conditions (a)-(d) but does not satisfy condition (e), then  $\mathcal{D}$  is not a coherent design because the design must satisfy all five conditions in order to be coherent.

### 3.2.2 Coherence for Bivariate Binary Targeted Agents

Table 3.2: Incoherent moves for Targeted Agents with Bivariate Binary Outcomes

	Toxic Outcome	Nontoxic Outcome
Response	Incoherent escalation: toxicity	No incoherent moves
Nonresponse	Incoherent escalation: toxicity	No incoherent moves

The idea of coherence needs to be refined for the type of agent being explored in a clinical trial. The previous section holds only for cytotoxic agents and not for targeted agents. When targeted agents are being analyzed, the assumption of monotonicity on the efficacy response may not hold. Therefore, it can be justified that a de-escalation may be appropriate when the previous patient does not experience an efficacious response since the peak of the dose-response curve may occur at any of the possible  $I$  dose levels. As a result of the lack of monotonicity in efficacious response there is no incoherent escalation or de-escalation due to the response outcome of the previous patient. Possible incoherent moves are summarized in Table 3.2 and only depend on the toxic outcome of the previous patient. Therefore, a design  $\mathcal{D}$  is

coherent for targeted agents only if with probability one Conditions (a) and (c) hold from Section 3.2.1. Thus,  $\mathcal{D}$  is coherent if  $\Pr_{\mathcal{D}}(x_j - x_{j-1} > 0 | Y_{j-1} = 1) = 0$  for all  $j$  and only depends on the toxicity outcomes. Cheung (2011) proved that the Bayesian CRM is coherent based only on Conditions 2 and 3 of Section 1.3.3. Therefore, the same proof can be applied to the method presented here and the design presented here is a coherent design.

## Chapter 4

# Application to Delayed Response

The proposed design in Chapter 2 is most appropriate when both binary toxicity and binary efficacy endpoints can be observed in a reasonably similar time frame. Therefore, a drawback of the method proposed in Chapter 2 is that it requires an efficacy response that occurs in a time frame similar to that of toxicity. Yin, Zheng, and Xu (2011) and Yuan and Yin (2009) note that in some practical situations, this may not be possible due to the fact that efficacy may occur much later than toxicity. For example, in a targeted agent oncology trial, toxicity outcomes can often be observed in a relatively timely manner after treatment administration but efficacy is observed in a relatively longer time frame. If this delay is the case, then the proposed method will fail in its existing form because the trial would have to pause before each patient is enrolled in order to fully observe the efficacy response for the previous patient. If the delay is particularly long, then it will cause the duration of the trial to be much too long and wastes resources (Yuan and Yin 2009). However, a modified approach can be applied to account for the delayed efficacy response. This would create a situation where we would be estimating toxicity probabilities based on more patient observations than efficacy probabilities. Since we have demonstrated that the association between toxicity and efficacy can

be excluded when modeling the responses and we are therefore modeling toxicity and efficacy independently, we can fit the likelihood for each response based on different amounts of data, and simply utilize the efficacy data we have available, even though it may be less than that of the toxicity data. This idea is explored with the use of modifications to the proposed methodology to handle such practical issues. Along these same lines, we incorporated time-to-event outcomes (TITE) for an effective extension of the method as a means of handling delayed response.

## 4.1 Lagged Response

One way of accounting for a delayed efficacy response is to assume that the efficacy response simply lags behind the toxicity outcome by a certain amount of time. For example, if it is assumed that the efficacy response takes five times as long to be observed over the toxicity outcome, then for the first five patients we allocate the next patient to a dose based solely on the toxicity model and the prior assumptions on efficacy. Therefore, the trial will be enrolling patients before the efficacy data has been fully observed on enrolled patients. Once an efficacious response has been observed, the method will switch to modeling both toxicity and efficacy with a slight tweak to the modeling approach. In Chapter 2, we assumed that the toxicity data was in the form of  $\Omega_j = \{(x_l, y_l) : l < j\}$  and the efficacy data in form of  $\mathcal{D}_j = \{(x_l, z_l) : l < j\}$ . But in the delayed response setting, we have toxicity data of the same form but the efficacy data now is  $\mathcal{D}_j = \{(x_l, z_l) : l < j - 5\}$  for  $j > 6$ . Therefore, we are modeling toxicity using data from all the previous  $j - 1$  patients but only modeling efficacy on the first  $j - 5$  patients. We let the number of patients that the efficacy response is delayed behind the toxicity response be denoted as  $n_{lag}$ . Thus, at the enrollment of the  $j^{th}$  patient, we have toxicity data in the form of  $\Omega_j = \{(x_l, y_l) : l < j\}$  and efficacy data in the form of  $\mathcal{D}_j = \{(x_l, z_l) : l < j - n_{lag}\}$

for  $j > n_{lag} + 1$ .

In Table C.1, a simulation study is presented where the parameter  $n_{lag}$  is varied from where there is no delay between the toxicity and efficacy outcomes to where there is a 15 patient lag between the follow-up for the efficacy and toxicity outcomes. The total sample size is set to  $n = 36$  with an AR sample size of  $n_R = 12$ . The simulation results align with our expectations that the longer efficacy lags behind that of efficacy, generally, the worse the method does at identifying the true OBD. In particular, the modified method performs poorly when  $n_{lag} = 15$  with a 48.6% PCS compared to a PCS of 87.2% for no lag and 60.4% for a 10 patient lag. The cause of this significant difference in performance can be attributed to the fact that the OBD is at the fourth dose level for this scenario and the efficacy model only has 21 patients to take into account in the modeling. Therefore, when the difference between the follow-up times for the toxicity and efficacy outcomes grows, the larger the sample size required to obtain similar operating characteristics to the method proposed in Chapter 2.

## 4.2 Review of TITE CRM

It was noted at the beginning of Chapter 4 that if a trial has a long duration of follow-up for each patient, then the trial duration using the original CRM may be impractical because the CRM requires the use of complete follow-up information. Therefore, a modification to the CRM incorporating the use of partial follow-up information was proposed by Cheung and Chappell (2000). The general idea is to incorporate all available information including the length of time a patient is on study but does not experience a toxicity into the model. Recall from Section 1.3.3 that  $\Omega_j = \{(x_l, y_l) : l < j\}$ . Let  $w_{l,j}$  denote the weight given to the  $l^{th}$  patient's observation prior to the entry of patient  $j$  which represents the proportion of follow-



up time completed and let  $Y_{l,j}$  be the toxicity indicator for patient  $l$  at the time of the enrollment of the  $j^{th}$  patient. Then the weighted likelihood of the model can be expressed through the equation

$$L(\beta; w|\Omega_j) = \prod_{l=1}^{j-1} \{w_{l,j} F(x_l, \beta)\}^{Y_{l,j}} \{1 - w_{l,j} F(x_l, \beta)\}^{1-Y_{l,j}}.$$

Then, as the  $j^{th}$  patient enters the trial, the parameter is estimated by

$$\hat{\beta}_j^w = \frac{\int_{\mathcal{B}} \beta L(\beta; w|\Omega_j) g(\beta) d\beta}{\int_{\mathcal{B}} L(\beta; w|\Omega_j) g(\beta) d\beta}$$

and then allocating the next patient to the dose such that

$$x_j = \arg \min_{d_i} |F(d_i, \hat{\beta}_j^w) - \phi_T|.$$

Various weight functions can be used in order to account for the times-to-event of the patients in the weighted likelihood. The weight  $w_{l,j}$  has the restriction of being in the interval  $[0, 1]$  and indicates the amount of information contributed to the the likelihood by patient  $l$ . When patient  $l$  has completed the follow-up period, the weight  $w_{l,j}$  equals 1. Therefore, if a new patient is only enrolled when all of the current patients have been completely followed, then the weighted likelihood simplifies to the regular binomial likelihood and as a result the TITE-CRM to the original CRM. Two weight functions were considered by Cheung and Chappell (2000) including a simple linear weight function defined as

$$w_{l,j} = w(t; T, x_l) = \min \left( \frac{t}{T}, 1 \right)$$

where  $t$  is the enrollment time of the  $l^{th}$  patient at the entry of the  $j^{th}$  patient and

$T$  is the maximum follow-up time for each patient. However, this weight function at first glance seems overly simplistic since it assumes an identical shape for all dose levels. An alternative then is to use an adaptive weighting scheme that updates the shape of the weight function based on the accrued observations. At the enrollment of the  $j^{th}$  patient, we define

$$w_{l,j}(t; T, x_l) = \frac{\kappa(t)}{c+1} + \frac{1}{c+1} \left\{ \frac{t - t_{(\kappa(t))}}{t_{(\kappa(t)+1)} - t_{(\kappa(t))}} \right\} \text{ for } t \leq T$$

where  $c$  is the total number of toxic outcomes regardless of dose level, the times-to-toxicity are ordered such that  $0 = t_{(0)} < t_{(1)} \leq \dots \leq t_{(c)} < t_{(c+1)} = T$ , and  $\kappa(t) = \max\{a \in [0, c] : t \geq t_{(a)}\}$ . When there are no toxic outcomes, the adaptive weight function simplifies to the linear weight function. Cheung and Chappell's (2000) study of the weight functions determined that the linear weight was sufficient in most situations and the use of the adaptive weight function does not improve the accuracy of the TITE-CRM. They noted that the use of the adaptive weighting scheme only effected the in-trial allocation of patients to doses in some situations where the delay of the toxicity is especially prominent. The adaptive weights can help reduce the amount of exposure to patients to overly toxic doses. However, in the targeted agents case this will not be an issue since there does not tend to be significant toxicity in these agents.

## 4.3 Extension of TITE to Bivariate Outcomes

### 4.3.1 TITE Model on Efficacy

In this section, a method is proposed in which the efficacy response has a much longer observation window than that of toxicity. Therefore, the toxicity response is still a bivariate outcome and the probability of toxicity is modeled through the

original CRM as detailed in Section 1.3.3 and the efficacy probability is modeled through the methods of Section 4.2 using the linear weight function. However, we still must incorporate the uncertainty of where the peak of the dose-efficacy curve is located through the introduction of the  $2I - 1$  models as discussed in Section 2.2. Therefore, the likelihood for model  $k$  at the enrollment of the  $j^{th}$  patient is given through the equation

$$L_k(\theta; w | \mathcal{D}_j) = \prod_{l=1}^{j-1} \{w_{l,j} G_k(x_l, \theta)\}^{z_{l,j}} \{1 - w_{l,j} G_k(x_l, \theta)\}^{1-z_{l,j}}$$

where  $\mathcal{D}_j = \{(x_l, z_l) : l < j\}$  as defined in Section 2.2. Then, we proceed in a similar manner by choosing the model with the highest posterior model probability model given by:

$$\Pr(M_k | \mathcal{D}_j) = \frac{\tau(k) \int_{\Theta} L_k(\theta; w | \mathcal{D}_j) g(\theta) d\theta}{\sum_{k=1}^{2I-1} \tau(k) \int_{\Theta} L_k(\theta; w | \mathcal{D}_j) g(\theta) d\theta}.$$

Each time a new patient enters the trial, the single model,  $k^*$ , with the largest posterior probability is chosen such that

$$k^* = \arg \max_k \Pr(M_k | \mathcal{D}_j), \quad k = 1, \dots, 2I - 1.$$

We then take the working model,  $G_{k^*}(d_i, \theta)$  associated with the model and obtain estimates of the efficacy probabilities at each dose through

$$\hat{\pi}_E(d_i) = G_{k^*}(d_i, \hat{\theta}_{jk^*})$$

Table 4.1: Example from the bortezomib trial as presented in Cheung (2011) to demonstrate the linear weight function used in the TITE CRM approach when the 5<sup>th</sup> patient enters the trial on April 16, 2004.

Patient	$x_l$	Entry Date	$Y_{l,5}$	$t_j$	$w_{l,5}$
1	3	2/3/04	0	73	$73/126 = 0.58$
2	3	2/10/04	0	66	$66/126 = 0.52$
3	3	3/12/04	0	35	$35/126 = 0.28$
4	3	3/19/04	0	28	$28/126 = 0.22$

where

$$\hat{\theta}_{jk^*}^w = \frac{\int_{\Theta} \theta L_{k^*}(\theta; w | \mathcal{D}_j) g(\theta) d\theta}{\int_{\Theta} L_{k^*}(\theta; w | \mathcal{D}_j) g(\theta) d\theta}.$$

Then, the algorithm proceeds in the same manner that is presented in Section 2.3 to determine dose allocation based on the estimated toxicity probabilities, estimated efficacy probabilities, and the toxicity tolerance,  $\phi_T$ .

As an example of the calculations of the linear weights, consider the example presented in Table 4.1. At the enrollment of the fifth patient, we have observed no toxicities out of the four already enrolled patients. The first patient was enrolled into the trial 73 days before the arrival of the current patient. Since the follow-up time for patients is pre-specified prior to the trial as 126 days,  $w_{1,5} = \frac{t_j}{T} = \frac{73}{126} = 0.58$ . Note that the units of the follow-up do not matter as long as both the follow-up time and the enrollment time are in the same units. Since a clinical trial is sequential in nature, we note here that the weights for a given patient will not decrease. For example, when the sixth patient enrolls, the first patient will have been followed for a longer period of time than when the fifth patient enrolls, so we know that for a given patient,  $l$ ,  $w_{l,j+1} > w_{l,j}$ .

A couple of different assumptions can be made on the arrival times of the patients in a clinical trial. First, we can assume that the patients will enter the trial

in a fixed manner. For example, we can assume that we have an eligible patient to enter the study every month. Another example of the distribution of arrival times is that we can consider them random events following a Poisson process with a rate parameter of (rate of patient accrual)/(length of patient follow-up). Since we do not know the underlying distribution of the TITE variable, we can simulate various underlying truths for the time-to-efficacy distribution. Two distributions were explored: a conditionally uniform model where the time of the efficacy event for a particular patient is generated only if the patient in question experienced an efficacious response. Specifically, the TITE would be randomly chosen from the interval  $(0, T)$  when a patient experiences an efficacious event. A different assumption of the underlying TITE distribution is that it follows a Weibull distribution with a fixed shape parameter of 4 as in Cheung and Chappell (2000). We take the same approach with the scale parameter,  $\gamma$ , as Cheung and Chappell (2000) and Braun (2006) took and set it to make the cumulative distribution function at  $T$ , the length of patient follow-up, equal to the probability of efficacy of the current dose level. Mathematically, we set  $\gamma$  equal to the value such that:

$$F(x; \gamma) = \int_0^T \frac{4}{\gamma} x^3 \exp^{-(x/\gamma)^4} dx = \pi_E(x_j).$$

Therefore, the failure time of the time-to-efficacy for the  $j^{th}$  patient will be a Weibull distribution with a shape parameter of

$$\gamma = \left[ \frac{-(T^4)}{\ln[1 - \pi_E(x_j)]} \right]^{1/4}.$$

Notice that the times-to-efficacy are set up in a way which naturally censors the observations. For example, the efficacy variable in the TITE setting is redefined to

be

$$Z_j = \begin{cases} 0 & \text{censored without efficacy at } t_{l,j} \\ 1 & \text{observed with efficacy at } t_{l,j} \end{cases}$$

where  $t_{l,j}$  is the follow-up time for the  $l^{th}$  patient at the enrollment of the  $j^{th}$  patient. The weight function automatically handles this notion of censoring. Specifically, when an enrolled patient, patient  $l$ , has not experienced an efficacious response and has not completed the entire follow-up period when a new patient arrives, then the patient's response is censored at  $t_{l,j}$ . If the patient experiences an efficacious event before the next patient is enrolled, then the patient's response will no longer be censored.

In order to examine the operating characteristics the performance of the proposed method using TITE outcomes for efficacy and the original CRM to model the toxicity probabilities, we conducted an extensive simulation study. The scenarios presented in Table C.2 compare the efficacious TITE outcome method, denoted TITE<sub>E</sub>, under various model assumptions to the approach with bivariate binary outcomes proposed in Section 2.3, denoted WT, and are presented in Yin, Zheng, and Xu (2013). Similar to their specifications, we set fixed the sample size at  $n = 60$  and a toxicity tolerance of  $\phi_T = 1/3$ . Since in Section 2.6.1 it was recommended to set the AR phase sample to a third of the total patients, we chose an AR sample size of  $n_R = 20$ . We performed 1,000 simulated trials for each scenario and the results are presented in Table C.2. Let  $p_1 = 0.10$ ,  $p_2 = 0.20$ ,  $p_3 = 0.30$ ,  $p_4 = 0.40$ ,  $p_5 = 0.50$ , and  $p_6 = 0.60$  and we consider  $2 \times 6 - 1 = 11$  possible models for efficacy:

1.  $q_{i1} = \{0.10, 0.20, 0.30, 0.40, 0.50, \mathbf{0.60}\}$
2.  $q_{i2} = \{0.10, 0.20, 0.30, 0.45, \mathbf{0.60}, 0.45\}$
3.  $q_{i3} = \{0.20, 0.30, 0.45, \mathbf{0.60}, 0.45, 0.30\}$

4.  $q_{i4} = \{0.30, 0.45, \mathbf{0.60}, 0.45, 0.30, 0.20\}$
5.  $q_{i5} = \{0.45, \mathbf{0.60}, 0.45, 0.30, 0.20, 0.10\}$
6.  $q_{i6} = \{\mathbf{0.60}, 0.50, 0.40, 0.30, 0.20, 0.10\}$
7.  $q_{i7} = \{\mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}\}$
8.  $q_{i8} = \{0.30, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}\}$
9.  $q_{i9} = \{0.20, 0.30, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}\}$
10.  $q_{i10} = \{0.20, 0.30, 0.40, \mathbf{0.60}, \mathbf{0.60}, \mathbf{0.60}\}$
11.  $q_{i11} = \{0.20, 0.30, 0.40, 0.50, \mathbf{0.60}, \mathbf{0.60}\}$ .

Since we have no reason to that believe that one model is more likely than another we will assume that, a priori,  $\tau(1) = \dots = \tau(11) = \frac{1}{11}$ . In each of the seven scenarios, we vary the method at which patients arrive to participate in the trial (fixed accrual rate vs. arrival as a Poisson Process) and the underlying distribution of the TITE efficacy outcome (Conditionally uniform vs. Weibull).

In each of the scenarios presented in Table C.2, the approach modeling efficacy as a TITE outcome performs similarly to the method proposed in Chapter 2. This fact is quite encouraging since we would expect the results of the TITE<sub>E</sub> method to diminish slightly compared to those of the proposed method as is seen in the second scenario where the proposed method has a PCS of 69.1% as compared with the TITE<sub>E</sub> method having PCS's varying from 62.5% to 68.3%. The last column of the table demonstrates the clear advantage of using the TITE<sub>E</sub> method over the proposed method: the trial duration is substantially shorter for the TITE<sub>E</sub> method. Since the proposed method requires the complete follow-up of each patient before the next patient can be enrolled into the trial, the duration of the trial is significantly long, 180 months (15 years), since the observation window is 3 months long for each

patient. The ability of the  $\text{TITE}_E$  method to make use of partial information and in turn enroll patients in a staggered fashion allows for the significant reduction of trial length to about 62 months (or roughly 5 years). Another advantage of the method proposed in this section is the robustness against the underlying TITE distribution and the arrival times. Whether the efficacy variable followed a conditionally Uniform or Weibull distribution with the parameters outlined in the previous section did not effect the operating characteristics of the method. Similarly, the fact that the true arrival times followed a uniform or a Poisson process did not diminish the ability of the  $\text{TITE}_E$  method to identify the OBD. This robustness is demonstrated by the fact that the largest difference between the two distributions is only a difference of about 6% in PCS for all of the scenarios.

### 4.3.2 TITE Model on Toxicity and Efficacy

A limitation of the method proposed in Section 4.3.1 is the different way toxicity and efficacy are modeled. However, we wished to assess the possible loss of information when modeling the toxicity response using TITE methods. Since the toxicity follow-up for a majority of early-phase clinical trials is only 28 days, which is a manageable time to assess follow-up, this was a reasonable approach. Note that the TITE CRM developed by Cheung and Chappell (2000) was proposed for trials in radiation oncology where the follow-up time for each patient was 90 days. A decrease in the duration of the trial was still able to be achieved by only modeling the efficacy response through TITE methods. However, in some cases the trial duration is still too long to be practical for an early-phase trial. One way to decrease the trial duration is to incorporate both toxicity and efficacy endpoints as TITE outcomes. Similar to using TITE to model the efficacy response, modeling the toxic response as a TITE outcome would allow the use of partial toxicity information and would



therefore allow for more than one patient to be monitored for a toxic response at a time. For example, if the length of follow-up for toxicity is 1 month, as was the case in the previous section, then a method requiring complete toxicity information would only enroll a single patient each month. On the other hand, a method modeling toxicity and efficacy as TITE outcomes can enroll multiple patients per month in order to significantly decrease trial duration.

Specifically, at the enrollment of the  $j^{th}$  patient, toxicity is modeled in the same manner as Section 4.2 to obtain the parameter estimate

$$\hat{\beta}_j^w = \frac{\int_{\mathcal{B}} \beta L(\beta; w | \Omega_j) g(\beta) d\beta}{\int_{\mathcal{B}} L(\beta; w | \Omega_j) g(\beta) d\beta}.$$

Then, we obtain the toxicity probability estimates

$$\hat{\pi}_{T_i} = F(d_i, \hat{\beta}_j^w)$$

and define the acceptable set as

$$\mathcal{A}_j = \{d_i : \hat{\pi}_{T_i} \leq \phi_T\}.$$

Efficacy is modeled through the methods proposed in Section 4.3.1 and then the proposed method uses the algorithm presented in Section 2.3 to determine the dose given to the  $j^{th}$  patient.

Table C.3 presents a simulation study examining the ability of the proposed method using TITE outcomes for both toxicity and efficacy, denoted  $\text{TITE}_{TE}$ , to identify the OBD under various assumptions. As in Table C.2, two ways of determining how patients arrive to participate in trial are examined: fixed accrual and Poisson arrival times. The survival distribution was varied between a condition-

ally Uniform and a Weibull distribution. Since we are now modeling toxicity and efficacy in the TITE setting, cases are presented where the survival distributions for toxicity and efficacy are the same and cases where they are different from one another. The similarities with regards to PCS in the results in Tables C.3 and C.2 are also quite encouraging. There is no substantial diminishment of the results in the TITE<sub>TE</sub> method when compared to the TITE<sub>E</sub> method. For example, in Scenario 6, where the OBD is located at  $d_6$ , the TITE<sub>E</sub> method recommended it in 31.0% of trials while the TITE<sub>TE</sub> method recommended it in 32.1% of trials when efficacy followed a conditionally Uniform survival distribution and the enrollment of patients was at fixed intervals. The last column of the table demonstrates the clear advantage of using the TITE<sub>TE</sub> approach over the TITE<sub>E</sub> approach: the trial duration is substantially shorter for the TITE<sub>TE</sub> method. Since the TITE<sub>E</sub> method requires complete follow-up for toxicity, the duration of the trial is significantly long, about 62 months, since the observation window is 1 month long for toxicity for each patient. The ability of the TITE<sub>TE</sub> method to make use of partial information for both toxicity and efficacy and in turn enroll patients in a staggered fashion allows for the significant reduction of trial length to about 22 months (or roughly 2 years), which is a much more reasonable trial duration for an early-phase clinical trial than the 62 months duration for the TITE<sub>E</sub> method.

Another advantage of the proposed TITE<sub>TE</sub> method is the robustness against the underlying TITE distribution and the arrival times. Whether the efficacy and toxicity variables followed a conditionally Uniform or Weibull distribution with the parameters outlined in Section 4.3.1 did not effect the operating characteristics of the method. Similarly, the fact that the true arrival times followed a uniform or a Poisson process did not diminish the ability of the TITE<sub>TE</sub> method to identify the OBD. This robustness is demonstrated by the fact that the largest difference between the two distributions is only a difference of about 7% in PCS for all of

the scenarios. Even in the cases where the survival distributions for toxicity and efficacy are different, the  $\text{TITE}_{TE}$  method performs similarly to when the survival distributions are the same.

# Chapter 5

## Conclusions and Future Research

### 5.1 Introduction

In this dissertation, we proposed a collection of new designs for effectively estimating the optimal biological dose in early phase trials that consider bivariate binary data in targeted agents. LoRusso, Boerner, and Seymour (2010), among others, have emphasized the need for new clinical trial designs for targeted agents different from those designed for cytotoxic agents. The proposed methods were developed within the framework of the CRM and partial order CRM by enumerating all the possible models for the probability of efficacy as dictated by the lack of monotonicity assumption on the dose-efficacy curve in targeted agents. The overall strategy of the methods is to lean upon Bayesian inference, prior clinical knowledge about the dose-toxicity and dose-efficacy curves, and allow the accumulating data to provide an estimate of the correct model from within a reasonable set of models. Then the methods lean upon the considerable work performed on the CRM to provide an efficient estimate of the OBD. Though generally not recommended, the method can handle cases where toxicity is not considered to be clinically significant as well as cases where a close eye must be simultaneously kept on toxicity when estimat-

ing the probability of efficacy. Simulation studies demonstrated this approach to perform competitively against existing methods and shows the ability of the design to correctly identify the OBD in a large portion of trials involving bivariate binary endpoints.

An extensive simulation study was performed over a wide choice of possible true unknown dose-toxicity and dose-efficacy relationships to demonstrate the behavior of the methodology presented in Chapter 2. Various design specifications were also varied to demonstrate the design's robustness. In particular, a simulation study was presented to demonstrate the design's robustness against different underlying correlation coefficients between toxicity and efficacy. After the broad simulation study, the theoretical properties were derived to provide confidence in using the proposed method. Furthermore, the design was extended to situations where the follow-up times of the toxicity and efficacy events can be different from each other. Therefore, the design is able to incorporate the use of partial follow-up information to make interim dose decisions.

## 5.2 Future Research

The methods presented in this dissertation were developed in the bivariate binary case where there are four possible outcomes. However, as discussed in Section 1.4.2, we can think of the possible outcomes in the bivariate outcomes case to be (1) no efficacy and no toxicity, (2) efficacy and no toxicity, and (3) toxicity. O'Quigley, Hughes, and Fenton (2001) proposed a method to extend the CRM to handle tri-nary outcomes in the case of an HIV trial. Their design was motivated by an HIV trial where a subject could experience either a toxicity, or a viral failure, or a success of reducing the viral load. The method proposes to use the CRM to model the probability of toxicity and the conditional probability of acceptable viral reduc-

tion given no toxicity and then using an independence assumption estimates the probability of success at each dose level to allocate the next patient to the dose with the highest estimated probability of success. However, in their paper they assume that the probability of an efficacious response given no toxicity is monotonic. This monotonicity assumption may not hold for targeted agents because of the reasons discussed in Section 1.4.1.

It is important to note that in the trinary outcome case, the goal of a trial is no longer one of finding the OBD, but of finding the Most Successful Dose (MSD). The MSD is defined as the dose with the highest overall success rate as estimated by the utility function,  $\mathcal{U} = \pi_E(d_i, \theta)\{1 - \pi_T(d_i, \beta)\}$ . In this case, we model the probability of toxicity as  $F(d_i, \beta)$ , the  $\Pr(\text{Efficacy}|\text{No Toxicity}) = G_k(d_i, \theta)$ , the probability of an efficacy response with no toxicity as  $G_k(d_i, \theta) \times (1 - F(d_i, \beta))$  where  $k = 1, \dots, 2 \times I - 1$  represent the working models for the different possible orderings for the probabilities of  $\Pr(\text{Efficacy}|\text{No toxicity})$ . The outcome for the  $j^{th}$  patient, denoted  $Y_j$ , follows a multinomial distribution where  $Y_j = 0$  if the  $j^{th}$  patient experiences no efficacy and no toxicity,  $Y_j = 1$  if the patient experiences an efficacious response but no toxicity, and  $Y_j = 2$  if the patient experiences a toxicity. Then, after the inclusion of  $j$  patients, the likelihood function of  $\beta$  and  $\theta$  is given by

$$L_j(\beta, \theta) = \prod_{l=1}^j p_0^{Y_{l0}} p_1^{Y_{l1}} (1 - p_0 - p_1)^{1 - Y_{l0} - Y_{l1}}$$

where  $Y_{l0}$  is an indicator function taking the value of 1 if patient  $l$  experiences no toxicity and no efficacy and 0 otherwise. The variable  $Y_{l1}$  is an indicator function taking the value of 1 if patient  $l$  experiences an efficacious response but no toxicity. Therefore, we can express  $p_0$  and  $p_1$  through the working models described above

and represented below.

$$\begin{cases} p_0 = \Pr(Y = 0) = (1 - F(d_i, \beta))(1 - G_k(d_i, \theta)) \\ p_1 = \Pr(Y = 1) = (1 - F(d_i, \beta))G_k(d_i, \theta) \\ p_2 = \Pr(Y = 2) = F(d_i, \beta) \text{ and } \Pr(Y = 1|Y \neq 2) = G_k(d_i, \theta) \end{cases}$$

The main advantage of using this parameterization is that the likelihood can be factored into two distinct parts: one that depends only on toxicities

$$\varphi_j(\beta) = \prod_{l=1}^j (1 - F(x_l, \beta))^{Y_{l0}+Y_{l1}} F(x_l, \beta)^{1-Y_{l0}-Y_{l1}}$$

and one that depends on efficacy given no toxicity

$$\varphi_{jk}(\theta) = \prod_{l=1}^j (1 - G_k(x_l, \theta))^{Y_{l0}} G_k(x_l, \theta)^{Y_{l1}}.$$

so that  $L_j(\beta, \theta) = \varphi_j(\beta)\varphi_{jk}(\theta)$ . This factorization allows for quicker evaluation of the posterior densities for  $\theta$  and  $\beta$  because only one-dimensional integrals are needed. Applying the same general idea used in the method proposed in Chapter 2 we place priors  $g(\beta)$  and  $g(\theta)$  on  $\beta$  and  $\theta$ , respectively. Then we obtain the posterior probability through the evaluation of

$$f(\beta|\mathcal{D}_j) = \frac{\varphi_j(\beta)g(\beta)}{\int_{\mathcal{B}} \varphi_j(\beta)g(\beta)d\beta}$$

Then, the estimates of the toxicity probabilities are obtained through  $\hat{\pi}_T(d_i) = \int_{\mathcal{B}} f(\beta)\beta d\beta$ . Similarly, the estimate of the posterior probabilities are obtained for

each of the  $2I - 1$  working models through the evaluation of

$$f_k(\theta|\mathcal{D}_j) = \frac{\varphi_{jk}(\theta)g(\theta)}{\int_{\Theta} \varphi_{jk}(\theta)g(\theta)d\theta}.$$

Then, we can choose the working model corresponding to the highest posterior probability given in the equation

$$\Pr(M_k|\mathcal{D}_j) = \frac{\tau(k) \int_{\Theta} \varphi_{jk}(\theta)g(\theta)d\theta}{\sum_{k=1}^{2I-1} \tau(k) \int_{\Theta} \varphi_{jk}(\theta)g(\theta)d\theta}$$

and obtain the estimates of the probabilities of the probability of efficacy given no toxicity through  $\hat{\pi}_{E|T^C}(d_i) = G_{k^*}(d_i, \hat{\theta}_{jk^*})$  where

$$\hat{\theta}_{jk^*} = \int_{\Theta} f_{k^*}(\theta)\theta d\theta$$

and

$$k^* = \arg \max_k \Pr(M_k|\mathcal{D}_j), k = 1, \dots, 2I - 1.$$

Once the probabilities of efficacy given no toxicities have been estimated, the method will then choose the dose given to the  $j^{th}$  patient,  $x_j$ , with the highest estimated probability of success through the equation

$$\Pr(\text{Success}) = \hat{\pi}_{E|T^C}(d_i) \times (1 - \hat{\pi}_T(d_i)).$$

Since the method uses features of the CRM, we would expect that the method would be competitive with existing methods in the trinary case. However, a rigorous exploration of the properties of the trinary outcomes method needs to be performed to provide justification for the use of such a design in practice. Various design



specifications would need to be explored as well including the use of an acceptable set similar to what is proposed in this dissertation, various skeleton specifications, and the choice of the working model.

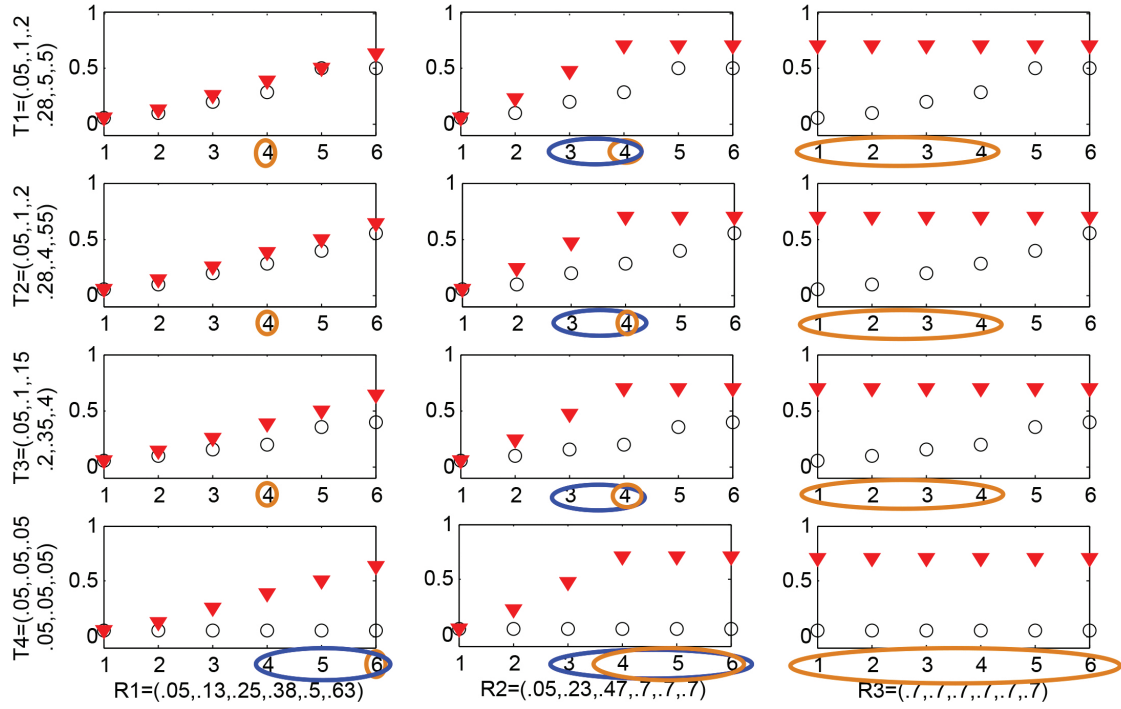
### **5.3 Concluding Remarks**

Another important extension of the methodology presented in this dissertation occurs when a combination of two targeted agents is used as a single treatment in a clinical trial. In these combination cases, the monotonicity assumption may or may not hold for toxicity. Therefore, we must take this uncertainty into account in the model through the use of the partial order CRM to model the probability of toxicity at each dose level.

# Appendix A

## Toxicity and Efficacy Scenarios

Figure A.1: Twelve scenarios used by Hoering et al. (2013). The Best doses are ones circled in orange and the good doses are circled in blue. If there is only orange, that indicates that the Best and Good doses are the same.



# Appendix B

## Simulation Results: Proposed Method

This section contains the simulation study results described in Chapter 2. In all tables, the OBD in each table is indicated by bold type. Table B.1 compares an existing method to a variant proposed method which only efficacy is modeled. Tables B.2 (comparison to Thall and Cook (2004), denoted by TC) and B.3 compare the operating characteristics of the proposed method (denoted WT) against existing methods in the literature. The second set of simulations, Tables B.4-B.10, explore the robustness of the proposed method under various design specifications including the AR sample size (Table B.4), skeleton choice (B.5), working model choice (B.6), two-parameter logistic (Table B.7), utility function (Table B.8), adaptive elimination of models (Table B.9), and model selection techniques (Table B.10).

Table B.1: *Selection Percentages for scenarios presented in Zang, Lee, and Yuan (2014) at each dose comparing the proposed method using only the efficacy response, denoted ‘WTE’, and the various methods in a previous version of the authors’ paper after 1000 simulation runs with a sample size of  $n = 30$  and  $n_R = 10$ . Peak and plateau curves were considered as possible models in this simulation.*

Scenario	True efficacy probability				
	0.20	0.40	0.60	<b>0.80</b>	0.55
Scenario 1					
Slope-sign	3.9	11.6	12.4	<b>48.0</b>	24.2
L-Logistic	2.4	6.5	16.0	<b>64.8</b>	10.4
Isotonic	12.7	13.0	12.8	<b>55.9</b>	5.6
WTE	1.4	6.4	21.3	<b>64.0</b>	6.9
Scenario 2	0.60	<b>0.80</b>	0.65	0.40	0.20
Slope-sign	9.5	<b>42.7</b>	37.8	9.0	0.8
L-Logistic	11.1	<b>59.5</b>	26.8	2.5	0.1
Isotonic	13.0	<b>65.8</b>	19.8	1.3	0.1
WTE	17.3	<b>63.7</b>	16.2	2.3	0.5
Scenario 3	0.05	0.25	0.45	0.65	<b>0.80</b>
Slope-sign	1.3	12.4	6.6	17.7	<b>62.0</b>
L-Logistic	0.6	3.8	8.3	23.0	<b>64.3</b>
Isotonic	6.1	12.2	13.6	20.6	<b>47.5</b>
WTE	0.2	2.0	7.7	25.4	<b>64.8</b>
Scenario 4	0.10	0.30	<b>0.50</b>	0.50	0.50
Slope-sign	2.5	14.3	<b>23.8</b>	27.3	32.2
L-Logistic	1.8	9.8	<b>33.8</b>	30.8	23.8
Isotonic	9.9	14.2	<b>35.5</b>	25.1	15.2
WTE	1.6	9.4	<b>35.9</b>	31.3	21.9

Table B.1 – continued from previous page

	True efficacy probability				
Scenario 5	0.10	0.30	0.45	<b>0.60</b>	0.60
Slope-sign	2.6	16.7	12.9	<b>27.3</b>	40.5
L-Logistic	1.8	8.7	17.1	<b>37.5</b>	35.0
Isotonic	10.1	16.1	18.8	<b>33.3</b>	21.7
WTE	1.3	6.9	19.1	<b>41.8</b>	31.0
Scenario 6	0.20	<b>0.40</b>	0.40	0.40	0.40
Slope-sign	6.6	<b>31.3</b>	22.2	19.0	20.9
L-Logistic	5.5	<b>29.9</b>	27.1	22.0	15.5
Isotonic	14.0	<b>34.1</b>	26.4	16.6	9.0
WTE	7.4	<b>29.4</b>	27.1	21.0	15.2
Scenario 7	0.40	0.60	<b>0.80</b>	0.70	0.55
Slope-sign	3.4	11.6	<b>35.5</b>	32.5	17.0
L-Logistic	4.6	11.7	<b>51.1</b>	26.8	5.9
Isotonic	12.0	14.0	<b>48.9</b>	22.1	2.9
WTE	4.2	15.4	<b>56.0</b>	21.3	3.1
Scenario 8	<b>0.8</b>	0.6	0.4	0.3	0.2
Slope-sign	<b>45.7</b>	34.7	15.8	3.0	0.9
L-Logistic	<b>69.5</b>	24.3	4.7	1.2	0.2
Isotonic	<b>77.2</b>	19.9	2.3	0.6	0.1
WTE	<b>74.6</b>	16.1	4.6	3.2	1.6

Table B.2: *Selection Percentages for each dose of Thall and Cook (2004), denoted TC, and the proposed method, denoted ‘WT’, after 1000 simulation runs with a sample size of  $n = 36$ ,  $n_R = 12$ , and  $\phi_T = 0.40$ . The true value of the association parameter used to simulate observations is given by  $\psi$ . Only Peak models were considered as candidate models in this simulation.*

Scenario 1	True (toxicity, efficacy) probability			
	(0.05, 0.02)	(0.12, 0.30)	<b>(0.30, 0.55)</b>	(0.80, 0.65)
TC ( $\psi = 0$ )	0	17.3	<b>80.3</b>	2.4
WT ( $\psi = 0$ )	0.6	6.5	<b>85.4</b>	7.5
WT ( $\psi = -2.049$ )	1.3	6.2	<b>86.0</b>	6.5
WT ( $\psi = -0.814$ )	0.8	6.2	<b>85.6</b>	7.4
WT ( $\psi = 0.814$ )	0.6	6.3	<b>85.7</b>	7.4
WT ( $\psi = 2.0486$ )	0.3	7.5	<b>84.8</b>	7.4
Scenario 2	(0.05, 0.02)	(0.10, 0.28)	(0.16, 0.50)	<b>(0.22, 0.80)</b>
TC ( $\psi = 0$ )	0	0.5	13.6	<b>86</b>
WT ( $\psi = 0$ )	0.1	0.9	11.8	<b>87.2</b>
WT ( $\psi = -2.049$ )	0.2	0.8	11.6	<b>87.3</b>
WT ( $\psi = -0.814$ )	0.0	0.9	11.8	<b>87.2</b>
WT ( $\psi = 0.814$ )	0.0	0.6	11.4	<b>87.9</b>
WT ( $\psi = 2.0486$ )	0.0	0.8	11.2	<b>88.0</b>
Scenario 3	(0.05, 0.25)	<b>(0.15, 0.65)</b>	(0.42, 0.50)	(0.65, 0.05)
TC ( $\psi = 0$ )	5.2	<b>81.3</b>	13.1	0.2
WT ( $\psi = 0$ )	1.5	<b>81.8</b>	16.4	0.2
WT ( $\psi = -2.049$ )	1.5	<b>81.6</b>	16.4	0.1
WT ( $\psi = -0.814$ )	1.5	<b>82.0</b>	16.8	0.1
WT ( $\psi = 0.814$ )	1.9	<b>82.1</b>	16.4	0.1
WT ( $\psi = 2.0486$ )	1.3	<b>82.1</b>	16.5	0.2
Scenario 4	<b>(0.05, 0.80)</b>	(0.10, 0.50)	(0.16, 0.28)	(0.22, 0.02)
TC ( $\psi = 0$ )	<b>92.5</b>	4.8	2.1	0.4
WT ( $\psi = 0$ )	<b>91.9</b>	7.5	0.6	0.0
WT ( $\psi = -2.049$ )	<b>92.1</b>	7.5	0.4	0.0
WT ( $\psi = -0.814$ )	<b>92.6</b>	6.7	0.6	0.0
WT ( $\psi = 0.814$ )	<b>92.2</b>	7.0	0.7	0.1
WT ( $\psi = 2.0486$ )	<b>91.8</b>	7.4	0.8	0.0
Scenario 5	<b>(0.05, 0.45)</b>	(0.45, 0.50)	(0.70, 0.55)	(0.85, 0.60)
TC ( $\psi = 0$ )	<b>72.9</b>	26.9	0.1	0.0
WT ( $\psi = 0$ )	<b>46.5</b>	53.3	0.2	0.0
WT ( $\psi = -2.049$ )	<b>44.4</b>	55.4	0.2	0.0
WT ( $\psi = -0.814$ )	<b>44.1</b>	55.9	0.2	0.0
WT ( $\psi = 0.814$ )	<b>41.2</b>	58.6	0.2	0.0
WT ( $\psi = 2.0486$ )	<b>42.3</b>	57.5	0.1	0.0

Table B.3: *Operating characteristics of the proposed design and Hoering et al. (2013). Overall probability of selecting a best or good dose level as the OBD after 1000 simulation runs with a maximum sample size of  $n = 64$  patients,  $n_R = 16$ , and  $\phi_T = 1/3$ . Peak and plateau models were considered as candidate models in this simulation.*

Efficacy Scenario	Toxicity Scenario	Best Dose Level	Good Dose Level	Probability of Picking <i>Best</i> Dose with Our Proposed Design	Probability of Picking <i>Good</i> Dose with Our Proposed Design	Probability of Picking <i>Best</i> Dose with Hoering et al. Design	Probability of Picking <i>Good</i> Dose with Hoering et al. Design
R1	T1	4	4	<b>0.60</b>	<b>0.60</b>	0.21	0.21
R1	T2	4	4	<b>0.52</b>	<b>0.52</b>	0.20	0.20
R1	T3	4	4	<b>0.40</b>	<b>0.40</b>	0.31	0.31
R1	T4	6	4-6	0.54	<b>0.93</b>	<b>0.63</b>	0.89
R2	T1	4	3, 4	<b>0.66</b>	<b>0.90</b>	0.30	0.68
R2	T2	4	3, 4	<b>0.58</b>	<b>0.82</b>	0.30	0.65
R2	T3	4	3, 4	<b>0.49</b>	0.59	0.42	<b>0.71</b>
R2	T4	4-6	3-6	<b>0.94</b>	<b>1.00</b>	0.91	0.94
R3	T1	1-4	1-4	<b>0.99</b>	<b>0.99</b>	0.93	0.93
R3	T2	1-4	1-4	<b>0.99</b>	<b>0.99</b>	0.90	0.90
R3	T3	1-4	1-4	<b>0.97</b>	<b>0.97</b>	0.84	0.85
R3	T4	1-6	1-6	<b>1.00</b>	<b>1.00</b>	0.97	0.97

Table B.4: *Recommendation Percentages for the Proposed Method under various Adaptive Randomization sample sizes after 1000 simulation runs with a sample size of  $n = 36$  and  $\phi_T = 0.40$ . Peak and plateau models were considered as candidate models in this simulation.*

	True (toxicity, efficacy) probability					
Scenario 1	(0.05, 0.05)	(0.10, 0.13)	(0.20, 0.25)	<b>(0.28, 0.38)</b>	(0.50, 0.50)	(0.50, 0.63)
$n_R = 0$	1.3	6.1	34.6	<b>52.9</b>	5.1	0.0
$n_R = 12$	0.5	6.6	34.3	<b>51.1</b>	7.3	0.2
$n_R = 18$	0.9	6.8	31.9	<b>51.4</b>	8.8	0.2
$n_R = 24$	1.2	5.3	36.2	<b>48.1</b>	8.9	0.3
Scenario 2	(0.05, 0.05)	(0.10, 0.23)	(0.20, 0.47)	<b>(0.28, 0.70)</b>	(0.50, 0.70)	(0.50, 0.70)
$n_R = 0$	0.2	6.1	45.5	<b>46.1</b>	1.9	0.2
$n_R = 12$	0.1	6.4	43.2	<b>48.8</b>	1.4	0.1
$n_R = 18$	0.0	5.9	43.7	<b>47.2</b>	3.0	0.2
$n_R = 24$	0.0	4.2	45.3	<b>47.6</b>	2.9	0.0
$n_R = 36$	0.0	5.9	49.3	<b>41.9</b>	2.9	0.0
Scenario 3	(0.05, 0.05)	(0.05, 0.13)	(0.05, 0.25)	(0.05, 0.38)	(0.05, 0.50)	<b>(0.05, 0.63)</b>
$n_R = 0$	0.8	2.9	8.0	20.2	34.1	<b>34.0</b>
$n_R = 12$	0.4	2.4	9.9	19.7	32.6	<b>35.0</b>
$n_R = 18$	0.3	1.1	7.5	19.6	31.2	<b>40.3</b>
$n_R = 24$	0.1	1.5	7.2	20.6	32.2	<b>38.4</b>
$n_R = 36$	0.1	2.1	10.0	27.4	34.7	<b>25.7</b>
Scenario 4	(0.05, 0.05)	(0.10, 0.23)	(0.20, 0.47)	<b>(0.28, 0.70)</b>	(0.40, 0.70)	(0.55, 0.70)
$n_R = 0$	0.0	6.7	44.4	<b>45.4</b>	3.4	0.1
$n_R = 12$	0.2	5.2	41.6	<b>48.3</b>	4.7	0.0
$n_R = 18$	0.0	3.3	46.2	<b>46.7</b>	3.5	0.3
$n_R = 24$	0.0	3.7	44.6	<b>46.1</b>	5.4	0.2
$n_R = 36$	0.1	4.8	49.6	<b>41.8</b>	3.6	0.1



Table B.5: *Selection Percentages for each dose for the proposed method under various skeleton specifications after 1000 simulation runs with a sample size of  $n = 36$ ,  $n_R = 12$ , and  $\phi_T = 0.40$ . Peak and plateau models were considered as candidate models in this simulation.*

Scenario	True (toxicity, efficacy) probability			
	(0.05, 0.02)	(0.12, 0.30)	<b>(0.30, 0.55)</b>	(0.80, 0.65)
Scenario 1				
$\alpha_1, Q_1$	0.1	21.4	<b>74.9</b>	3.6
$\alpha_2, Q_1$	0.0	17.7	<b>80.5</b>	1.8
$\alpha_3, Q_1$	0.2	18.3	<b>80.5</b>	1.0
$\alpha_4, Q_1$	0.2	15.3	<b>78.8</b>	5.7
$\alpha_1, Q_2$	0.6	22.4	<b>73.4</b>	3.6
Scenario 2	(0.05, 0.02)	(0.10, 0.28)	(0.16, 0.50)	<b>(0.22, 0.80)</b>
$\alpha_1, Q_1$	0.1	6.2	19.5	<b>74.2</b>
$\alpha_2, Q_1$	0.0	8.5	24.0	<b>67.5</b>
$\alpha_3, Q_1$	0.2	6.4	22.7	<b>70.7</b>
$\alpha_4, Q_1$	0.2	5.4	15.6	<b>78.8</b>
$\alpha_1, Q_2$	0.1	7.6	18.3	<b>74.0</b>
Scenario 3	(0.05, 0.25)	<b>(0.15, 0.65)</b>	(0.42, 0.50)	(0.65, 0.05)
$\alpha_1, Q_1$	4.4	<b>85.0</b>	10.6	0.0
$\alpha_2, Q_1$	2.7	<b>86.5</b>	10.8	0.0
$\alpha_3, Q_1$	4.3	<b>85.0</b>	10.7	0.0
$\alpha_4, Q_1$	5.9	<b>84.3</b>	9.8	0.0
$\alpha_1, Q_2$	3.9	<b>90.6</b>	5.5	0.0
Scenario 4	<b>(0.05, 0.80)</b>	(0.10, 0.50)	(0.16, 0.28)	(0.22, 0.02)
$\alpha_1, Q_1$	<b>90.7</b>	7.8	1.5	0.0
$\alpha_2, Q_1$	<b>92.1</b>	6.1	1.8	0.0
$\alpha_3, Q_1$	<b>90.7</b>	7.3	2.0	0.0
$\alpha_4, Q_1$	<b>90.8</b>	7.0	2.1	0.1
$\alpha_1, Q_2$	<b>94.4</b>	4.9	0.7	0.0

Table B.6: *Selection Percentages for each dose of the proposed method under various working model selections after 1000 simulation runs with a sample size of  $n = 36$ ,  $n_R = 12$ , and  $\phi_T = 0.40$ . Peak and plateau models were considered as candidate models in this simulation.*

	True (toxicity, efficacy) probability			
Scenario 1	(0.05, 0.02)	(0.12, 0.30)	<b>(0.30, 0.55)</b>	(0.80, 0.65)
$p_i^{\exp(\beta)}$	0.1	21.4	<b>74.9</b>	3.6
Logistic ( $\alpha_0 = 3$ )	0.8	25.3	<b>71.1</b>	2.8
2-parameter Logistic	24.8	34.5	<b>39.0</b>	1.7
Scenario 2	(0.05, 0.02)	(0.10, 0.28)	(0.16, 0.50)	<b>(0.22, 0.80)</b>
$p_i^{\exp(\beta)}$	0.1	6.2	19.5	<b>74.2</b>
Logistic ( $\alpha_0 = 3$ )	1.0	10.6	19.0	<b>69.4</b>
2-parameter Logistic	12.8	7.0	22.0	<b>58.2</b>
Scenario 3	(0.05, 0.25)	<b>(0.15, 0.65)</b>	(0.42, 0.50)	(0.65, 0.05)
$p_i^{\exp(\beta)}$	4.4	<b>85.0</b>	10.6	0.0
Logistic ( $\alpha_0 = 3$ )	8.5	<b>84.3</b>	7.2	0.0
2-parameter Logistic	27.3	<b>65.0</b>	7.7	0.0
Scenario 4	<b>(0.05, 0.80)</b>	(0.10, 0.50)	(0.16, 0.28)	(0.22, 0.02)
$p_i^{\exp(\beta)}$	<b>90.7</b>	7.8	1.5	0.0
Logistic ( $\alpha_0 = 3$ )	<b>90.3</b>	8.0	1.7	0.0
2-parameter Logistic	<b>92.9</b>	6.8	0.3	0.0
Scenario 5	<b>(0.05, 0.45)</b>	(0.45, 0.50)	(0.70, 0.55)	(0.85, 0.60)
$p_i^{\exp(\beta)}$	<b>46.5</b>	53.3	0.2	0.0
Logistic ( $\alpha_0 = 3$ )	<b>66.5</b>	33.0	0.5	0.0
2-parameter Logistic	<b>92.9</b>	6.8	0.3	0.0

Table B.7: *Selection Percentages for each dose of the proposed method under various degrees of the use of the 2-parameter logistic function after 1000 simulation runs with a sample size of  $n = 36$ ,  $n_R = 12$ , and  $\phi_T = 0.40$ . Peak and plateau models were considered as candidate models in this simulation. The Logistic ( $\alpha_0 = 3$ ) row corresponds to the same row in Table B.6. ‘Only Tox’ represents a method in which toxicity is modeled through the 2-parameter logistic working model but efficacy is modeled through the 1-parameter logistic working model, vice versa for ‘Only Eff’, and ‘Both’ represents a method in which both toxicity and efficacy are modeled through the 2-parameter logistic working model.*

Scenario	True (toxicity, efficacy) probability			
	(0.05, 0.02)	(0.12, 0.30)	<b>(0.30, 0.55)</b>	(0.80, 0.65)
Scenario 1				
Logistic ( $\alpha_0 = 3$ )	0.8	25.3	<b>71.1</b>	2.8
Only Tox	8.3	40.6	<b>49.4</b>	1.7
Only Eff	14.1	24.5	<b>59.0</b>	2.4
Both	24.8	34.5	<b>39.0</b>	1.7
Scenario 2	(0.05, 0.02)	(0.10, 0.28)	(0.16, 0.50)	<b>(0.22, 0.80)</b>
Logistic ( $\alpha_0 = 3$ )	1.0	10.6	19.0	<b>69.4</b>
Only Tox	4.4	7.8	14.8	<b>73.0</b>
Only Eff	7.7	10.6	24.1	<b>57.6</b>
Both	12.8	7.0	22.0	<b>58.2</b>
Scenario 3	(0.05, 0.25)	<b>(0.15, 0.65)</b>	(0.42, 0.50)	(0.65, 0.05)
Logistic ( $\alpha_0 = 3$ )	8.5	<b>84.3</b>	7.2	0.0
Only Tox	18.3	<b>77.2</b>	4.5	0.0
Only Eff	12.3	<b>77.6</b>	9.4	0.4
Both	27.3	<b>65.0</b>	7.7	0.0
Scenario 4	<b>(0.05, 0.80)</b>	(0.10, 0.50)	(0.16, 0.28)	(0.22, 0.02)
Logistic ( $\alpha_0 = 3$ )	<b>90.3</b>	8.0	1.7	0.0
Only Tox	<b>90.3</b>	7.1	2.5	0.1
Only Eff	<b>89.6</b>	8.9	1.3	0.2
Both	<b>85.2</b>	11.3	3.3	0.2
Scenario 5	<b>(0.05, 0.45)</b>	(0.45, 0.50)	(0.70, 0.55)	(0.85, 0.60)
Logistic ( $\alpha_0 = 3$ )	<b>66.5</b>	53.3	0.2	0.0
Only Tox	<b>89.6</b>	10.2	0.2	0.0
Only Eff	<b>64.9</b>	34.9	0.2	0.3
Both	<b>92.9</b>	6.8	0.3	0.0

Table B.8: *Recommendation Percentages for the Proposed Method and the Success Probability Utility Function after 1000 simulation runs with a sample size of  $n = 36$ ,  $n_R = 12$ , and  $\phi_T = 0.40$ . Peak and plateau models were considered as candidate models in this simulation. ‘WT’ represents the method proposed in Chapter 2, ‘Success Probability’ represents a method method which uses the success probability utility function, and ‘Success Probability (A)’ represents a method which uses the success probability utility function coupled with an acceptable set.*

Scenario 1	True (toxicity, efficacy) probability			
	(0.05, 0.02)	(0.12, 0.30)	<b>(0.30, 0.55)</b>	(0.80, 0.65)
WT	0.1	21.4	<b>74.9</b>	3.6
Success Probability	0.1	7.5	<b>41.8</b>	50.6
Success Probability (A)	0.5	20.9	<b>75.1</b>	3.5
Scenario 2	(0.05, 0.02)	(0.10, 0.28)	(0.16, 0.50)	<b>(0.22, 0.80)</b>
WT	0.1	6.2	19.5	<b>74.2</b>
Success Probability	0.2	4.3	12.6	<b>82.9</b>
Success Probability (A)	0.1	9.2	16.9	<b>73.8</b>
Scenario 3	(0.05, 0.25)	<b>(0.15, 0.65)</b>	(0.42, 0.50)	(0.65, 0.05)
WT	4.4	<b>85.0</b>	10.6	0.0
Success Probability	4.4	<b>76.4</b>	19.1	0.1
Success Probability (A)	4.2	<b>87.1</b>	8.6	0.1
Scenario 4	<b>(0.05, 0.80)</b>	(0.10, 0.50)	(0.16, 0.28)	(0.22, 0.02)
WT	<b>90.7</b>	7.8	1.5	0.0
Success Probability	<b>90.7</b>	7/9	1.3	0.1
Success Probability (A)	<b>92.0</b>	6.6	1.4	0.0

Table B.9: *Selection Percentages for each dose of the proposed method under various adaptive elimination of orders after 1000 simulation runs with a sample size of  $n = 36$ ,  $n_R = 12$ , and  $\phi_T = 0.40$ . Peak and plateau models were considered as candidate models in this simulation. ‘WT’ represents the method proposed in Chapter 2, ‘E 0.5 at  $n_R$ ’ represents a method which eliminates half of the candidate models at  $n_R$ , ‘E twice’ represents a method which eliminates one-fourth of the candidate models at  $n_R$  and then eliminates another one-fourth at  $3/4$  of the way through the trial, and ‘E at  $3/4$  way’ represents a method which eliminates half of the candidate models at  $3/4$  of the way through the trial.*

	True (toxicity, efficacy) probability			
Scenario 1	(0.05, 0.02)	(0.12, 0.30)	<b>(0.30, 0.55)</b>	(0.80, 0.65)
WT	0.1	21.4	<b>74.9</b>	3.6
E 0.5 at $n_R$	5.0	34.3	<b>43.7</b>	17.0
E twice	4.9	34.2	<b>44.6</b>	16.3
E at $3/4$ way	0.7	26.3	<b>67.7</b>	5.3
Scenario 2	(0.05, 0.02)	(0.10, 0.28)	(0.16, 0.50)	<b>(0.22, 0.80)</b>
WT	0.1	6.2	19.5	<b>74.2</b>
E 0.5 at $n_R$	3.5	14.7	24.4	<b>57.4</b>
E twice	1.9	15.7	24.8	<b>57.6</b>
E at $3/4$ way	0.2	8.5	21.4	<b>69.9</b>
Scenario 3	(0.05, 0.25)	<b>(0.15, 0.65)</b>	(0.42, 0.50)	(0.65, 0.05)
WT	4.4	<b>85.0</b>	10.6	0.0
E 0.5 at $n_R$	14.5	<b>66.3</b>	18.2	1.0
E twice	14.0	<b>69.1</b>	16.4	0.5
E at $3/4$ way	5.0	<b>84.8</b>	10.1	0.1
Scenario 4	<b>(0.05, 0.80)</b>	(0.10, 0.50)	(0.16, 0.28)	(0.22, 0.02)
WT	<b>90.7</b>	7.8	1.5	0.0
E 0.5 at $n_R$	<b>87.3</b>	9.1	3.5	0.1
E twice	<b>84.9</b>	11.4	3.5	0.2
E at $3/4$ way	<b>91.6</b>	6.8	1.5	0.1

Table B.10: *Selection Percentages for each dose of the proposed method under various model selection procedures after 1000 simulation runs with a sample size of  $n = 36$ ,  $n_R = 12$ , and  $\phi_T = 0.40$ . Peak and plateau models were considered as candidate models in this simulation. ‘WT’ represents the proposed method in Chapter 2, ‘BMA’ represents a method which uses Bayesian Model Averaging to obtain the probability estimates for efficacy, and ‘BMAO’ represents a method which uses Bayesian Model Averaging but excludes models that are deemed to be far worse than the best model in terms of  $\delta$ .*

	True (toxicity, efficacy) probability			
Scenario 1	(0.05, 0.02)	(0.12, 0.30)	<b>(0.30, 0.55)</b>	(0.80, 0.65)
WT	0.1	21.4	<b>74.9</b>	3.6
BMA	0.2	13.2	<b>82.1</b>	4.5
BMAO, $\delta = 0.6$	0.2	16.0	<b>79.9</b>	3.9
BMAO, $\delta = 0.7$	0.1	18.4	<b>77.9</b>	3.6
Scenario 2	(0.05, 0.02)	(0.10, 0.28)	(0.16, 0.50)	<b>(0.22, 0.80)</b>
WT	0.1	6.2	19.5	<b>74.2</b>
BMA	0.0	2.7	23.7	<b>73.6</b>
BMAO, $\delta = 0.6$	0.2	3.8	17.0	<b>79.0</b>
BMAO, $\delta = 0.7$	0.0	5.5	17.3	<b>77.2</b>
Scenario 3	(0.05, 0.25)	<b>(0.15, 0.65)</b>	(0.42, 0.50)	(0.65, 0.05)
WT	4.4	<b>85.0</b>	10.6	0.0
BMA	1.0	<b>81.1</b>	17.9	0.0
BMAO, $\delta = 0.6$	2.4	<b>86.4</b>	11.0	0.2
BMAO, $\delta = 0.7$	2.6	<b>87.8</b>	9.6	0.0
Scenario 4	<b>(0.05, 0.80)</b>	(0.10, 0.50)	(0.16, 0.28)	(0.22, 0.02)
WT	<b>90.7</b>	7.8	1.5	0.0
BMA	<b>76.5</b>	21.9	1.6	0.0
BMAO, $\delta = 0.6$	<b>90.2</b>	7.8	2.0	0.0
BMAO, $\delta = 0.7$	<b>89.7</b>	8.7	1.6	0.0

# Appendix C

## Simulation Results: Delayed Response

This section contains the simulation study results described in Chapter 4. Table C.1 shows the results of the proposed method under varying degrees of lag between the efficacy and toxicity responses. Table C.2 examines a method in which efficacy is modeled using TITE methods and Table C.3 examines a method where both toxicity and efficacy are modeled using TITE methods. As before, the OBD in each table is indicated by bold type. Each of the three simulations presented in this Appendix consider peak and plateau models as candidate models. In the tables that follows, ‘TITE<sub>E</sub>’ denotes a method which models toxicity as a binary outcome and efficacy as a TITE outcome, ‘TITE<sub>TE</sub>’ denotes a method which models toxicity and efficacy as TITE outcomes, ‘U’ represents that the TITE outcome follows a conditionally Uniform distribution, ‘W’ represents that the TITE outcomes follows a Weibull distribution, ‘f’ represents that the arrival of patients occurs at fixed intervals, and ‘P’ represents that the arrival of patients is a random event which follows a Poisson process.

Table C.1: *Selection Percentages for each dose the proposed method under different amounts of delay after 1000 simulation runs with a sample size of  $n = 36$ ,  $n_R = 12$ , and  $\phi_T = 0.40$ . When  $n_{lag} = 0$ , the method is the same as proposed in Chapter 2 and corresponds to  $\alpha_1$ ,  $Q_1$  in Table B.5.*

	True (toxicity, efficacy) probability			
Scenario 1	(0.05, 0.02)	(0.12, 0.30)	<b>(0.30, 0.55)</b>	(0.80, 0.65)
$n_{lag} = 0$	0.1	21.4	<b>74.9</b>	3.6
$n_{lag} = 3$	1.5	7.5	<b>88.0</b>	4.5
$n_{lag} = 5$	0.1	8.1	<b>85.3</b>	6.6
$n_{lag} = 8$	0.3	9.2	<b>83.6</b>	7.0
$n_{lag} = 10$	0.2	8.8	<b>84.5</b>	6.6
$n_{lag} = 15$	0.1	15.1	<b>78.6</b>	6.3
Scenario 2	(0.05, 0.02)	(0.10, 0.28)	(0.16, 0.50)	<b>(0.22, 0.80)</b>
$n_{lag} = 0$	0.1	6.2	19.5	<b>74.2</b>
$n_{lag} = 3$	0.0	1.5	23.5	<b>75.1</b>
$n_{lag} = 5$	0.1	1.5	28.5	<b>70.0</b>
$n_{lag} = 8$	0.0	1.6	32.5	<b>65.9</b>
$n_{lag} = 10$	0.1	2.7	36.9	<b>60.4</b>
$n_{lag} = 15$	0.0	7.8	44.2	<b>48.6</b>
Scenario 3	(0.05, 0.25)	<b>(0.15, 0.65)</b>	(0.42, 0.50)	(0.65, 0.05)
$n_{lag} = 0$	4.4	<b>85.0</b>	10.6	0.0
$n_{lag} = 3$	1.9	<b>82.3</b>	15.8	0.1
$n_{lag} = 5$	2.4	<b>81.1</b>	16.1	0.5
$n_{lag} = 8$	2.8	<b>79.5</b>	17.0	0.8
$n_{lag} = 10$	2.8	<b>77.7</b>	18.7	0.9
$n_{lag} = 15$	4.3	<b>74.3</b>	20.4	1.1
Scenario 4	<b>(0.05, 0.80)</b>	(0.10, 0.50)	(0.16, 0.28)	(0.22, 0.02)
$n_{lag} = 0$	<b>90.7</b>	7.8	1.5	0.0
$n_{lag} = 3$	<b>93.4</b>	6.3	0.3	0.1
$n_{lag} = 5$	<b>92.4</b>	7.0	0.6	0.0
$n_{lag} = 8$	<b>92.7</b>	6.7	0.5	0.1
$n_{lag} = 10$	<b>92.1</b>	7.3	0.5	0.2
$n_{lag} = 15$	<b>93.6</b>	5.8	0.5	0.2



Table C.2: *Recommendation Percentages for the Proposed Method using TITE Efficacy Outcomes and Bivariate Toxicity Outcomes after 1000 simulation runs with a sample size of  $n = 60$ ,  $n_R = 20$ , and  $\phi_T = 1/3$ .*

Method, TITE - Arrival	True (toxicity, efficacy) probability						Trial Duration
Scenario 1	<b>(0.02, 0.40)</b>	(0.03, 0.25)	(0.05, 0.20)	(0.06, 0.15)	(0.07, 0.10)	(0.07, 0.05)	
WT	<b>58.5</b>	17.4	11.2	7.4	4.6	0.0	180
TITE <sub>E</sub> , U-f	<b>58.8</b>	16.7	12.2	6.6	4.5	1.2	61.52
TITE <sub>E</sub> , U-P	<b>60.3</b>	14.3	12.1	8.0	4.3	1.0	61.95
TITE <sub>E</sub> , W-f	<b>60.0</b>	17.6	10.5	7.6	3.7	0.6	62.13
TITE <sub>E</sub> , W-P	<b>62.4</b>	14.3	12.0	6.4	3.4	1.5	62.32
Scenario 2	(0.02, 0.15)	<b>(0.04, 0.45)</b>	(0.07, 0.30)	(0.09, 0.25)	(0.10, 0.16)	(0.15, 0.10)	
WT	1.6	<b>69.1</b>	17.3	9.8	1.9	0.3	180
TITE <sub>E</sub> , U-f	5.7	<b>62.5</b>	18.8	9.9	2.3	0.8	61.46
TITE <sub>E</sub> , U-P	4.2	<b>65.3</b>	17.7	9.6	2.5	0.7	60.99
TITE <sub>E</sub> , W-f	4.2	<b>68.3</b>	15.8	7.9	3.3	0.5	62.09
TITE <sub>E</sub> , W-P	5.2	<b>65.8</b>	17.8	8.0	2.5	0.7	62.32
Scenario 3	(0.01, 0.10)	(0.02, 0.15)	<b>(0.03, 0.35)</b>	(0.04, 0.18)	(0.05, 0.12)	(0.06, 0.07)	
WT	5.2	12.0	<b>63.7</b>	11.5	6.1	1.5	180
TITE <sub>E</sub> , U-f	6.0	14.2	<b>60.2</b>	10.8	6.5	2.3	61.61
TITE <sub>E</sub> , U-P	6.4	9.9	<b>66.8</b>	10.9	4.5	1.5	61.70
TITE <sub>E</sub> , W-f	7.4	11.3	<b>63.3</b>	10.4	5.4	2.2	62.18
TITE <sub>E</sub> , W-P	6.6	12.4	<b>60.8</b>	10.6	7.1	2.5	62.05
Scenario 4	(0.01, 0.05)	(0.02, 0.15)	(0.04, 0.30)	<b>(0.06, 0.45)</b>	(0.08, 0.35)	(0.10, 0.30)	
WT	0.4	5.6	24.8	<b>52.6</b>	13.4	3.2	180
TITE <sub>E</sub> , U-f	0.1	5.1	22.1	<b>53.4</b>	13.8	5.5	61.44
TITE <sub>E</sub> , U-P	0.7	6.3	20.7	<b>52.5</b>	13.5	6.3	61.65
TITE <sub>E</sub> , W-f	1.2	6.1	23.2	<b>49.6</b>	15.2	4.7	62.06
TITE <sub>E</sub> , W-P	1.1	5.7	22.8	<b>49.1</b>	16.0	5.3	61.99
Scenario 5	(0.01, 0.15)	(0.02, 0.25)	(0.03, 0.33)	(0.05, 0.47)	<b>(0.06, 0.60)</b>	(0.07, 0.40)	
WT	0.9	7.5	11.4	25.4	<b>51.4</b>	3.4	180
TITE <sub>E</sub> , U-f	2.4	7.3	10.2	24.4	<b>51.9</b>	3.8	61.21
TITE <sub>E</sub> , U-P	2.4	6.8	11.7	23.6	<b>50.5</b>	5.0	61.46
TITE <sub>E</sub> , W-f	1.9	6.8	12.5	24.4	<b>50.7</b>	3.7	61.98
TITE <sub>E</sub> , W-P	2.9	7.5	10.4	26.3	<b>49.2</b>	3.7	62.42
Scenario 6	(0.01, 0.05)	(0.02, 0.15)	(0.03, 0.30)	(0.04, 0.35)	(0.05, 0.40)	<b>(0.06, 0.50)</b>	
WT	0.6	6.7	21.6	21.3	21.4	<b>28.4</b>	180
TITE <sub>E</sub> , U-f	0.6	6.3	22.4	18.3	22.6	<b>29.8</b>	61.44
TITE <sub>E</sub> , U-P	0.6	4.3	22.0	16.7	23.1	<b>33.3</b>	61.37
TITE <sub>E</sub> , W-f	0.7	4.3	19.5	20.7	23.4	<b>31.4</b>	62.09
TITE <sub>E</sub> , W-P	1.2	3.8	20.2	21.6	21.6	<b>31.6</b>	61.99
Scenario 7	(0.05, 0.01)	(0.10, 0.02)	<b>(0.16, 0.03)</b>	(0.35, 0.05)	(0.40, 0.07)	(0.50, 0.09)	
WT	2.9	7.0	<b>31.7</b>	52.8	5.6	0.0	180
TITE <sub>E</sub> , U-f	3.0	6.1	<b>36.7</b>	47.4	6.7	0.1	61.94
TITE <sub>E</sub> , U-P	2.4	7.6	<b>33.2</b>	51.0	5.8	0.0	62.11
TITE <sub>E</sub> , W-f	3.3	6.2	<b>36.7</b>	48.6	5.2	0.0	62.15
TITE <sub>E</sub> , W-P	3.3	8.8	<b>36.1</b>	46.4	5.3	0.1	62.36

Table C.3: *Recommendation Percentages for the Proposed Method using TITE Toxicity and Efficacy Outcomes after 1000 simulation runs with a sample size of  $n = 60$ ,  $n_R = 20$ , and  $\phi_T = 1/3$ .*

Method, TITE - Arrival	True (toxicity, efficacy) probability						Trial Duration
Scenario 1	(0.02, 0.40)	(0.03, 0.25)	(0.05, 0.20)	(0.06, 0.15)	(0.07, 0.10)	(0.07, 0.05)	
WT	<b>58.5</b>	17.4	11.2	7.4	4.6	0.0	180
TITE <sub>TE</sub> , U-f	<b>61.4</b>	15.9	11.7	6.4	4.0	0.6	22.26
TITE <sub>TE</sub> , U-P	<b>60.7</b>	16.3	11.9	6.2	4.4	0.5	22.25
TITE <sub>TE</sub> , W-f	<b>63.6</b>	16.7	11.6	4.8	2.8	0.5	21.47
TITE <sub>TE</sub> , W-P	<b>57.9</b>	19.7	13.0	5.6	3.4	0.4	21.53
TITE <sub>TE</sub> , Tox-U/Eff-W,f	<b>65.0</b>	17.0	9.6	5.0	2.4	1.0	21.43
TITE <sub>TE</sub> , Tox-W/Eff-U,f	<b>61.4</b>	16.1	11.5	6.9	3.0	1.1	22.23
Scenario 2	(0.02, 0.15)	(0.04, 0.45)	(0.07, 0.30)	(0.09, 0.25)	(0.10, 0.16)	(0.15, 0.10)	
WT	1.6	<b>69.1</b>	17.3	9.8	1.9	0.3	180
TITE <sub>TE</sub> , U-f	5.7	<b>65.6</b>	18.1	8.3	1.8	0.5	22.11
TITE <sub>TE</sub> , U-P	7.2	<b>62.9</b>	19.1	8.3	2.3	0.2	22.31
TITE <sub>TE</sub> , W-f	6.3	<b>60.6</b>	22.6	7.6	2.4	0.5	21.60
TITE <sub>TE</sub> , W-P	6.1	<b>60.9</b>	21.9	8.8	1.9	0.4	21.29
TITE, Tox-U/Eff-W,f	6.1	<b>64.7</b>	18.7	7.9	1.9	0.7	21.50
TITE <sub>TE</sub> , Tox-W/Eff-U,f	5.5	<b>61.8</b>	18.9	9.5	2.9	1.4	22.18
Scenario 3	(0.01, 0.10)	(0.02, 0.15)	(0.03, 0.35)	(0.04, 0.18)	(0.05, 0.12)	(0.06, 0.07)	
WT	5.2	12.0	<b>63.7</b>	11.5	6.1	1.5	180
TITE <sub>TE</sub> , U-f	7.6	10.2	<b>64.1</b>	11.3	5.2	1.6	22.30
TITE <sub>TE</sub> , U-P	5.5	10.1	<b>67.3</b>	11.8	3.7	1.6	22.38
TITE <sub>TE</sub> , W-f	7.4	14.1	<b>63.2</b>	9.5	4.2	1.6	21.40
TITE <sub>TE</sub> , W-P	7.2	14.1	<b>61.0</b>	9.6	6.5	1.6	21.52
TITE <sub>TE</sub> , Tox-U/Eff-W,f	8.2	12.6	<b>60.8</b>	10.7	6.0	1.7	21.37
TITE <sub>TE</sub> , Tox-W/Eff-U,f	6.3	10.7	<b>63.8</b>	12.5	5.2	1.5	22.33
Scenario 4	(0.01, 0.05)	(0.02, 0.15)	(0.04, 0.30)	(0.06, 0.45)	(0.08, 0.35)	(0.10, 0.30)	
WT	0.4	5.6	24.8	<b>52.6</b>	13.4	3.2	180
TITE <sub>TE</sub> , U-f	0.6	5.1	26.1	<b>48.4</b>	15.2	4.6	22.18
TITE <sub>TE</sub> , U-P	1.0	4.7	25.1	<b>48.8</b>	15.7	4.7	22.06
TITE <sub>TE</sub> , W-f	0.3	5.6	19.6	<b>50.5</b>	17.6	6.4	21.55
TITE <sub>TE</sub> , W-P	0.8	5.0	21.8	<b>49.7</b>	18.2	4.5	21.59
TITE, Tox-U/Eff-W,f	0.8	5.7	20.4	<b>52.7</b>	14.9	5.5	21.55
TITE <sub>TE</sub> , Tox-W/Eff-U,f	0.8	5.8	24.0	<b>49.6</b>	14.8	5.0	22.19
Scenario 5	(0.01, 0.15)	(0.02, 0.25)	(0.03, 0.33)	(0.05, 0.47)	(0.06, 0.60)	(0.07, 0.40)	
WT	0.9	7.5	11.4	25.4	<b>51.4</b>	3.4	180
TITE <sub>TE</sub> , U-f	2.6	9.1	12.0	24.0	<b>47.3</b>	5.0	22.02
TITE <sub>TE</sub> , U-P	2.4	7.5	11.9	24.5	<b>49.6</b>	4.1	21.96
TITE <sub>TE</sub> , W-f	1.5	5.7	10.4	27.1	<b>50.0</b>	5.3	21.60
TITE <sub>TE</sub> , W-P	2.8	7.4	9.6	29.0	<b>46.5</b>	4.7	21.74
TITE <sub>TE</sub> , Tox-U/Eff-W,f	2.8	6.6	11.8	26.5	<b>46.8</b>	5.5	21.65
TITE <sub>TE</sub> , Tox-W/Eff-U,f	3.1	7.3	10.7	24.8	<b>48.7</b>	5.4	22.00
Scenario 6	(0.01, 0.05)	(0.02, 0.15)	(0.03, 0.30)	(0.04, 0.35)	(0.05, 0.40)	(0.06, 0.50)	
WT	0.6	6.7	21.6	21.3	21.4	<b>28.4</b>	180
TITE <sub>TE</sub> , U-f	0.4	5.4	21.1	20.4	19.4	<b>33.3</b>	22.14
TITE <sub>TE</sub> , U-P	1.2	6.1	21.7	21.3	19.3	<b>30.4</b>	22.29
TITE <sub>TE</sub> , W-f	0.9	5.1	19.2	17.7	21.9	<b>35.2</b>	21.62
TITE <sub>TE</sub> , W-P	0.5	4.4	19.8	20.5	20.8	<b>34.0</b>	21.48
TITE <sub>TE</sub> , Tox-U/Eff-W,f	0.7	4.7	21.0	21.6	19.5	<b>32.5</b>	21.52
TITE <sub>TE</sub> , Tox-W/Eff-U,f	0.5	4.4	21.9	18.7	19.6	<b>34.9</b>	22.16
Scenario 7	(0.05, 0.01)	(0.10, 0.02)	(0.16, 0.03)	(0.35, 0.05)	(0.40, 0.07)	(0.50, 0.09)	
WT	2.9	7.0	<b>31.7</b>	52.8	5.6	0.0	180
TITE <sub>TE</sub> , U-f	2.8	7.3	<b>34.1</b>	50.8	4.9	0.1	22.61
TITE <sub>TE</sub> , U-P	4.5	8.3	<b>39.3</b>	42.9	5.0	0.0	22.69
TITE <sub>TE</sub> , W-f	1.6	6.6	<b>35.1</b>	51.9	4.7	0.1	20.91
TITE <sub>TE</sub> , W-P	3.6	7.0	<b>37.7</b>	46.6	5.0	0.1	20.95
TITE <sub>TE</sub> , Tox-U/Eff-W,f	4.3	6.2	<b>36.7</b>	46.6	6.2	0.0	20.87
TITE, Tox-W/Eff-U,f	2.3	6.5	<b>37.8</b>	47.3	5.8	0.3	22.60

# Bibliography

- [1] Akaike, H. (1973), Information theory and an extension of the maximum likelihood principle. In B.N. Petrov and F. Csaki (Eds.), *2nd International Symposium on Information Theory* (pp.267-281). Akad. Kiado, Budapest.
- [2] Bartroff, J., Lai, T. L., and Shih, M. (2013), *Sequential Experimentation in Clinical Trials*. New York: Springer.
- [3] Bickel, P. and Yahav, J. A. (1966), The asymptotic theory of Bayes solutions. Technical Report, Department of Statistics, Stanford University.
- [4] Bozdogan, H. (1987), Model selection and Akaike's information criterion (AIC): the general theory and its analytical extensions. *Psychometrika* **52**, 345-370.
- [5] Braun, T. (2002), The bivariate continual reassessment method: Extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials* **23**, 240-256.
- [6] Braun, T. (2006), Generalizing the TITE-CRM to adapt for early- and late-onset toxicities. *Statistics in Medicine* **25**, 2071-2083.
- [7] Burnham, K. P. and Anderson, D. R. (2002), *Model Selection and inference: a practical information - theoretic approach, 2nd ed.* New York: Springer-Verlag.

- [8] Cheung, Y. K. (2011), *Dose Finding by the Continual Reassessment Method*. New York: Chapman & Hall/CRC Press.
- [9] Cheung, Y. K. and Chappell, R. (2000), Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56**, 1177-1182.
- [10] Cheung, Y. K. and Chappell, R. (2002), A Simple Technique to Evaluate Model Sensitivity in the Continual Reassessment Method. *Biometrics* **58**, 671-674.
- [11] Chevret, S. (2006), *Statistical Methods in Dose-Finding Experiments*, New York: Wiley.
- [12] Conaway, M. R., Dunbar, S., and Peddada, S. (2004), Designs for single- or multiple-agent phase I trials. *Biometrics* **60**, 661-669.
- [13] Daimon, S., Zohar, S., and O'Quigley, J. (2011), Posterior maximization and averaging for Bayesian working model choice in the continual reassessment method. *Statistics in Medicine* **30**, 1563-1573.
- [14] Doob, J. L. (1949), Applications of the theory of martingales. *Colloque International Centre Nat. Rech. Sci. Paris*, 22-28.
- [15] Dragalin, V. and Federov, V. V. (2006), Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* **136**, 1800-1823.
- [16] Durham, S. D. and Flournoy, N. (1994), Random walks for quantile estimation, In: *Statistical decision theory and related topics, V* (West Lafayette, IN, 1992). Springer, New York: 467-476
- [17] Faries, D. (1994), Practical modifications of the continual reassessment method for phase I cancer trials. *Journal of Biopharmaceutical Statistics* **4** 147-164.

- [18] Federov, V. V. and Leonov, S. L. (2013), *Optimal design for nonlinear response models*. New York: Chapman and Hall/CRC Press.
- [19] Garrett-Mayer, E. (2006), The continual reassessment method for dose finding studies: a tutorial. *Clinical Trials* **3**, 57-71.
- [20] Gerber, D. E., Stopeck, A. T., Wong, L., Rosen, L. S., Thorpe, P. E., Shan, J. S., and Ibrahim, N. K. (2011), Phase I safety and pharmacokinetic study of bavituximab, a chimeric phosphatidylserine-targeting monoclonal antibody, in patients with advanced solid tumors. *Clinical Cancer Research* **17**, 6888-6896.
- [21] Goodman, S. N., Zahurak, M. L., and Piantadosi, S. (1995), Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine* **12**, 1149-1161.
- [22] Guo, R. and Speckman, P. (2009), Bayes factor consistency in linear models. In The 2009 International Workshop on Objective Bayes Methodology in Philadelphia, PA, June 5-9, 2009.
- [23] Heyd, J. M. and Carlin, B. P. (1998), Adaptive design improvements in the continual reassessment method for phase I studies. *Statistics in Medicine* **18**, 1307-1321.
- [24] Hoering, A., LeBlanc, M., and Crowley, J. (2011), Seamless phase I/II trial design for assessing toxicity and efficacy for targeted agents. *Clinical Cancer Research* **17**, 640-46.
- [25] Hoering, A., Mitchell, A., LeBlanc, M., and Crowley, J. (2013), Early phase trial design for assessing several dose levels for toxicity and efficacy for targeted agents. *Clinical Trials* **0**, 1-8.

- [26] Hunsberger, S., Rubinstein, L., Dancey, J., and Korn, E. (2005), Dose Escalation trial designs based on a molecularly targeted endpoint. *Statistics in Medicine* **24**, 2171-2181.
- [27] Hwang, J. and Peddada, S. (1994), Confidence interval estimation subject to order restrictions. *Annals of Statistics* **22**, 67-93.
- [28] Iasonos, A., Wilton, A., Riedel, E., Venkatraman, S., and Spriggs, D. (2008), A comprehensive comparison of the continual reassessment method to the standard 3 + 3 dose escalation scheme in Phase I dose-finding studies. *Clinical Trials* **5**, 465-477.
- [29] Jain, R. K., Lee, J. J., Hong, D., Markman, M., Gong, J., Naing, A., Wheeler, J., and Kurzrock, R. (2010), Phase I oncology studies: evidence that in the era of targeted therapies on lower doses do not far worse. *Clinical Cancer Research* **16**, 1289-1297.
- [30] Kapur, J. N., and Kesavan, H. K. (1992), *Entropy optimization principles with applications*. Academic Press, London.
- [31] Kass, R. E., and Wasserman, L. (1995), A Reference Bayesian Test for Nested Hypotheses and Its Relationship to the Schwarz Criterion. *Journal of the American Statistical Association* **90** 928-934.
- [32] Korn, E. L. (2004), Nontoxicity endpoints in phase I trial designs for targeted, noncytotoxic agents. *Journal of National Cancer Institute* **96**, 977-978.
- [33] Korn, E. L., Arbuck, S. G., Pluda, J. M., Simon, R., Kaplan, R. S., and Christian, M. C. (2001), Clinical trial designs for cytostatic agents: are new approaches needed? *Journal of Clinical Oncology* **19**(12), 265-272.

- [34] Korn, E. L., Midthune, D., Chen, T. T., Rubinstein, L. V., Christian, M. C., and Simon, R. M. (1994), A comparison of two phase I trial designs. *Statistics in Medicine* **13**, 1799-1806.
- [35] Lee, S. M. and Cheung, Y. K. (2009), Model calibration in the continual reassessment method. *Clinical Trials*, **6**, 227-238.
- [36] LoRusso P. M., Boerner S. A., and Seymour, L. (2010), An overview of the optimal planning, design, and conduct of phase I studies of new therapeutics. *Clinical Cancer Research* **16**, 1710-1718.
- [37] O’Quigley, J. (2002), Continual reassessment designs with early termination. *Biostatistics* **3**(1), 87-99.
- [38] O’Quigley, J. (2006), Theoretical study of the continual reassessment method. *Journal of Statistical Planning and Inference* **136**, 1765-1780.
- [39] O’Quigley, J., Hughes, M. D., and Fenton, T. (2001), Dose-Finding Designs for HIV Studies. *Biometrics* **57**, 1018-1029.
- [40] O’Quigley, J., Pepe, M. and Fisher, L. (1990), Continual reassessment method: a practical design for phase I clinical trials. *Biometrics* **46**, 33-48.
- [41] O’Quigley, J. and Reiner, E. (1998), A stopping rule for the continual reassessment method. *Biometrika* **85**, 741-748.
- [42] O’Quigley, J. and Shen, Z. L. (1996), Continual reassessment method: a likelihood approach. *Biometrics* **52**, 673-684.
- [43] O’Quigley, J. and Zohar, S. (2010), Retrospective robustness of the continual reassessment method. *Journal of Biopharmaceutical Statistics*, **20**, 1013-1025.

- [44] Parulekar, W. R. and Eisenahuer, E. A. (2002), Novel endpoints and design of early clinical trials. *Annals of Oncology* **13**(Supplement 4), 139-143.
- [45] R Core Team. (2013), R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing. Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.R-project.org>.
- [46] Shen, L. Z. and O'Quigley, J. (1996), Consistency of continual reassessment method under model misspecification. *Biometrika* **83**(2), 395-405.
- [47] Shiryaev, A. N. (1984), *Probability*. (Trans. R.P. Boas) New York: Springer-Verlag.
- [48] Storer, B. E. (1989), Design and analysis of Phase I clinical trials. *Biometrics* **45**, 925-937.
- [49] Thall, P. F. and Cook, J. D. (2004), Dose-Finding Based on Efficacy-Toxicity Trade-Offs. *Biometrics* **60**, 684-693.
- [50] Thall, P. F. and Nguyen, H. Q. (2012), Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *Journal of Biopharmaceutical Statistics* **22**, 785-801.
- [51] Thall, P. F. and Russell, K. T. (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.
- [52] Wages, N. A., Conaway, M. R., and O'Quigley, J. (2011), Continual Reassessment Method for Partial Ordering. *Biometrics* **67**, 1555-1563.
- [53] Wasserman, L. (2000), Bayesian Model Selection and Model Averaging. *Journal of Mathematical Psychology* **44**, 92-107.



- [54] Whitehead, J., Zhou, Y., Stevens, J., Blakey, G., Price, J., and Leadbetter, J. (2006), Bayesian decision procedures for dose-escalation based on evidence of undesirable events and therapeutic benefit. *Statistics in Medicine* **25**, 37-53.
- [55] Yin, G. (2012), *Clinical Trial Design: Bayesian and Frequentist Adaptive Methods*. Hoboken, New Jersey: John Wiley & Sons.
- [56] Yin, G., Li, Y., and Ji, Y. (2006), Bayesian dose finding in phase I/II clinical trials using toxicity and efficacy odds ratio. *Biometrics* **62**, 777-787.
- [57] Yin, G. and Yuan, Y. (2009), Bayesian Model Averaging Continual Reassessment Method in Phase I Clinical Trials. *Journal of the American Statistical Association* **104**, 954-968.
- [58] Yin, G., Zheng, S., and Xu, J. (2013), Two-stage dose finding for cytostatic agents in phase I oncology trials. *Statistics in Medicine* **32**, 644-660.
- [59] Yuan, Y. and Yin, G. (2009), Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society Series C* **58**(5), 719-736.
- [60] Yuan, Y. and Yin, G. (2011), Bayesian hybrid dose-finding design in phase I oncology clinical trials. *Statistics in Medicine* **30**, 2098-2108.
- [61] Zang, Y., Lee, J., and Yuan, Y. (2014), Phase I Dose-finding Trial Designs for Identifying Optimal Biological Dose for Molecularly Targeted Agents. *Clinical Trials* **11**(3), 319-327.