Statistical Methods for Bayesian Adaptive Early-Phase Clinical Trial Designs

by

Jin Zhang

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biostatistics)
in The University of Michigan
2013

Doctoral Committee:

Associate Professor Thomas M. Braun, Chair
Assistant Professor Veronica Berrocal
Professor John E. Levine
Associate Professor Susan Murray

To my family & my advisor

ACKNOWLEDGEMENTS

I am eternally grateful to my thesis advisor Dr. Thomas M. Braun for his unreserved guidance and support during my Ph.D. study and research and for his patience, motivation, enthusiasm, and immense knowledge. I am also thankful for his funding support through his National Institute of Health R01 grant (grant no. 5R01CA148713). Without his encouragement, help and support, I would never have been able to finish my dissertation. I could not have imagined having a better advisor for my research.

I would like to thank the rest of my thesis committee members, Dr. Veronica Berrocal, Dr. John E. Levine and Dr. Susan Murray, for their critical input and valuable suggestions on my dissertation; I enjoyed discussing my research with them. I also want to thank Dr. Min Zhang for serving as my academic advisor while a master's student. She encouraged me to pursue a doctoral degree and offered numerous academic advice.

I am extremely indebted to the Biostatistics Department in providing teaching assistantships for me when I was a master's student in the department. I also greatly appreciate the University of Michigan Rackham Graduate School for providing me a One-Term Dissertation Fellowship, which permitted me to work full-time on the final stages of my dissertation.

I thank Nicole Fenech, Fatma-Zohra Nedjari, Wendy Mashburn, Amanda Meyer and other staff members in the Biostatistics Department. Without their help, my stay in the department would not have been so smooth. Also, I would like to thank

the library staff for their support.

Finally, I want to express my deepest gratitude to my family who encouraged me throughout all the difficulties I experienced. Without them, I could not have completed my work.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

ABSTRACT

Statistical Methods for Bayesian Adaptive Early-Phase Clinical Trial Designs

by

Jin Zhang

Chair: Thomas M. Braun

This dissertation develops new methods for unaddressed issues in the design of Bayesian adaptive Phase I and Phase I/II oncology clinical trials, which are trials that seek to identify the optimal dose and/or schedule of a new cytotoxic agent in a small group of patients either based on dose-limiting toxicity (DLT) alone or both toxicity and efficacy.

Our first project focuses on methods to calibrate the prior variance assumed for the parameter in the Continual Reassessment Method (CRM). We propose three systematic approaches to adaptively calibrate the prior variance continually throughout the trial and compare those approaches to existing methods that calibrate the variance only at the beginning of a trial. Computer simulations show that our approaches have the ability to perform better than the existing methods under various scenarios.

In our second project, we extend the traditional Phase I dose-schedule-finding design that only optimizes dose and schedule among patients by adaptively re-evaluating and, if necessary, varying the intra-patient dose-schedule assignment as the study proceeds. Our design is based on a Bayesian non-mixture cure rate model that incorporates multiple administrations each patient receives with the per-administration

dose included as a covariate. Simulations indicate that our design identifies correct dose and schedule combinations as well as the traditional method that does not allow for intra-patient doses-schedule reassignments, but with a larger number of patients assigned to those combinations. The method is illustrated by application to a bone marrow transplantation trial for acute myelogenous leukemia (AML).

In our third project, we generalize our method in the second project by jointly modeling toxicity and efficacy as time-to-event outcomes in a Phase I/II clinical trial. We adopt a non-mixture cure rate model for the marginal distributions. A copula is then assumed to obtain a bivariate time-to-event distribution. To ensure an ethical trial, dose-schedule regimes are selected for successive patient cohorts based on the proposed safety and efficacy acceptability criteria at each decision-making time. Through simulations we show that the proposed design has a high probability of making correct decisions and treats most patients at desirable treatment regimes.

CHAPTER I

Introduction

1.1 Background

Early-phase clinical trials are first-in-human studies for a new agent. This thesis focuses on the development of novel statistical methods for Bayesian adaptive dose- and/or schedule-finding designs in Phase I and Phase I/II oncology trials.

Anti-cancer drugs are naturally toxic in order to kill cancer cells or suppress their growth. Thus, the study subjects in oncology trials are cancer patients, instead of healthy volunteers that might be used in other therapeutic areas. In oncology trials, especially for cytotoxic drugs, the general belief of oncologists is that the more toxic a regimen (a single drug, combined drugs or a dose-schedule combination) is, the more efficacious it will be. In addition, it is often reasonable to assume that the higher the dose is, the more toxic the regimen should be. In a typical early-phase oncology trial, the clinician suggests several doses and/or schedules regimes for investigation and tries to identify the optimal regimen. Hence, the number of regimes under study is actually finite rather than lying along continuum.

In Phase I oncology trials, the investigators are interested in determining a treatment regimen that is not only safe but also likely to be efficacious in a small group of patients, typically ranging from 15-40 patients. A reasonable approach is to specify a target toxicity rate η that should be sufficiently low to indicate safety and sufficiently

high to indicate efficacy, where η usually falls in the interval $[0.2, 0.4]$. Toxicity here does not include all types of adverse events but is a dose-limiting toxicity (DLT), which, even though varying among trials , often includes Grade 3 or higher toxicity according to the National Cancer Institute. The dose with dose-limiting toxicity rate closest to η is usually defined as the maximum tolerated dose (MTD), although a penalty could be imposed to select a lower dose for over-dose control and trial safety. One of the main goals of Phase I clinical trials in oncology is to establish the MTD that will be examined further in a Phase II trial for efficacy.

A limitation of a typical Phase I trial design is that efficacy is ignored and dose finding is based on toxicity alone, although efficacy information is still collected. Ignoring efficacy outcomes might result in an inefficient design that tends to target a suboptimal regimen or requires a larger sample size. One obvious reason is that the MTD might not be estimated in a reliable way in a Phase I trial due to the small sample size. Second, for cytostatic agents, efficacy effect may not necessarily increase with the dose. Therefore, it might be a better strategy to incorporate both toxicity and efficacy in a seamless phase I/II design. Third, for trials that aim to identify optimal dose combinations or dose-schedule combinations, there would be no obvious ordering in terms of toxicity or efficacy for the combinations in the dose-dose or dose-schedule matrix. It is very likely that two regimes could be very similar in the DLT rate but differs substantially in response rates. Hence, it would be desirable to consider both toxicity and efficacy for dose and/or schedule finding, necessitating a seamless Phase I/II trial design. In general, the main goal of a Phase I/II clinical trial is to identify a treatment regimen that is both safe and efficacious by imposing some conditions on the toxicity and response rate when selecting the best regimen for a new cohort of patients.

Since anti-cancer agents are quite toxic, a Phase I or Phase I/II trial should (1) minimize the number of patients treated at sub-optimal regimes, including those

overly toxic regimes and those with unacceptably low toxicity rates or low response rates; (2) stop a Phase I trial early if it is very likely that no acceptably safe regimen and stop a Phase I/II trial early if there is sufficient evidence that no regimen under investigation is both safe and efficacious. To meet the above ethical constraints, patients are rarely randomized to all the investigational treatment regimes. Instead, patients are enrolled in cohorts and the assignment for the next cohort is adaptively determined based on the observed accumulating data from the enrolled patients. This procedure will be repeated until the maximum number of patients is reached or the trial is stopped early.

The focus of this thesis is on model-based methods and the optimal regimen will be determined based on the proposed dose-toxicity and/or dose-response model. We adopt a Bayesian approach because early-phase trials have relatively small sample sizes and a Bayesian approach is useful for the decisions to make at each interim analysis. For example, the Bayesian Continual Reassessment Method (CRM) in Phase I dose-finding trials assumes a single-parameter dose-toxicity model, updates the model parameter whenever a new cohort is enrolled, and selects the current MTD, which is defined to the one with estimated DLT rate close to the target η, to assign to the next cohort (o'Quigley, Pepe, and Fisher 1990).

Due to the ethical constraints mentioned above, it would be ideal that we assign most of our patients at or around the optimal treatment regimen. As a result, the trial is typically severely unbalanced, which might cause problems in estimating the DLT or response rates of those regimes that very few patients are assigned to. However, since the main goal is to identify the optimal regimen instead of finding a good model, minimizing the bias/variance of the model parameter estimate is not quite relevant. Instead, we are more interested in the local model fit for the dose-toxicity/dose-response curve around the optimal regimen. In light of the above design features and the small sample size in an early phase trial, a parsimonious model is usually

preferred to a complex model. For example, the CRM assumes a single-parameter model that has been shown to perform well in numerous simulation studies and a two-parameter model does not improve the operating characteristics(Shu and O'Quigley 2008; Paoletti and Kramar 2009).

1.2 Motivation and Significance

We develop methods for three unaddressed issues in Bayesian adaptive Phase I and Phase I/II trial designs. The details will be discussed in the ensuing chapters. In Chapter II, we focus on methods to adaptively calibrate the prior variance assumed for the parameter in the Phase I Bayesian CRM in order to improve its operating characteristics. Due to the small sample size in a Phase I study, the CRM can be sensitive to the amount of the prior variance. Although methods have emerged to adaptively select skeletons and to calibrate the prior variance only at the beginning of a trial, there has not been any approach developed to adaptively calibrate the prior variance throughout a trial. We propose three systematic approaches to adaptively calibrate the variance of the prior distribution during a trial in the CRM and compare those approaches via simulation to existing methods that calibrate the variance only at the beginning of a trial (Zhang, Braun, and Taylor 2012).

In Chapter III, we develop methods for dose-schedule-finding designs in which we generalize the methods of Liu and Braun (2009) to simultaneously optimize the dose and schedule assigned to each patient. A second limitation of existing model-based adaptive Phase I designs is that they only determine the assignment for the next patient or group of patients by using the most recent model estimates. What these designs fail to do is to re-examine the assignments of patients who are still receiving treatment and may benefit from a change to their assignment, such as a higher dose at the next administration or increasing the number of planned administrations at the current dose. It is important to introduce intra-patient dose and/or schedule

reassignment when necessary, especially for the patients enrolled early in a trial since they are more likely to receive a suboptimal dose or schedule. Although the model of Braun et al. (2007) could allow the possibility that the patients planned dose for each administration to be changed, both the benefit of intra-patient dose change and how to reassign intra-patient dose and/or schedules are areas that have not been studied. The second contribution of our work is to adaptively optimize the dose and schedule assignments both among patients and within patients.

In Chapter IV, we generalize the method of Phase I dose-schedule-finding in Chapter III to allow for both toxicity and efficacy outcomes by modeling times to toxicity and efficacy outcomes jointly. There are several challenges for designing such a trial. First, the ordering information in terms of either toxicity or response rates is not completely known for two dose-schedule regimes, which excludes the use of 3+3 and other traditional escalation/de-escalation methods that rely on the ordering information. Secondly, patient responses usually take a relatively long time to assess and an efficient design should be able to incorporate incomplete follow-up for the outcomes. However, most of current Phase I/II designs require complete follow-up for each enrolled patience because they model binary outcomes instead of time-to-event outcomes. As a result, the total duration of a trial might be overly long. Third, the definition of the schedule can vary between studies. A typical schedule can be either a nested schedule or a non-nested schedule. However, there has not been any systematic approach to handle both of them. Lastly, intra-patient dose modification is common in practice but most of current Phase I/II designs fail to accommodate intra-patient dose variation. In order to address all the issues above, we adopt the marginal survival functions using the non-mixture cure rate models that incorporate the cumulative effect due to multiple administration with each per-administration dose as a covariate. A copula is then assumed to obtain a bivariate time-to-event distribution. To ensure an ethical trial, adaptive safety and efficacy acceptability

conditions are imposed on the dose-schedule regimes.

In Chapter V, we present a brief summary of our proposed work, and discuss future research areas that can be explored further based on our current work.

CHAPTER II

Adaptive Prior Variance Calibration in the Bayesian Continual Reassessment Method

2.1 Introduction

Phase I clinical trials are studies of human subjects aimed at estimating the maximum-tolerated dose (MTD) with the sample size typically in the range of 20-40 subjects. The MTD is the dose at which the probability of having a dose-limiting toxicity (DLT) is near a predefined target $0 < \eta < 1$. Since the dose identified as the MTD will be further investigated for efficacy in Phase II trials, it is important to obtain an accurate estimate for the MTD. Due to the severity of most DLTs, patient safety dictates that the study begins at low doses and escalates doses as patients are accrued so that exposure of patients to doses above the MTD is minimized. However, escalation of doses should also occur as quickly as possible as lower doses are also expected to be ineffective for treating or preventing recurrence of cancer.

A vast amount of methodology exists for the design of Phase I trials. The 3 + 3 design is the standard algorithmic design using cohorts of three patients. While algorithmic designs are simple to understand and implement, their resulting MTD estimates have large bias and variance. Also, many subjects are likely to be treated at doses below the MTD (Storer 1989; Rosenberger and Haines 2009).

A preferred design would incorporate a parametric model for the association of dose and probability of DLT. One popular model-based method is the Continual Reassessment Method (CRM) which provides the MTD estimate from a fixed set of dose levels using a one-parameter model for the dose-toxicity relationship. The parameter estimate is updated every time a new subject or cohort completes its follow-up either using Bayesian methods as proposed by O'Quigley, Pepe, and Fisher (1990) or maximum likelihood methods as proposed by O'Quigley and Shen (1996).

In the Bayesian CRM, one must determine *a priori* DLT rates for each dose, referred to as a skeleton, and the first subject is assigned to the dose whose skeleton value is closest to η. Faries (1994), Korn et al. (1994) and Moller (1995) proposed modifications to the original CRM to promote patient safety and slow dose escalation. Specifically, the modified CRM suggests that the first patient be assigned to the lowest dose, regardless of the skeleton, and that skipping of doses during dose escalation should not be allowed. Numerous extensions to the CRM have been published since the original CRM manuscript, including the time-to-event CRM (TITE-CRM) of Cheung and Chappell (2000) to account for incomplete follow-up of patients and the later generalization of Braun (2005) to adapt the TITE-CRM for early- and late-onset DLTs. Yin and Yuan (2009) proposed the Bayesian Model Averaging CRM (BMA-CRM) to allow for the incorporation of multiple skeletons, and Yuan and Yin (2011) developed a hybrid design to combine rule-based methods and the CRM. Lee and Cheung (2009; 2011) suggested a systematic but computationally intensive approach to calibrate the skeleton through the use of indifference intervals.

The work of Lee and Cheung (2011) also proposed methods to determine the value of the variance given to the prior distribution of the parameter in the dose-toxicity model at the onset of the trial. In general Bayesian applications, a large, i.e. vague, prior variance usually connotes a less-influential prior distribution, and Chevret (1993) suggested using a vague prior variance with the Bayesian CRM, although the

specific definition of vagueness is controversial. Lee and Cheung (2011) proposed a least-informative prior variance, defined as that value of the prior variance that results in all doses being *a priori* equally likely of being the MTD. The value of the least-informative prior variance tends to be much smaller than what is traditionally considered a vague prior variance.

It is also not appreciated that the level of vagueness of the prior variance is dependent upon the values selected for the skeleton. The aggressive behavior of the CRM in the case studies of Moller (1995) and Neuenschwander et al. (2008) can be entirely explained by the dependence between the prior variance and the skeleton, so that the prior variance used in each study was too small for the chosen skeleton. As a specific example, O'Quigley et al. (1990) suggested using a standard exponential distribution. Consider two skeletons for five dose levels: the original skeleton used by O'Quigley et al. (1990) and a skeleton developed using the methods of Lee and Cheung (2011). Both skeletons specify the third dose as the MTD. The target probability is 0.20, the true MTD is dose 6, the maximum number of enrolled patients is 25 and the dose-response model is the hyperbolic model defined by O'Quigley et al. (1990). From the results presented in Scenario 1 in Table 2.1, we see a notable difference in the dose selected as the MTD under these two skeletons even though both use the same prior distribution for the parameter. The prior distribution works well with the second skeleton but may be too small for the first skeleton. If we increase the prior variance by using a Gamma distribution with mean 1 and variance 4, the first skeleton now gives results comparable to the second skeleton. Hence, the vagueness of a prior variance heavily depends on the skeleton used.

For a specific skeleton used in a trial, the choice of the prior variance also depends on the relative location of the true MTD and the MTD defined by the skeleton. If the MTD specified by the skeleton is close to the true MTD, a small prior variance could help find the correct MTD more efficiently. However, if the skeleton does not match

Table 2.1: A simulation study comparing the impact of the prior variance and skeleton on the ability to identify the MTD for the traditional CRM with a fixed prior variance. Numbers 1-6 in the first row stand for doses 1 to 6. Other numbers in the table stand for the proportion of simulations that select each dose as the MTD. Skeleton A denotes the original skeleton used by O'Quigley *et al.*: {0.05, 0.10, **0.20**, 0.30, 0.50,0.70} and Skeleton B denotes the skeleton used by Lee and Cheung: {0.05, 0.11, **0.20**, 0.31, 0.42, 0.53}. Numbers in bold indicate which dose is the MTD.

Scen	Skeleton	Prior		1	2	3	4	5	6
1			DLT rates:	0.00	0.00	0.03	0.05	0.11	0.22
	A	$Exp(1)$		0	0	0	6	65	**29**
	B	$Exp(1)$		0	0	0	4	36	**60**
	A	$Gamma(1/4, 4)$		0	0	0	3	36	**61**
2			DLT rates:	0.02	0.05	0.10	0.20	0.30	0.50
	A	$Exp(1)$		0	2	20	**56**	23	0
	A	$Gamma(1/4, 4)$		0	2	18	**47**	31	1
	B	$Exp(1)$		0	2	22	**50**	23	2
	B	$Gamma(1/4, 4)$		0	2	20	**49**	26	3

the truth well, then a larger prior variance is needed to help find the MTD. In Table 2.1, for the same Skeleton A, we find that a larger prior variance works better when the true MTD is dose 6 and that a smaller prior variance works better when the true MTD is dose 4. For both scenarios, the *a priori* MTD is dose 3. Again the vagueness of a prior variance depends on the skeleton used. To achieve similar performance with using Skeleton A and standard exponential distribution as the prior, one might need to further reduce the prior variance used for Skeleton B. It seems that the CRM using a constant prior variance could perform well in specific scenarios but might not perform well in other scenarios, no matter which value is selected for the constant prior variance.

The above motivating example indicates that the traditional methods to calibrate the prior variance may not work well in many scenarios, since the traditional approaches try to find a prior variance based on the skeleton at the onset of a trial and keep it constant during a trial but fail to take into account the relative location of the

true MTD and *a priori* MTD for the specific scenario. However, the accumulating data during a trial might provide us information regarding the relative distance between the truth and the skeleton. Hence, we consider the prior variance as a tuning parameter that should be adaptively calibrated during the entire study to determine whether or not the variance chosen at the beginning of the study should be modified. We introduce three systematic approaches for adaptively calibrating the prior variance throughout a Phase I trial. In Section 2.2, we review the CRM and the work of Lee and Cheung (2011). In Section 2.3, we present the details for our three variance calibration approaches. In Section 2.4, we apply our methods to two hypothetical settings and compare operating characteristics with current approaches. In Section 2.5, we conclude with some discussion.

2.2 Existing Methods

2.2.1 Continual Reassessment Method

Under the assumption that the probability of DLT increases monotonically with dose, the CRM procedure updates the dose-response relationship throughout the trial as new observations are available. Patients are assigned to the dose whose estimated DLT rate is the closest to the target probability η, subject to possible restrictions. Let J denote the number of doses examined and let N denote the number of subjects enrolled by the end of the trial. For each dose j, $j = 1, \ldots, J$, there is a skeleton value p_j, denoting the *a priori* DLT rate for dose j. The response y_i of patient i is binary: $y_i = 1$, if there is DLT or $y_i = 0$, if there is no DLT, $i = 1, \ldots, N$. The CRM uses a one-parameter model given by $\pi_i = \psi(x_i; \beta)$, where β is some unknown parameter, ψ is a monotonic function with the range $[0, 1]$ and x_i denotes the rescaled value of the assigned dose for subject i. Here, we consider two commonly used models: (1) a logistic model with intercept 3 given by $\psi(x_i; \beta) = 1/\{1 + \exp[-3 - \exp(\beta)x_i]\}$ and (2) a power model given by $\psi(x_i; \beta) = x_i^{\exp(\beta)}$. In both models, we place a normal

prior on β with mean zero and variance σ^2.

The rescaling of doses attempts to mirror the investigators' prior assumptions and provides a good fit over the skeleton probabilities for the dose levels under the study (O'Quigley et al. 1990). Specifically, x_i can take one of the rescaled values x_j^* that are determined from the equations

$$p_j = \int \psi(x_j^*; \beta) g(\beta) d\beta \qquad j = 1, \ldots, J,$$

where $g(\beta)$ is the prior distribution for β. In practice, this computation is replaced with the simplified formula

$$x_j^* = \psi^{-1}_{\beta = E(\beta)}(p_j). \tag{2.1}$$

Let $\mathbf{Y}_n = \{y_1, \ldots, y_n\}$ denote the observed DLT responses for subjects $1, \ldots, n, 1 \leq n \leq N$, after subject n has completed follow-up for DLT. Then the likelihood function for $\mathbf{Y}_n$ is given by

$$L(\mathbf{Y}_n|\beta) = \prod_{i=1}^{n} \{\psi(x_i; \beta)\}^{y_i} \{1 - \psi(x_i; \beta)\}^{1-y_i}$$

By Bayes' theorem, the posterior mean of the DLT rate at dose d_j given the observed data is given by

$$\tilde{\pi}_j = E(\psi(\beta; x_i = x_j^*)|\mathbf{Y}_n) = \int \psi(\beta; x_i = x_j^*) \frac{L(\mathbf{Y}_n|\beta) g(\beta)}{\int L(\mathbf{Y}_n|\beta) g(\beta) d\beta} d\beta.$$

In practice, the plug-in estimator, $\psi(x_j^*; \tilde{\beta})$ where $\tilde{\beta} = E(\beta|\mathbf{Y}_n)$, is commonly used to simplify the calculation for $\tilde{\pi}_j$. Based on the updated posterior DLT rates $\tilde{\pi}_j$, $j = 1, \ldots, J$, the recommended dose for the next patient is chosen as the one with a DLT rate closest to the target η. So the next subject is assigned to dose level j such

that

$$j = \underset{j\in(1,\ldots,J)}{\arg\min} |\tilde{\pi}_j - \eta|. \tag{2.2}$$

The CRM usually does not allow dose skipping during dose escalation. The trial either progresses until the total number of subjects N is reached or is terminated if a certain stopping rule is satisfied. The MTD is determined at the end of the trial by simply selecting dose j according to (2.2) based upon $\mathbf{Y}_N$.

In order to address the ethical concern of overdosing subjects, many authors have developed stopping rules for dose-finding studies that halt a study if all doses under study are too toxic, including Korn et al. (1994), O'Quigley (1992), and O'Quigley and Reiner (1998). In our simulations presented in Section 2.4, we used a variant of the stopping rule proposed by Thall and Russell (1998), in which the trial is stopped and no dose is selected as the MTD once the posterior probability that the DLT rate of the lowest dose is higher than the target probability is larger than a pre-specified value.

2.2.2 Least Informative Prior Variance of Lee and Cheung (2011)

We first briefly review the concept of indifference intervals proposed by Cheung and Chapell (2002) in the context of the CRM. The parameter space of β can be divided into J intervals: $I_1 = [b_l, b_1), I_j = (b_j, b_{j+1})$ for $j = 1, \ldots, J-2$ and $I_J = (b_{J-1}, b_u)$, where $b_1, \ldots, b_{J-1}$ are solved from

$$\psi(x_j^*; b_j) + \psi(x_{j+1}^*; b_j) = 2\eta, \quad \text{for} \quad j = 1, \ldots, J-1.$$

It is obvious that the CRM would assign dose j to the next subject if and only if the estimate $\tilde{\beta}$ falls in the interval I_j, $j = 1, \ldots, J$. Although $\beta \in (-\infty, \infty)$, finite values for b_l and b_u are used in practice to avoid computational difficulty.

The least informative prior variance, denoted as σ_{LI}^2, is the prior variance that results in β being equally likely of belonging to any of the J intervals, i.e., all doses

being *a priori* equally likely of being the MTD. These J probabilities can be regarded as being from a discrete uniform distribution, although it is usually not possible to make them exactly equal. Instead, Lee and Cheung (2011) defined σ_{LI}^2 as the prior variance such that the variance of the J probabilities matches $(J^2-1)/12$, the variance of a discrete uniform distribution. Although σ_{LI}^2 is uninformative in terms of the prior model-based MTD distribution, the value of σ_{LI}^2 is usually not large with respect to what is usually considered to be an uninformative variance.

For example, in the setting where there are five dose levels, the skeleton is {0.05, 0.10, 0.20, 0.35, 0.50}, the target η is 0.20 and a logistic model with intercept 3 is used, the resulting five intervals of β in which doses 1 to 5 are the MTD are $I_1 = (-\infty, -0.23), I_2 = (-0.23, -0.08), I_3 = (-0.08, 0.10), I_4 = (0.10, 0.29)$ and $I_5 = (0.29, \infty)$, respectively. The least informative prior variance σ_{LI}^2 is 0.32^2, which would usually be regarded as an informative prior variance in general Bayesian applications.

2.3 Methods for Adaptive Variance Calibration

2.3.1 Defining a Large Prior Variance σ_{HI}^2

When the MTD defined by the skeleton is not the first or last dose, a prior variance larger than σ_{LI}^2 would result in a U-shaped distribution of the *a prior* model-based MTD (Lee and Cheung 2011). As a result, dose 1 and J would be more likely to be selected as the MTD. Hence, σ_{LI}^2 could perform poorly when the MTD is the lowest or highest dose and the MTD defined by the skeleton lies elsewhere, at least when no stopping rule is used. Therefore, we further define a larger prior variance, σ_{HI}^2, as the prior variance that satisfies $Pr(\beta \in I_1 \cup I_J) = 0.8$, producing a U-shaped distribution for the model-based MTD. Presumably, σ_{HI}^2 could perform well when σ_{LI}^2 performs poorly. A value other than 0.80 can certainly be used to determine the value of σ_{LI}^2 . However, values larger than 0.80 will place more mass in the tails of the MTD distribution and may be too aggressive in situations when the MTD is

not the highest dose. Conversely, values smaller than 0.80 will place less mass in the tails of the MTD distribution and will lessen the ability to find the MTD when it is the highest dose. We found that 0.80 was a good compromise between these two situations.

2.3.2 CRM-VC1: Increasing the Prior Variance With the Sample Size

We denote CRM-VC1 as our first approach to adaptively calibrate the prior variance in the CRM. Since the sample size is small early in a trial, it may be appropriate to use σ^2_{LI} at the beginning of a trial so that each dose is *a priori* equally likely to be selected as the MTD. However, it would not be desirable for the prior to dominate the data (Iasonos and OQuigley 2011), especially when the MTD is the lowest or highest dose. Hence, a sufficiently large prior may be preferred later in a trial. One natural approach is to start the prior variance at σ^2_{LI} and increase it to σ^2_{HI} at a rate based upon n, the number of currently enrolled patients. We have selected five different functions explaining how the prior variance increases with n so that different rates of change could be captured:

1. $\sigma_n^2 = \sigma^2_{LI} + (\sigma^2_{HI} - \sigma^2_{LI})(n-1)^4/(N-1)^4$
2. $\sigma_n^2 = \sigma^2_{LI} + (\sigma^2_{HI} - \sigma^2_{LI})(n-1)^2/(N-1)^2$
3. $\sigma_n^2 = \sigma^2_{LI} + (\sigma^2_{HI} - \sigma^2_{LI})(n-1)/(N-1)$
4. $\sigma_n^2 = \sigma^2_{LI} + (\sigma^2_{HI} - \sigma^2_{LI})\log(2n-1)/\log(2N-1)$
5. $\sigma_n^2 = \sigma^2_{LI} + 2N(\sigma^2_{HI} - \sigma^2_{LI})(n-1)/(N^2-1) - (\sigma^2_{HI} - \sigma^2_{LI})(n-1)^2/(N^2-1)$.

Figure 2.1 displays the five different patterns when $N = 30$, $\sigma_{LI} = 0.33$ and $\sigma_{HI} = 1.08$, a setting we will further explore in our simulations. These five functions represent typical variance-sample size relationships: (a) the prior variance increases slowly at first and then quickly reaches σ^2_{HI}; (b) the prior variance increases with n with a

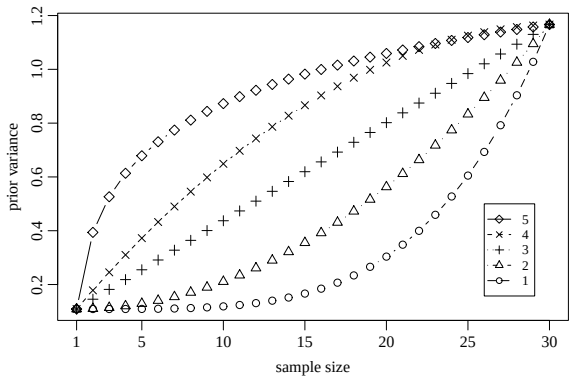

Figure 2.1: CRM-VC1: the prior variance increases with the sample size in five different patterns.

constant rate, and (c) the prior variance increases quickly at first and slowly reaches σ^2_{HI}.

2.3.3 CRM-VC2: A Hypothesis Testing Approach

If we start with a certain prior variance in the CRM, it would be ideal if the accumulating data could help determine whether the current prior variance should change. If the skeleton specifies the correct MTD, the prior variance should be small and the prior information is incorporated to enhance estimation of the MTD. Otherwise, it is preferable to change the prior variance if the data indicate that the skeleton has misidentified the MTD. This is the motivation for CRM-VC2.

A trial based on CRM-VC2 starts with the prior variance σ^2_{LI}. When the data favors the hypothesis that the MTD is the highest dose but the MTD defined by the skeleton lies elsewhere, CRM-VC2 increases the prior variance to σ^2_{HI}, since a large prior variance would increase the probability of selecting the tail dose levels due to the U-shaped distribution. We do not increase the prior variance if the MTD is dose 1, since the use of a stopping rule makes it unnecessary. However, when the MTD defined by the skeleton coincides with the highest dose, the prior variance determined by CRM-VC2 would remain at σ^2_{LI} since increasing the prior variance is no longer helpful when the prior information is correct.

The decision to switch from σ^2_{LI} to σ^2_{HI} involves a hypothesis testing approach,

similar to what Yuan and Yin (2011) proposed for their hybrid design. We propose three hypotheses: $H_1 : \beta \in I_1$, $H_2 : \beta \in I_2 \cup I_3 \ldots \cup I_{J-1}$ and $H_3 : \beta \in I_J$. We also propose two reasonable bounds b_l and b_u for β to avoid technical difficulties, that is, $\beta \in [b_l, b_u]$ rather than $(-\infty, \infty)$. Specifically, b_l satisfies $\psi(x_1, b_l) = \eta + 0.05$ and b_u satisfies $\psi(x_J, b_u) = \eta - 0.05$. Although it is guaranteed that b_l is smaller than b_u for our model parameterization, such a result may not hold true for all models, in which case one would switch b_l with b_u. Via simulation, we also examined using bounds defined by $\eta \pm 0.025$ and $\eta \pm 0.10$ and found little change in the operating characteristics when using $\eta \pm 0.05$ (results not shown). Actually, when the true β falls outside $[b_l, b_u]$, the true DLT rates for all the J doses would be far away from the target η, implying that the doses examined would be either too toxic or overly safe. A trial would hence either be terminated by a stopping rule or quickly find the highest dose as the MTD.

To be objective, we assign a uniform prior distribution under each hypothesis: $\beta|H_1 \sim \text{Unif}[b_l, b_1)$, $\beta|H_2 \sim \text{Unif}[b_1, b_J)$ and $\beta|H_3 \sim \text{Unif}[b_J, b_u]$. The marginal distribution of $\mathbf{Y}_n$ under H_1 is then given by

$$p(\mathbf{Y}_n|H_1) = \int_{b_l}^{b_1} \prod_{i=1}^{n} \{\psi(x_i;\beta)\}^{y_i} \{1 - \psi(x_i;\beta)\}^{1-y_i} \frac{1}{b_1 - b_l} \, \mathrm{d}\beta.$$

Similarly, we can compute $p(\mathbf{Y}_n|H_2)$ and $p(\mathbf{Y}_n|H_3)$. The posterior probability of H_k, $k = 1, 2, 3$, is given by

$$p(H_k|\mathbf{Y}_n) = \frac{p(H_k)p(\mathbf{Y}_n|H_k)}{p(H_1)p(\mathbf{Y}_n|H_1) + p(H_2)p(\mathbf{Y}_n|H_2) + p(H_3)p(\mathbf{Y}_n|H_3)}.$$

If we let $BF_{hk} = p(\mathbf{Y}_n|H_h)/P(\mathbf{Y}_n|H_k)$, $h = 1, 2, 3$, denote the Bayes factor for com-

paring H_h and H_k, then

$$p(H_k|\mathbf{Y}_n) = \frac{p(H_k)}{p(H_1)BF_{1k} + P(H_2)BF_{2k} + P(H_3)BF_{3k}}.$$

We specify $p(H_1) = P(H_2) = P(H_3) = 1/3$ and use Jeffreys' rule that $\log_{10}(BF_{kk'}) > 1/2$ indicates substantial evidence in favor of H_k against $H_{k'}$ (Jeffreys, 1961). This rule translates to the criterion that if $p(H_3|\mathbf{Y}_n) > 0.61$, then there is substantial evidence that $\beta \in I_J$. Once such evidence exists, the prior variance would increase to σ^2_{HI}; otherwise, the prior variance stays at σ^2_{LI}.

2.3.4 CRM-VC3: Adaptively Changing Skeletons

Instead of changing the prior variance during a trial to make the MTD more likely to be selected, CRM-VC3, our third approach to calibrate the prior variance, is to modify the skeleton adaptively but keep the prior variance constant. Consequently, the intervals $I_1, \ldots, I_J$ would also change, because the intervals $I_1, \ldots, I_J$ only depend on the skeleton and the model used. If we can properly adjust these intervals, more mass of the prior distribution could be placed over the interval that results in selecting the correct MTD.

For CRM-VC3, a trial starts with the prior variance σ^2_{LI} and once the new estimates for π_j, are obtained, the dose values would be rescaled again and used as the new dose values. Let $\tilde{\beta}_n$ denote the posterior mean of β after n subjects have finished follow-up for DLT and let $x^*_{j,0} = x^*_j$. The updated skeleton p^n_j is set equal to the current DLT probabilities based on $\tilde{\beta}_{(n)}$, i.e.,

$$p^n_j = \psi(x^*_{j,(n-1)}; \tilde{\beta}_{(n)}), \tag{2.3}$$

and, similar to (2.1), the updated rescaled dose value for dose j is,

$$x^*_{j,n} = \psi^{-1}_{\beta=E(\beta)}(p^n_j). \tag{2.4}$$

Skeletons and rescaled dose values are updated according to (2.3) and (2.4) during a trial. All other facets of the design, including the model and prior distribution, remain the same. The resulting intervals $I_1 \ldots I_J$ would change adaptively with the updating of $\tilde{\beta}$. Consider the setting where there are five dose levels, the skeleton is $\{0.05, 0.10, 0.20, 0.35, 0.50\}$, the target η is 0.20 and the model used is a logistic model with intercept 3. Figure 2.2 shows that we could assign more mass to tail areas adaptively if $\tilde{\beta}$ falls in I_1 or I_J. After the first subject is observed, if $\tilde{\beta} = 0$, the new skeleton will be the same with the original skeleton, indicating the prior information is close to the truth. As a result, the J intervals and the resulting areas under the prior density curve for the J intervals do not change since we are using σ^2_{LI} as the prior variance; see Figure 2.2 (a). This is reasonable because the data suggests that the MTD does not lie in the tail. If $\tilde{\beta} = \log(3/2)$, suggesting the MTD is dose 5. The prior density will place more mass in I_5 after rescaling the dose values with the area under $I_5 = 0.64$; see Figure 2.2 (b). If $\tilde{\beta} = \log(2/3)$, suggesting the MTD is dose 1, the prior density places more mass on I_1 after rescaling the dose values with the area under $I_1 = 0.70$; see Figure 2.2 (c). As more subjects enter the study, $\tilde{\beta}$ becomes more accurate and the resulting updated skeleton is driven by the data, avoiding the effect that a misspecified skeleton would have in the conventional CRM. One may be concerned that the updated skeleton values may be unstable early in the study when little data exists and restrict the use of CRM-VC3 after a minimum sample size has been accrued. However, restricting any skipping of doses during escalation will alleviate any possible instability. Furthermore, we examined the mean dose assigned to the first ten subjects in the settings presented in Section 2.4 (results not shown) and found that using CRM-VC3 was no more or less stable than the other methods.

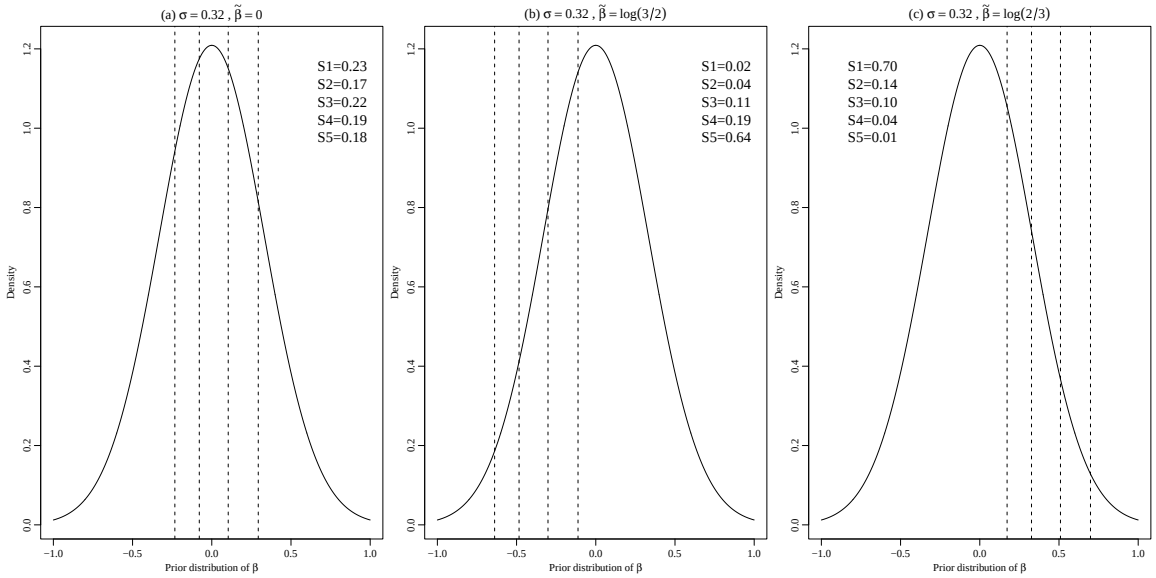

Figure 2.2: Areas under the Normal prior density curve of β with mean zero and variance $\sigma_{LI}^2 = 0.32^2$ for intervals I_1 to I_5: $S1$ to $S5$. They are also the prior probabilities for selecting each of the five doses as the MTD. Rescaling the doses sequentially could change the area under the curve for each interval. The four vertical lines stand for the boundaries for the five intervals.

2.4 Simulation Results

2.4.1 Rules Used in Simulation

We use a cohort size of one subject in our study. Like most dose-finding studies, we restrict dose escalation to be no more than one dose above the assignment of the most recent subject. However, we do not impose any restriction on dose de-escalation. Also, the first subject is always assigned to the lowest dose. However, in the simulations of the hypothetical trial of Lee and Cheung (2011), the third dose is assigned to the first subject in order to make our results comparable to theirs.

The prior variance (CRM-VC1 and CRM-VC2) or skeleton (CRM-VC3) will be updated after a new subject finishes follow-up for DLT. For patient safety, we will stop the trial if at least two out of the first three patients experience DLT, or, if $Pr(\pi_1 > \eta|\mathbf{Y}_n) > 0.9$ after four or more patients have been enrolled. In order to reduce the sensitivity to the prior variance and instability due to small sample size, our stopping rule mimics the 3+3 method for the first three subjects and then

switches to use of the posterior probability that the DLT rate for the lowest dose is above the target. The threshold of 0.9 was found to work well in simulations but could be adjusted depending upon how great the need for early stopping is. We performed $2,000$ simulations in each scenario; all simulations were done in the statistical package `R` (R Development Core Team 2008), the code for which is available upon request.

2.4.2 A Hypothetical Trial

In our hypothetical clinical trial of $N = 30$ subjects with five dose levels and the target DLT rate $\eta = 0.20$, we used the logistic model with intercept 3. The prior distribution for β was normal with mean 0 and variance σ_n^2. We considered two commonly used skeletons that specify the *a priori* MTDs to be the middle dose and the highest dose. Specifically, Skeleton 1 is $\{0.05, 0.10, 0.20, 0.35, 0.5\}$ ($\sigma_{LI} = 0.32$; $\sigma_{HI} = 1.04$) and Skeleton 2 is $\{0.01, 0.04, 0.07, 0.11, 0.20\}$ ($\sigma_{LI} = 0.35$; $\sigma_{HI} = 0.68$). We examined the performances of CRM-VC1, CRM-VC2 and CRM-VC3 in five different scenarios under both Skeleton 1 and 2 based on the percentage of simulations selected as the MTD and the average number of patients assigned to each dose. The true MTDs are doses 1 to 5 for the scenarios 1 to 5, respectively. We also performed the traditional CRM using the fixed prior variances σ_{LI}^2 and σ_{HI}^2 in the same five scenarios for comparison.

The performances under each scenario are summarized in the final eleven columns of Table 2.2 and Table 2.3 for Skeleton 1 and 2, respectively. The first five of the eleven columns display the percentage of simulations in which each dose was identified as the MTD at the end of the study, and the last five columns display the average number of patients assigned to each dose. For each scenario, we list the true toxicity probabilities in the first row, the results for the CRM using the fixed prior variance σ_{LI}^2 and σ_{HI}^2 in rows 2-3, the results obtained by CRM-VC1 using functions (1) and (2) in Section 3.2 in rows 4-5, the results obtained by CRM-VC2 in row 6 and the results obtained by CRM-VC3 in row 7. We do not present the results for CRM-VC1

using functions (3) to (5) as they did no better or worse than functions (1) and (2). We note that in general, we have found that all five functions perform similarly in terms of finding the correct MTD, indicating that there is no real need to choose among them in application.

In Table 2.2, we note that the traditional CRM is sensitive to the value of the prior variance. Using the prior variance σ^2_{LI}, a similar approach with Method A1 in Lee and Cheng (2011), performs better than using σ^2_{HI} in Scenarios 2, 3 and 4, where the true MTD is close to the MTD defined by the skeleton. But using σ^2_{LI} performs poorly relative to σ^2_{HI} in Scenario 5 where the true MTD is the highest dose but the MTD defined by the skeleton is dose 3. Overall, the CRM using σ^2_{HI} is more robust than using σ^2_{LI} in finding the MTD except when the true MTD is at or close to the skeleton MTD. We also see that using a large prior variance could produce more unnecessary early stopping. For example, in Scenario 1, using σ^2_{HI} results in 44% early stopping compared with 36% when using σ^2_{LI}. This is why the prior variance does not increase to σ^2_{HI} in CRM-VC2 when there is evidence that the true MTD is dose 1.

CRM-VC1 gives comparable results with the traditional CRM using σ^2_{LI} when the true MTD is similar to that specified by the skeleton, but performs much better when the true MTD is the highest dose. Compared with the traditional CRM using the prior variance σ^2_{HI}, CRM-VC1 performs slightly better in Scenario $1, 2$ and 3 and has comparable performance in scenarios 4 and 5. CRM-VC2 performs consistently well overall, even though it performs slightly worse than other competing methods in Scenario 4. CRM-VC2 performs as well as the CRM using the prior variance σ^2_{LI} in Scenario 1, 2 and 3, but CRM-VC2 performs much better in Scenario 5 where the true MTD is dose 5. Compared with the traditional CRM using σ^2_{HI}, CRM-VC2 also demonstrates a better ability in identifying the MTD: 61% versus 56% in Scenario 2 and 65% versus 59% in Scenario 3. CRM-VC3 performs similarly with the traditional

CRM using the prior variance σ^2_{HI}.

We also see that the design giving a higher probability of selecting the MTD also assigns more patients to the correct dose. Hence, similar results are obtained if we compare mean dose assignments among the methods examined. Overall, CRM-VC2 performs best among all the methods examined across the five scenarios. It is also common in practice that the MTD determined by the skeleton is the highest dose. Skeleton 2 is one such skeleton. When using Skeleton 2 while keeping other facets of the design unchanged, we notice that the results presented in Table 2.3 are quite similar among the five scenarios for most of the methods examined. However, CRM-VC2 and the traditional CRM using prior variance σ^2_{LI} slightly outperform other methods across the five scenarios.

2.4.3 A Hypothetical Trial in Lee and Cheung (2011)

In this setting, there are $N = 25$ subjects, six dose levels and the target DLT rate is $\eta = 0.20$. The first patient is assigned to dose 3 and the model used is the power model. In order to make our results comparable to those of Lee and Cheung, we used the skeleton $\{0.05, 0.11, 0.20, 0.31, 0.42, 0.53\}$ ($\sigma_{LI} = 0.68$; $\sigma_{HI} = 2.45$) used with Method A1 of Lee and Cheung, which we denote as LC-A1. Table 2.4 contains results obtained from the traditional CRM, LC-A1, CRM-VC1, CRM-VC2 and CRM-VC3 in the similar scenarios examined by Lee and Cheung. Note that LC-A1 is equivalent to the traditional CRM using the prior variance σ_{LI} in this example because the skeleton is the same for both methods.

In Scenario 1, all the methods work similarly. The traditional CRM using prior variance σ^2_{HI} results in more trials being terminated than using σ^2_{LI}. In Scenarios 2 and 3, LC-A1, CRM-VC2 and the traditional CRM using σ^2_{LI} perform slightly better than other approaches; all three approaches correctly identify the MTD in approximately 55% of simulations in contrast to 49-53% for the other approaches. However, LC-A1 and the traditional CRM using σ^2_{LI} perform poorly relative to CRM-VC2 in Scenario

Table 2.2: Simulation study comparing the CRM-VC1, CRM-VC2 and CRM-VC3 with the traditional CRM under Skeleton 1: {0.05, 0.10, 0.20, 0.35, 0.5}. "None" denotes the proportion of trials that stop early. Numbers 1-5 in the first row stand for dose 1 to 5. Numbers in bold indicate which dose is the MTD.

			Percentage of simulations selected as MTD						Mean number of subjects assigned				
Scen			None	1	2	3	4	5	1	2	3	4	5
1	Pr(DLT)			**0.20**	0.30	0.35	0.45	0.50					
	CRM	σ^2_{LI}	17	**44**	29	9	1	0	11	9	5	1	0
		σ^2_{HI}	32	**36**	23	8	1	0	11	6	3	1	1
	CRM-VC1	1	23	**47**	21	7	1	0	12	8	4	2	0
		2	26	**43**	24	9	1	0	11	9	5	1	0
	CRM-VC2		18	**44**	28	9	1	0	10	8	6	2	0
	CRM-VC3		34	**36**	20	8	1	0	10	6	4	1	1
2	Pr(DLT)			0.10	**0.20**	0.35	0.45	0.50					
	CRM	σ^2_{LI}	4	15	**62**	16	1	0	6	14	7	2	0
		σ^2_{HI}	6	18	**56**	19	2	0	8	12	6	2	1
	CRM-VC1	1	4	22	**57**	22	2	0	7	13	7	2	0
		2	4	20	**56**	19	1	0	7	13	7	2	0
	CRM-VC2		4	15	**61**	19	1	0	6	14	7	2	0
	CRM-VC3		9	18	**55**	17	2	0	8	12	6	2	1
3	Pr(DLT)			0.05	0.10	**0.20**	0.35	0.45					
	CRM	σ^2_{LI}	0	1	19	**66**	13	0	2	7	15	5	0
		σ^2_{HI}	2	1	19	**59**	18	1	3	7	12	6	2
	CRM-VC1	1	0	1	23	**61**	14	1	2	7	14	6	1
		2	2	1	19	**60**	18	1	2	7	13	6	1
	CRM-VC2		0	1	20	**65**	14	1	2	7	15	5	1
	CRM-VC3		2	1	20	**61**	15	1	3	7	12	6	2
4	Pr(DLT)			0.02	0.05	0.10	**0.20**	0.35					
	CRM	σ^2_{LI}	0	0	1	28	**61**	10	1	2	10	14	3
		σ^2_{HI}	0	0	1	23	**60**	16	2	3	7	12	7
	CRM-VC1	1	0	1	1	25	**58**	16	1	2	9	14	4
		2	0	0	1	21	**61**	16	1	2	8	13	5
	CRM-VC2		0	0	1	27	**56**	16	1	2	9	12	5
	CRM-VC3		0	0	1	23	**60**	15	2	3	7	12	7
5	Pr(DLT)			0.01	0.04	0.07	0.11	**0.20**					
	CRM	σ^2_{LI}	0	0	0	7	37	**56**	1	2	5	12	11
		σ^2_{HI}	0	0	0	4	27	**69**	1	2	3	7	16
	CRM-VC1	1	0	0	0	6	28	**66**	1	2	4	11	12
		2	0	0	0	4	27	**68**	1	2	4	9	14
	CRM-VC2		0	0	0	6	26	**68**	1	2	4	8	14
	CRM-VC3		0	0	0	5	25	**70**	1	2	3	7	17

Table 2.3: Simulation study comparing the CRM-VC1, CRM-VC2 and CRM-VC3 with the traditional CRM under Skeleton 2: {0.01, 0.04, 0.07, 0.11, 0.20}. "None" denotes the proportion of trials that stop early. Numbers 1-5 in the first row stand for dose 1 to 5. Numbers in bold indicate which dose is the MTD.

			Percentage of simulations selected as MTD						Mean number of subjects assigned				
Scen			None	1	2	3	4	5	1	2	3	4	5
1	Pr(DLT)			**0.20**	0.30	0.35	0.45	0.50					
	CRM	σ^2_{LI}	18	**41**	30	8	1	0	9	9	4	2	1
		σ^2_{HI}	26	**44**	24	5	1	0	13	7	3	2	1
	CRM-VC1	1	21	**46**	26	6	1	0	10	8	4	2	1
		2	22	**45**	26	8	2	0	10	9	4	2	1
	CRM-VC2		18	**41**	30	8	1	0	9	9	4	2	1
	CRM-VC3		31	**43**	20	5	1	0	12	6	3	1	1
2	Pr(DLT)			0.10	**0.20**	0.35	0.45	0.50					
	CRM	σ^2_{LI}	6	16	**60**	17	2	0	5	13	7	3	1
		σ^2_{HI}	5	20	**57**	14	2	0	8	12	5	2	1
	CRM-VC1	1	6	20	**57**	16	2	0	5	13	6	3	1
		2	6	16	**57**	17	3	1	6	13	6	3	1
	CRM-VC2		6	16	**60**	17	2	0	5	13	7	3	1
	CRM-VC3		9	21	**55**	14	2	0	9	11	5	2	1
3	Pr(DLT)			0.05	0.10	**0.20**	0.35	0.45					
	CRM	σ^2_{LI}	1	1	22	**54**	21	2	2	7	11	7	3
		σ^2_{HI}	1	2	25	**49**	22	2	3	8	10	6	3
	CRM-VC1	1	1	1	29	**49**	20	2	2	8	11	7	3
		2	2	1	26	**49**	20	2	2	8	11	7	3
	CRM-VC2		1	1	22	**54**	21	2	2	7	11	7	3
	CRM-VC3		2	2	29	**50**	17	1	3	8	10	6	3
4	Pr(DLT)			0.02	0.05	0.10	**0.20**	0.35					
	CRM	σ^2_{LI}	0	0	3	22	**58**	18	1	2	6	12	8
		σ^2_{HI}	0	0	2	21	**55**	22	2	3	6	11	9
	CRM-VC1	1	0	0	3	22	**58**	18	1	2	6	12	8
		2	0	0	2	23	**54**	21	1	2	6	12	9
	CRM-VC2		1	0	3	22	**58**	18	1	2	6	12	8
	CRM-VC3		0	0	3	24	**55**	18	2	3	7	11	8
5	Pr(DLT)			0.01	0.04	0.07	0.11	**0.20**					
	CRM	σ^2_{LI}	0	0	0	2	23	**75**	1	1	2	7	19
		σ^2_{HI}	0	0	0	2	23	**75**	1	2	2	6	18
	CRM-VC1	1	0	0	0	3	23	**73**	1	1	2	7	18
		2	0	0	0	2	22	**75**	1	1	2	6	19
	CRM-VC2		0	0	0	2	23	**75**	1	1	2	7	19
	CRM-VC3		0	0	1	4	23	**73**	1	2	3	6	18

5 where the MTD is the highest dose.

As seen earlier, the traditional CRM using σ^2_{HI} does not perform as well as σ^2_{LI} in Scenarios 2 and 3 where the true MTD is close to the MTD defined by the skeleton, even though the difference is not large. CRM-VC1 and CRM-VC2 performs as well as or better than LC-A1 in Scenarios 2,3 and 4 but performs better than LC-A1 in Scenario 5. CRM-VC3 performs similarly with the traditional CRM using σ^2_{HI}. We also see that the design giving a higher probability of selecting the MTD also assigns more patients to the correct dose. Hence, similar results are obtained if we compare mean dose assignments among the methods examined. Overall, in this setting, CRM-VC2 and CRM-VC3 seem to perform best among all the methods examined in terms of the ability of identifying the MTD across the five scenarios.

2.5 Discussion

In the present project, we relax the assumption of a fixed prior variance in the traditional CRM and propose three systematic approaches to adaptively calibrate the prior variance continually throughout the trial. Our approaches have the ability to perform better than the traditional CRM using a constant prior variance as well as methods that calibrate the prior variance only at the beginning of the trial.

Although Lee and Cheung (2011) suggested using σ^2_{LI} after first calibrating the skeleton at the beginning of the trial, this approach does not perform well when the true MTD is far away from the MTD defined by the skeleton. Although Lee and Cheung proposed an alternate, computationally intensive design, which we refer to as LC-A2, we found that LC-A2 generally offers no improvement to the results of LC-A1. Our approaches, however, are able to improve upon the results of LC-A1 in scenarios where the MTD is the highest dose without sacrificing the performance much in other scenarios, and are less computationally expensive than LC-A2. However, as seen in our simulation results, our methods might be more aggressive than the CRM

Table 2.4: Simulation study comparing the CRM-VC1, CRM-VC2 and CRM-VC3 with the method LC-A1 and the traditional CRM in scenarios examined by Lee and Cheung. The column "None" denotes the proportion of trials that stop early. The numbers 1-6 in the first row denote dose 1 to 6. Numbers in bold indicate which dose is the MTD.

			Percentage of simulations selected as MTD							Mean number of subjects assigned					
Scen			None	1	2	3	4	5	6	1	2	3	4	5	6
1	Pr(DLT)			**0.20**	0.35	0.50	0.61	0.76	0.87						
	CRM	σ_{LI}^2	37	**44**	18	1	0	0	0	8	6	3	1	0	0
		σ_{HI}^2	43	**43**	14	1	0	0	0	11	4	4	2	1	0
	LC-A1		37	**44**	18	1	0	0	0	8	6	3	1	0	0
	CRM-VC1	1	44	**43**	13	1	0	0	0	8	5	3	1	0	0
		2	45	**42**	14	0	0	0	0	8	5	3	1	0	0
	CRM-VC2		39	**44**	16	1	0	0	0	8	6	3	0	0	0
	CRM-VC3		41	**44**	14	1	0	0	0	10	4	3	1	0	0
2	Pr(DLT)			0.05	0.10	**0.20**	0.30	0.50	0.70						
	CRM	σ_{LI}^2	5	1	21	**51**	22	1	0	1	5	11	6	1	0
		σ_{HI}^2	4	3	22	**49**	22	1	0	3	5	9	5	2	1
	LC-A1		5	1	21	**51**	22	1	0	1	5	11	6	1	0
	CRM-VC1	1	4	2	22	**47**	22	2	0	1	5	11	6	1	0
		2	5	1	22	**48**	22	2	0	1	6	10	6	1	0
	CRM-VC2		4	1	20	**52**	21	2	0	1	5	11	5	1	0
	CRM-VC3		4	2	22	**48**	22	1	0	2	5	10	5	2	1
3	Pr(DLT)			0.06	0.08	0.12	**0.18**	0.40	0.71						
	CRM	σ_{LI}^2	3	0	4	27	**55**	11	0	0	2	8	10	3	0
		σ_{HI}^2	2	1	7	27	**49**	14	0	2	2	7	9	4	1
	LC-A1		3	0	4	27	**55**	11	0	0	2	8	10	3	0
	CRM-VC1	1	3	1	5	23	**53**	16	0	0	2	8	10	3	0
		2	3	0	5	25	**52**	15	0	1	2	8	10	4	0
	CRM-VC2		3	0	5	27	**54**	12	0	0	2	9	10	3	1
	CRM-VC3		2	1	6	26	**52**	14	0	1	2	8	9	4	1
4	Pr(DLT)			0.05	0.06	0.08	0.11	**0.19**	0.34						
	CRM	σ_{LI}^2	1	0	1	8	33	**46**	11	0	1	5	8	8	3
		σ_{HI}^2	1	0	2	6	23	**47**	20	1	1	3	6	8	6
	LC-A1		1	0	1	8	33	**46**	11	0	1	5	8	8	3
	CRM-VC1	1	2	0	1	7	23	**45**	22	0	1	5	7	8	4
		2	1	0	1	7	25	**45**	20	0	1	5	7	8	5
	CRM-VC2		2	0	1	8	26	**43**	21	0	1	5	7	6	5
	CRM-VC3		1	0	2	7	23	**46**	21	1	1	4	5	7	6
5	Pr(DLT)			0	0	0.03	0.05	0.11	**0.22**						
	CRM	σ_{LI}^2	0	0	0	0	7	43	**50**	0	0	3	5	9	8
		σ_{HI}^2	0	0	0	0	4	30	**65**	0	0	2	3	7	13
	LC-A1		0	0	0	0	7	43	**50**	0	0	3	5	9	8
	CRM-VC1	1	0	0	0	0	4	29	**66**	0	0	3	4	8	10
		2	0	0	0	0	4	31	**64**	0	0	3	4	7	11
	CRM-VC2		0	0	0	0	6	29	**65**	0	0	3	5	6	12
	CRM-VC3		0	0	0	0	4	28	**68**	0	0	3	3	6	13

using σ_{LI}^2 in certain scenarios. Nonetheless, the CRM-VC2 and CRM-VC3 can be modified to be less conservative by simply changing some of the thresholds used in those methods. For example, we could increase the threshold of 0.61 for the posterior probability of H_3 in CRM-VC2 to a larger value, in which case the performance would be similar to that of Lee and Cheung. However, we note that there is no one value for the threshold that will work best in all scenarios and we feel our threshold is a good choice in most scenarios.

One reviewer questioned whether or not our designs are coherent in the sense that dose escalation is possible when the most recent patient experiences a DLT (Cheung 2005). In the scenarios of Section 2.4.3, escalation after a DLT never occurred, and in the scenarios of Section 2.4.2, the dose escalation never occurred after an observed DLT in more than 3% of simulations. Thus, although there is no guarantee of coherence of our designs in all settings, any deviation from coherence is quite small and patient safety is not compromised.

Our approaches could be extended to accommodate a wider range of applications. CRM-VC1 could be easily applied to more complex studies, including finding the most successful dose or the most tolerated schedule , once we determine σ_{LI}^2 and σ_{HI}^2. CRM-VC2 and CRM-VC3 rely on the skeleton and hence could be naturally extended to a study where a skeleton is specified and the dose values are rescaled, for example modeling the toxicity in the study of finding the MSD. For a model with more than two parameters, CRM-VC2 and CRM-VC3 are still applicable even though it may be hard to find the indifference regions in high-dimensional parameter space.

CHAPTER III

A Phase I Bayesian Adaptive Design to Simultaneously Optimize Dose and Schedule Assignments Both Between and Within Patients

3.1 Introduction

Traditional Phase I trials assign a dose of a therapeutic agent to each subject and the subject receives that dose in a single administration. However, if the agent is safe at that dose, reason suggests that the patient should be given additional administrations of the agent at the same, or perhaps different, doses in hopes of maximizing any efficacy the agent may have with regard to treating or preventing disease.

Such was the motivation of the Phase I trial described by de Lima, et al. (2010). Chemotherapy is often the first treatment given to patients with acute myelogenous leukemia (AML) or advanced myelodysplastic syndrome (MDS). If chemotherapy fails to force remission of a patient's cancer, the next course of treatment is an allogeneic hematopoietic stem cell transplant (HSCT). Although short-term complete remission (CR) of cancer frequently occurs after HSCT, long-term cancer recurrence is still quite prevalent in HSCT recipients. Therefore, researchers hope to find interventions that can be given not only in proximity to HSCT, in order to promote a short-

term CR, but also repeatedly after HSCT in order to maintain a CR for a longer period of time. One such intervention is azacitidine; however, the safety profile for multiple administrations of different doses of azacitidine in AML and MDS patients was unknown, thereby necessitating the design of the Phase I trial to address this question.

At the time that the azacitidine trial was being considered, there were no published designs for simultaneous dose- and schedule-finding in Phase I trials. Although Braun, Zheng, and Thall (2004) proposed a design in which the time to toxicity was modeled using a triangular hazard model for each administration, their model assumed a single dose was under study. Liu and Braun (2009) later developed a more flexible model for the cumulative hazard of a dose-limiting toxicity (DLT) by introducing a non-mixture cure rate model and a smooth hazard function. However, their methods also assumed a single dose was being considered.

To meet the needs of the azacitadine trial, Braun et al. (2007) developed the first design for dose- and schedule-finding by generalizing the work of Braun et al. (2004) to incorporate different triangular hazard functions for each dose. However, as noted by Liu and Braun (2009), the method of Braun et al. (2004) can be inflexible mainly due to its computational difficulty, the finite support of the triangular hazard function and the difficulties with including patient-level or administration-level covariates.

The azacitidine Phase I trial was designed to identify which combination of three doses and four administration schedules was the maximum tolerated combination (MTC) of dose and schedule, defined as the combination estimated to have the probability of a DLT within 116 days of starting treatment closest to 0.30. In the design, each patient was adaptively assigned to whichever dose and schedule combination was believed to be the MTC, based upon the data collected on previously enrolled patients. One important characteristic of this trial was that once enrolled, a patient's dose and/or schedule was to remain unchanged, except for reductions in dose and/or

number of administrations due to complications unrelated to azacitidine such as infection. In the actual trial, three patients received reduction in their assigned dose, and about half of the patients had reductions to their planned number of administrations.

However, one could also envision patients who were assigned to combinations that during the trial are determined to have DLT rates well below that of the MTC. Patients assigned to such combinations who have not completed all their administrations might benefit from increases to the dose and/or number of administrations they receive. Such changes are not necessarily expected to increase correct identification of the MTC at the end of the study, but should increase the number of subjects during the trial who are assigned to combinations near the MTC. Although the model of Braun et al. (2007) could allow the possibility that the patient's planned dose for each administration to be changed, the benefit of patient reassignments, as well as how and when appropriate reassignments of doses and/or schedules are determined, are areas that have not been studied.

These issues are the motivation of our current work. First, we generalize the methods of Liu and Braun (2009) to simultaneously optimize the dose and schedule assigned to each patient. Specifically, we extend their Bayesian non-mixture cure model by incorporating the per-administration dose as a covariate for modeling the cure fraction to allow for multiple dose levels. In addition, we derive a non-mixture cure rate model through a competing risks approach to accommodate multiple administrations one patient may receive. The second contribution of our work is to adaptively optimize the dose and schedule assignments both between patients and within patients. While new patients are given the most recent maximum-tolerated dose-schedule combination (MTC) estimate, our approach also re-evaluates the estimated DLT rate for the current assignment of each enrolled patient and automatically determines whether dose-schedule reassignment is needed. Patient accrual, data monitoring, and outcome-adaptive decision-making are done continuously throughout the

trial under a Bayesian formulation. We describe the probability model and the dose-schedule-finding algorithm in Section 3.2, and we illustrate the proposed design in the context of a real trial and present a simulation study in Section 3.3. We conclude with a brief discussion in Section 3.4.

3.2 Methods

3.2.1 Preliminary Notation

Typical dose-schedule finding trials aim to find the MTC within a $J \times K$ matrix consisting of J per-administration doses and K nested schedules. We denote the administration times for schedule k, $k = 1, \ldots, K$, as $\mathbf{s}^{(k)} = \{s_1, s_2, \ldots, s_{m_k}\}$ such that $\mathbf{s}^{(1)} \subset \mathbf{s}^{(2)} \subset \cdots \subset \mathbf{s}^{(K)}$ and $m_1 < m_2 < \cdots < m_K$, where m_k is the number of administrations for schedule k. We focus on nested schedules because they have natural ordering and hence are of interest to the clinicians.

In our motivating example, there are $J = 3$ doses and $K = 4$ nested schedules. A course of administrations corresponds to daily administrations for the first 5 days followed by 24 days of rest, which we denotes as (5+, 24-). The first schedule is comprised of one single course, so the administration times $\mathbf{s}^{(1)} = \{0, 1, 2, 3, 4\}$. Schedule 2 consists of two courses with the additional course starting 28 days after the beginning of $\mathbf{s}^{(1)}$, we have $\mathbf{s}^{(2)} = \{0, 1, 2, 3, 4, 28, 29, 30, 31, 32\} = \left\{\mathbf{s}^{(1)}, \mathbf{s}^{(1)} + 28\right\}$, and so on. Ideally, we plan to give dose $j = 1, 2, \ldots J$, at each administration in $\mathbf{s}^{(k)}$, and we let d_j denote the per-administration dose. The number of subjects enrolled by the end of the trial is N and each subject will be followed up to the maximum follow-up time $\omega = 116$ days, which is determined by the clinical investigators and is a clinically meaningful duration of time that is sufficiently late enough to observe DLTs attributed to the longest schedule. A target DLT rate $\eta = 0.30$ is also elicited from clinicians and is defined as the targeted probability of cumulative toxicity by ω.

Note that d_j and $\mathbf{s}^{(k)}$ represent the combinations of doses and schedule that are

possible assignments to each patient as they enter the study. In contrast, during a trial, the actual number of administrations and the dose at each administration for each patient may differ from each of those possible combinations. To make this concept distinct, we let $\mathbf{s}_i = \{s_{i,1}, \ldots, s_{i,m_i}\}, i = 1, \ldots, N$, denote the successive times at which patient i receives the agent and let $\mathbf{d}_i = \{d_{i,1}, \ldots, d_{i,m_i}\}$ where $d_{i,l} \in \{1, \ldots, J\}$ and $l = 1, \ldots, m_i$ denote the per-administration doses for patient i at the administration times $\mathbf{s}_i$.

3.2.2 Model for Time-to-DLT After a Single Administration

As noted by Liu and Braun (2009), a significant proportion of patients are "cured," i.e. never experience DLTs after a single administration. Thus, they chose to model the time-to-DLT for a single administration using the non-mixture cure model proposed by Chen, Ibrahim, and Sinha (1999). Specifically, we take a standard cumulative distribution function $F(\nu|\boldsymbol{\phi})$ with parameters $\boldsymbol{\phi}$, with a corresponding density function $f(\nu|\boldsymbol{\phi})$, and scale $F(\nu|\boldsymbol{\phi})$ by a parameter $\theta > 0$ to create the respective survival and hazard functions $S(\nu|\theta, \boldsymbol{\phi}) = \exp[-\theta F(\nu|\boldsymbol{\phi})]$ and $g(\nu|\boldsymbol{\phi}, \theta) = \theta f(\nu|\boldsymbol{\phi})$. We adopt $S(\nu|\theta, \boldsymbol{\phi})$ as the probability of no DLT by follow-up time ν after a single administration and interpret θ as a cure rate parameter because the cure fraction $S(\infty) = \exp(-\theta)$ is determined solely by θ. The Time-to-Event CRM (TITE-CRM) of Cheung and Chappell (2000) for traditional dose-finding can be viewed as a mixture cure model for the time-to-DLT, as outlined in Braun (2005). However, we have chosen to use a non-mixture cure rate model instead of a mixture cure model because the latter does not have a proportional hazards structure and is less feasible for Bayesian computations (Chen, Ibrahim, and Sinha 1999; Tsodikov, Ibrahim, and Yakovlev 2003).

However, the non-mixture cure model used by Liu and Braun (2009) must be extended to allow the cure fraction to vary by dose. To that end, we model the cure rate fraction for administration l of patient i as $\log(\theta_{i,l}) = \beta_0 + \exp(\beta_1) d_{i,l}, -\infty < \beta_0, \beta_1 <$

∞, so that $\theta_{i,l} > 0$. Note that we exponentiate β_1 to ensure that the probability of DLT after a single administration increases with dose. As a result, the respective hazard and survival functions for a single administration l are $g(\nu_{i,l}|\boldsymbol{\beta}, \boldsymbol{\phi}) = \theta_{i,l} f(\nu_{i,l}|\boldsymbol{\phi})$ and $S(\nu_{i,l}|\boldsymbol{\beta}, \boldsymbol{\phi}) = \exp[-\theta_{i,l} F(\nu_{i,l}|\boldsymbol{\phi})]$, in which $\boldsymbol{\beta} = (\beta_0, \beta_1)$. Thus, our proposed hazard is an increasing function of dose through the cure fraction, even though the parameters in $f(\cdot)$ do not involve dose. We feel that the log-linear model should be adequate in many settings for identifying the MTC since the sample size is usually small in Phase I trials and the overall model fit is not our primary interest (O'Quigley et al. 1990). However, if one were truly concerned about the log-linear assumption, one could always add more parameters in the model if needed. Although computationally challenging relative to the small sample size, one could propose several competing models and use Bayesian Model Averaging or select the best-fitting model at each interim analysis time as outlined by Raftery, Madigan, and Hoeting (1997) and Ying and Yuan (2009).

With regard to $f(\cdot)$, we adopt the model of Liu and Braun (2009), a two-parameter Weibull density $f(\nu_{i,l}|\boldsymbol{\phi}) = \exp(-\gamma)\alpha\nu_{i,l}^{\alpha-1} \exp\left[-\nu_{i,l}^{\alpha} \exp(-\gamma)\right]$ with $\boldsymbol{\phi} = (\alpha, \gamma)$. Such a choice has biologic appeal because the resulting hazard function increases with time to a certain time point and then attenuates afterward, as was suggested by clinical investigators in the azacitidine trial. Mathematically, we expect the mode of the hazard function to exist at $\exp(\gamma/\alpha)(1 - 1/\alpha)^{1/\alpha}$, and we assume $\alpha > 1$ so that the mode exists.

We did consider modeling $\boldsymbol{\phi}$ as a function of dose, but did not because doing so would eliminate the proportional hazards structure of our model and there will be no guarantee that the hazard will increase with dose if $\boldsymbol{\phi}$ is also a function of dose. Further support for our approach is given by Chen, Ibrahim and Sinha (1999), who examined the standard cure model without proportional hazards and found both models (with or without proportional hazards) led to similar point and interval estimates.

We also ran simulations (results not shown) with $\boldsymbol{\phi}$ varying with dose and found that this added level of complexity to our model offered no benefit to identification of the MTC. The main reason for this result is Phase I trials seek to estimate well the DLT rate of the MTC and not necessarily the DLT rates for all dose-schedule combinations. Thus, overall model fit is not the primary interest and we prefer a parsimonious model with reasonable flexibility.

Thus the distribution of DLT times is controlled by the four parameters β_0, β_1, α, and γ, whose interpretations are as follows. If we denote $p_{i,l}$ as the DLT rate by the maximum follow-up time ω for a single administration l of patient i, then $p_{i,l} = 1 - S(\omega|\boldsymbol{\beta}, \boldsymbol{\phi}) = 1 - \exp\{-\exp[\beta_0 + \exp(\beta_1)d_{i,l}]F(\omega|\boldsymbol{\phi})\}$. Thus, with infinite follow-up, we have $p_{i,l} = 1 - \exp\{-\exp[\beta_0 + \exp(\beta_1)d_{i,l}]\}$, which is a complementary log-log model that could be used in the CRM with binary DLT outcomes. The intercept β_0 quantifies the limiting probability of DLT for a single administration of a dose $d_{i,\ell} = 0$, while β_1 quantifies how the limiting probability varies with dose. The rate at which the limiting probability is reached for each dose is controlled by α and γ, in which α and γ determines the mode of the DLT times and increasing the value of γ quantifies later DLT times.

3.2.3 Model for Time-to-DLT After Multiple Administrations

We employ a competing risks cure rate model by treating $y_{i,l}$, the time to the DLT after administration l of patient i, as a latent variable to incorporate the multiple administrations received by each patient. The *patient* time when patient i experiences a DLT is then defined as the random variable $Y_i = \min\{(s_1 + y_{i,1}), \ldots, (s_{m_i} + y_{i,m_i})\}$. Therefore, under the assumption of independence of $y_{i,1}, \ldots, y_{i,m_i}$, the survival function for patient i at *patient* time t, is given by

$$\psi(t|\boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}_i, \mathbf{d}_i) = Pr(Y_i > t) = \prod_{l=1}^{m_i} S(\nu_{i,l}|\boldsymbol{\beta}, \boldsymbol{\phi}) = \exp\left\{-\sum_{l=1}^{m_i} \theta_{i,l} F(\nu_{i,l}|\boldsymbol{\phi})\right\} \quad (3.1)$$

and the density function is given by

$$q(t|\boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}_i, \mathbf{d}_i) = \left[\sum_{l=1}^{m_i} \theta_{i,l} F(\nu_{i,l}|\boldsymbol{\phi})\right] \exp\left\{-\sum_{l=1}^{m_i} \theta_{i,l} F(\nu_{i,l}|\boldsymbol{\phi})\right\} \tag{3.2}$$

where $\theta_{i,l} = \exp[\beta_0 + \exp(\beta_1) d_{i,l}]$ is the cure parameter and $\nu_{i,l} = t - s_{i,l}$ is the follow-up time for administration l of patient i. The hazard function is then given by $h(t; |\boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}_i, \mathbf{d}_i) = \sum_{l=1}^{m_i} \theta_{i,l} F(\nu_{i,l}|\boldsymbol{\phi})$, which indicates the cumulative effect of multiple administrations.

The assumption that the times-to-DLT after each administration, $y_{i,1}, \ldots, y_{i,m_i}$, are independent for the m_i administrations of the same patient i might not hold, although the actual amount of correlation is not testable (Tsiatis 1975). A more general model could be based on an Archimedean copula-type model or a frailty model with a cure fraction (Hougaard 2000). For example, the above survival function could be generalized to

$$\psi^o(t|\boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}_i, \mathbf{d}_i) = \exp\left\{-\left(\sum_{l=1}^{m_i} [\theta_{i,l} F(\nu_{i,l}|\boldsymbol{\phi})]^{1/\xi}\right)^{\xi}\right\},$$

which is a Gumbel copula model with a correlation parameter ξ, in which $\xi = 0$ indicates independent DLT times. Here, we will assume independence for our model since it is simple and we feel that copula models could possibly impose strong and untestable assumptions on the correlation structure of DLT times.

Define the observed *patient* time $T_i = \min(Y_i, U_i)$ and $C_i = I(Y_i \leq U_i)$, where U_i denotes the censoring time and $I(\cdot)$ is the indicator function. Hence, we observe a DLT for patient i if $C_i = 1$. Since we perform interim analyses whenever a new patient in enrolled, by the time patient $n + 1$ is enrolled, we denote the number of patients currently in the study as n and for each enrolled patient, we observe T_i, C_i, $\mathbf{s}_i$ and $\mathbf{d}_i$, where $\mathbf{s}_i$ and $\mathbf{d}_i$, as we defined previously, are the respective time and dose for

each administration patient i have received, $i = 1, \ldots, n$. Note $U_i = \min(W_{i,n+1}, \omega)$, where $W_{i,n+1}$ the inter-patient time between patient i and $n+1$. If $n = N$, the maximum number of patients for the study, then define $W_{i,n+1}$ as the time between patient i and the end of the study. It is reasonable to assume random censoring since $W_{i,n+1}$ is usually independent of the DLT time and ω is a fixed value. Based on the above information, Equations (3.1) and (3.2), the likelihood on the data $\mathcal{D}_i = (T_i = t_i, C_i = c_i, \mathbf{s}_i, \mathbf{d}_i)$ for patient i is given by

$$\mathcal{L}(T_i = t_i, C_i = c_i | \boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}_i, \mathbf{d}_i) = \psi(t_i | \boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}_i, \mathbf{d}_i)^{1-c_i} q(t_i | \boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}_i, \mathbf{d}_i)^{c_i} \tag{3.3}$$

After determining the prior distribution $p(\boldsymbol{\beta}, \boldsymbol{\phi})$, then the posterior distribution of $(\boldsymbol{\beta}, \boldsymbol{\phi})$ based on $\mathcal{D} = \{\mathcal{D}_i : i = 1, ..., n\}$ is

$$p(\boldsymbol{\beta}, \boldsymbol{\phi} | \mathcal{D}) \propto p(\boldsymbol{\beta}, \boldsymbol{\phi}) \prod_{i=1}^{n} \mathcal{L}(T_i = t_i, C_i = c_i | \boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}_i, \mathbf{d}_i).$$

We can compute posterior quantities via adaptive Markov Chain Monte Carlo (MCMC) methods (Rosenthal 2007). Those posterior quantities will be used to identify the dose-schedule assignment for a new patient and a possibly new dose-schedule assignment for an existing patient as described in Section 3.2.5.

3.2.4 Establishing Prior Distributions

For the two parameters of the cure fraction, $\boldsymbol{\beta}$, we assign independent Gaussian distributions with prior mean and prior variance (μ_0, σ_0^2) for β_0 and (μ_1, σ_1^2) for β_1. In order to determine values for the prior means μ_0 and μ_1, we ask the investigators to provide the "skeleton" $\mathcal{P}$, which is a $J \times K$ matrix of *a priori* estimates of the DLT rates by ω for all dose-schedule combinations, in which element (j, k), denoted P_{jk}, corresponds to the combination of dose j and schedule k. We then fit the linear regression model $\log(-\log[1 - P_{jk}]) = \log(m_k) + b_0 + \exp(b_1) d_j$ and use the ordinary

least square estimates $\hat{b}_0$ and $\hat{b}_1$ as the respective values for μ_0 and μ_1.

For the two parameters of the hazard, $\boldsymbol{\phi} = (\alpha, \gamma)$, we have chosen to make α fixed to maintain a parsimonious model and limit the number of parameters to estimate. In addition, preliminary simulations (results not shown) indicated that there was no meaningful change in operating characteristics when assigning a prior distribution to α. However, because the mode of the hazard for a single administration monotonically increases with γ, estimation of γ is important to the performance our algorithm. Therefore, we assign a Gaussian prior distribution for γ with mean μ_γ and variance σ_γ^2. To determine values for α and μ_γ, we apply the method outlined in Liu and Braun (2009) for each dose and calculate the average. If the resulting value for $\alpha < 1$, we set $\alpha = 1.01$, so that the mode exists.

It is important to carefully calibrate the prior variances σ_1^2, σ_2^2 and σ_γ^2 as we can imagine that Phase I trials are usually sensitive to prior variances due to the small sample size. The prior variances should not be too small, otherwise the prior information dominates the trial. However, they cannot be too large either since we hope to incorporate the prior information for possibly more accurate estimation. We recommend calibrating the prior variances through simulations using a few different skeletons and prior variances. The prior variance that is the most insensitive to skeletons and leads to the best operating characteristics will be used for a real trial. We present an example of variance calibration related to the simulations of Section 3.3.

3.2.5 Algorithm for Adaptive Assignments for New Patients and Reassignments for Enrolled Patients

The algorithm for assigning a dose-schedule combination to a new patient is similar to that used in the CRM and many other Phase I designs. When a new patient enters the study, for every combination of a dose j and a schedule k, we compute $\hat{p}_{jk} = 1 - \hat{\psi}(\omega|\boldsymbol{\phi}, \boldsymbol{\beta}, \mathbf{s}^{(k)}, d_j)$, the posterior estimate of the DLT rate by the maximum follow-up time ω. In Phase I studies, $\hat{\psi}$ is usually approximated by plugging in $\hat{\boldsymbol{\phi}}$ and

$\hat{\boldsymbol{\beta}}$, the respective posterior medians/means of $\boldsymbol{\phi}$ and $\boldsymbol{\beta}$. Given a desired DLT rate η, the dose-schedule combination that minimizes a distance measure $d(\hat{p}_{jk}, \eta)$, which we denote (j^*, k^*) is assigned to the next patient, subject to one restriction. Both j^* and k^* cannot simultaneously be respectively more than one dose higher than j_{i-1}, the dose assigned to the most recently enrolled patient, and k_{i-1}, the schedule assigned to the most recently enrolled patient. Even though we use a "no-skipping" rule for dose-schedule escalation among successive patients, there is no such rule when it comes to de-escalation. We place no restriction on escalation of dose and schedule *within* a patient, which some may view as overly aggressive. However, in the simulation results presented in Section 3.3.2, we see no evidence of a higher than desired rate of DLTs. We also ran simulations in which the between-patient restriction on escalation to also applied within-patient (results not shown). We saw little change to the results presented in Table 3.4, except that patient assignments to acceptable combinations tended to lessen with the restriction than without it.

We adopt the measure $d(\hat{p}_{jk}, \eta) = |\hat{p}_{jk} - \eta|$, although we could adopt a different metric in order to penalize selection of toxic regimens, like that proposed in the Escalation with Overdose Control (EWOC) design of Babb, Rogatko, and Zacks (1998). EWOC selects the dose that minimizes the distance $(\hat{p}_{jk} - \eta)(1 - \delta)I(\hat{p}_{jk} > \eta) + (\eta - \hat{p}_{jk})\delta I(\hat{p}_{jk} < \eta)$, and with $\delta < 0.5$, will penalize the selection of toxic dose-schedule combinations more than our metric $|\hat{p}_{jk} - \eta|$. Furthermore, a stopping rule for excessive toxicity could be easily incorporated into our design. For example, in our simulations, we use a stopping rule stating that a trial is halted if at least three patients have been enrolled and $\hat{p}_{11} > \eta + 0.15$, in which $\hat{p}_{11}$ is the DLT rate estimate for the lowest dose-schedule combination. A similar stopping rule when all DLT rates are too low could also be used.

We emphasize that reassignment of dose and/or schedule does not apply to patients who have experienced DLT, nor to those who finished their originally assigned

treatment, nor those whose treatment was terminated early. For the remaining $n^* \leq n$ patients who are still planning to receive additional administrations, we compute $\hat{p}_\ell = 1 - \hat{\psi}(\omega|\boldsymbol{\phi}, \boldsymbol{\beta}, \mathbf{s}_\ell, \mathbf{d}_\ell)$, which is the estimated DLT rate of the administrations received so far by patient $\ell = 1, 2, \ldots n^*$. We immediately terminate the treatment of any patient ℓ for whom $\hat{p}_\ell \geq \eta + 0.1$, as they have already received a combination that appears to be overly toxic and further treatment would be unethical. Once the treatment is terminated, no additional administrations will be given to the patient but this patient is still under follow-up until a DLT occurs or the maximum follow-up time ω is reached.

For each of the remaining patients for whom $\hat{p}_\ell < \eta + 0.1$, we need to consider how many more administrations might be given and which dose would be given at each of those administrations. Specifically, if m_ℓ is the number of administrations received so far, we compute $\delta_\ell^k = (m_k - m_\ell)$, for each schedule k, including the schedule to which the patient was originally assigned. Among all schedules with $\delta_\ell^k > 0$, let $\boldsymbol{s}_{\ell+}^{(k)}$ denote the *remaining* administration times for schedule k that could still be assigned to patient ℓ. We consider the combination of each $\boldsymbol{s}_{\ell+}^{(k)}$ with each dose j and let $\boldsymbol{d}_{\ell+}^{(jk)}$ denote the remaining dose assignments, which is a vector of δ_ℓ^k elements each with the value d_j. We then compute $P_\ell^{(jk)} = 1 - \hat{\psi}(\omega|\boldsymbol{\phi}, \boldsymbol{\beta}, [\mathbf{s}_\ell, \boldsymbol{s}_{\ell+}^{(k)}], [\mathbf{d}_\ell, \boldsymbol{d}_{\ell+}^{(jk)}])$, which is the probability of DLT by ω for patient ℓ for each of these possible reassignments appended to what he has already received. We will reassign patient ℓ according to whichever $P_\ell^{(jk)}$ is closest to the targeted DLT rate, η. We emphasize again that one of the possible "reassignments" is simply the assignment currently belonging to patient ℓ.

To clarify our notation, we consider a hypothetical study of $J = 3$ doses and $K = 5$ schedules in which schedule k is comprised of k consecutive $(5+, 24-)$ courses as described in Section 3.2.1. Imagine that a new patient is to be enrolled in the study and that we have an enrolled patient ℓ who was assigned to schedule

Table 3.1: Nine possible remaining dose and schedule assignments for a hypothetical patient who has not completed their originally assigned treatment and remains under observation without DLT. Vectors of remaining administration times are $\boldsymbol{s}_{\ell+}^{(3)} = \{58, 59, 60\}$, $\boldsymbol{s}_{\ell+}^{(4)} = \{\boldsymbol{s}_{\ell+}^{(3)}, 84, 85, 86, 87, 88\}$, and $\boldsymbol{s}_{\ell+}^{(5)} = \{\boldsymbol{s}_{\ell+}^{(4)}, 112, 113, 114, 115, 116\}$
.

Decision	Dose at Each Administration	Times of Administration
No change	32	$\boldsymbol{s}_{\ell+}^{(3)}$
Change dose only	8	$\boldsymbol{s}_{\ell+}^{(3)}$
	16	$\boldsymbol{s}_{\ell+}^{(3)}$
Change schedule only	32	$\boldsymbol{s}_{\ell+}^{(4)}$
	32	$\boldsymbol{s}_{\ell+}^{(5)}$
Change dose & schedule	8	$\boldsymbol{s}_{\ell+}^{(4)}$
	8	$\boldsymbol{s}_{\ell+}^{(5)}$
	16	$\boldsymbol{s}_{\ell+}^{(4)}$
	16	$\boldsymbol{s}_{\ell+}^{(5)}$

3, has not yet experienced a DLT, and has respective administration times and doses for each administration $\mathbf{s}_\ell = \{0, 1, 2, 3, 4, 28, 29, 30, 31, 32, 56, 57\}$ and $\mathbf{d}_\ell = \{8, 8, 8, 8, 8, 16, 16, 16, 16, 16, 32, 32\}$ mg/m^2. As each schedule had a total of five planned administrations, we see that patient ℓ has completed two courses and has three administrations remaining in her third course. Assuming that the treatment received so far does not have an estimated DLT rate 10 points above the target, Table 3.1 delineates the nine possible remaining assignments that could now be given to patient ℓ. Whichever of these nine combinations, when appended to $\mathbf{s}_\ell$ and $\mathbf{d}_\ell$, leads to an estimated DLT rate by ω closest to the target DLT rate is the reassignment given to patient ℓ.

This example emphasizes the fact that we attempt to keep the dose constant within a patient as much as possible, i.e. each administration for a patient will be at the same dose until a reassignment occurs. Thus, the hypothetical patient ℓ described

above had already received two previous changes to her assignment, as her dose was increased from 8 mg/m^2, then to 16 mg/m^2, and then again to 32 mg/m^2. Of course one could consider a setting in which the best treatment plan would be contrary to this, i.e. perhaps alternating back-and-forth between two doses. However, such a treatment plan, or one that considers any of the J doses at each administration period would be infeasible in practice and would likely lead to treatment errors if the treatment plans assigned to several patients were all different and impossible to remember.

Furthermore, we have chosen to only consider reassignments when a new patient is enrolled. This certainly is not the only benchmark at which we might consider reassigning patients. For example, we might instead (or also) re-evaluate the data collected so far each time a patient completes their follow-up, either by reaching ω without a DLT or experiencing a DLT sometime before ω. Or we could re-evaluate the data each time a patient completes a course, thereby allowing a course-by-course evaluation for every patient. And if we truly wanted to optimize the treatment of every patient in the study, it would seem most sensible to evaluate each patient after every single administration. However, most of these alternate approaches are unrealistic in practice as the frequency of the necessary computations would become administratively impossible. On the opposite end of the spectrum, we could administratively set times, i.e. every three months, when we might consider reassignments that have nothing to do with patient outcomes but makes the process of re-assignment known before the trial begins. However, we feel our approach of re-evaluating assignments when each new patient is enrolled is a good compromise between optimizing the treatment of each patient as much as possible and maintaining a feasible level of computation.

3.2.6 Conduct of the Trial

We plan on enrolling a maximum of N patients in the trial, and each patient will be followed for ω days after enrollment. The first patient is enrolled at study time $t = 0$ and is assigned to the shortest schedule ($k = 1$) with the lowest dose ($j = 1$). When patient $i = 2, \ldots, N$ is to be enrolled in the study at study time t, we perform the following steps:

(1) Place each enrolled patient $i^{'} = 1, 2, \ldots i-1$ into one of two groups, either those without DLT or those with DLT;

(2) For patients without DLT, record:

   (i) $C_{i'} = 0$, indicating no DLT

   (ii) $T_{i'} = \min\{W_{i',i}, \omega\}$, where $W_{i',i}$ is the inter-patient time between patient i' and i;

(3) For patients with DLT, record:

   (i) $C_{i'} = 1$, indicating DLT

   (ii) $T_{i'} = Y_{i'}$, the patient time when a DLT occurred;

(4) For all enrolled patients, record $\mathbf{s}_{i'}$, the vector of times of each administration received, and $\mathbf{d}_{i'}$, the vector of doses given at each administration;

(5) Use the information recorded from (2)-(4) above to compute the likelihood given in Equation (3.3). Specifically, patients without DLT will contribute an amount given in Equation (3.1) and patients with DLT will contribute an amount given in Equation (3.2);

(6) Combine the likelihood with the prior distributions described in Section 3.2.4 to compute the posterior medians of $\boldsymbol{\phi}$ and $\boldsymbol{\beta}$;

(7) Apply the methods described in Section 3.2.5 to determine whether to terminate the trial, and if not, determine for each patient whether to assign a new dose and/or schedule or terminate their treatment altogether;

(8) Determine the dose and schedule assignment for patient i using the methods described in Section 3.2.5;

(9) Once all N patients have been enrolled, use all accumulated data to compute final posterior estimates of the DLT rates of each dose and schedule combination and select the combination with estimated DLT rate closest to η as the optimal combinations

3.3 Application

3.3.1 Simulation Design

In the motivating azacitidine trial, there were $J = 3$ doses of interest: 8, 16 and 24 mg/m^2, and $K = 4$ schedules with respective numbers of administrations $m_1 = 5$, $m_2 = 10$, $m_3 = 15$ and $m_4 = 20$, for a total of 12 combinations. A course consists of five daily consecutive administrations followed by 28 days of rest as described in the example in Section 3.2.1, and schedule k consists of k consecutive courses. Investigators would like to determine which of the 12 combinations has a DLT rate close to $\eta = 0.30$. The maximum follow-up time for each patient is $\omega = 116$ days. We consider the dose-schedule combinations with DLT rates of $\eta \pm 0.10$ to be acceptable choices of the MTC, since a small deviation from η is acceptable for the investigators (Braun et al. 2007). A maximum of $N = 60$ patients will be enrolled.

We considered two skeletons that we feel would reflect those most commonly used in practice. Skeleton 1 specifies the *a priori* MTC to exist at middle combinations whereas Skeleton 2 specifies the highest combinations as the *a priori* MTC; the actual values of the skeletons can be found in Table 3.2. For each skeleton, we used the

Table 3.2: The two skeletons used in the simulation study. The boldfaced values correspond to acceptable combinations.

		Schedule			
Skeleton	Dose(mg/m^2)	1	2	3	4
1	8	0.03	0.12	**0.30**	0.50
	16	0.15	**0.30**	0.50	0.60
	24	**0.30**	0.50	0.60	0.75
2	8	0.02	0.06	0.15	**0.25**
	16	0.08	0.15	**0.25**	**0.30**
	24	0.15	**0.25**	**0.30**	**0.38**

methods described in Section 2.5 to calculate the prior means. For Skeleton 1, this leads to $\mu_0 = -4.80$, $\mu_1 = -0.32$, $\mu_\gamma = -0.818$ and $\alpha = 1.73$. The corresponding values for Skeleton 2 were $\mu_0 = -5.50$, $\mu_1 = -0.43$, $\mu_\gamma = -0.822$ and $\alpha = 1.72$. With either skeleton, the mode of the hazard function is around four days after administration. For both skeletons, we calibrated the prior variances of our model parameters through a process outlined in Section 3.5.2.

We examined our approach in 16 different scenarios that are summarized in Table 3.3. The true DLT rates of every combination of dose and schedule were not generated by the model used in our methods but were instead created using an approach outlined in Section 3.5.1.

Table 3.3 also contains three metrics that seek to measure how difficult finding the MTC might be in each scenario. The first value, N_c, is the number of combinations with DLT rates within 10 points of the target η, the second value, MSE, denotes the mean sum-of-squared-errors for the fit of the linear model $\log[-\log(1-p_{d_j})] = \beta_0 + \exp(\beta_1)d_j$, and the third value, SD, is the sample standard deviation of the 12 DLT rates. Thus, smaller values of N_c and SD would indicate greater difficulty of finding the MTC and larger values of MSE would indicate that the linearity assumed

Table 3.3: Summary of the 16 scenarios studied, including the actual DLT rates of each dose and schedule combination and three metrics that measure the difficulty of identifying an MTC. Boldfaced values indicate dose and schedule combinations with DLT rates within 10 points of the desired DLT rate $\eta = 0.30$.

			Schedule								
Method	Scenario	Dose(mg/m^2)	1	2	3	4	p_d	N_c	MSE	SD	ξ
Independent	1	8	0.05	0.10	0.14	0.18	0.010	3	0.03	0.08	n/a
		16	0.06	0.12	0.18	**0.23**	0.013				
		32	0.09	0.17	**0.24**	**0.30**	0.018				
	2	8	0.05	0.10	0.14	0.18	0.010	3	0.32	0.23	n/a
		16	0.09	0.17	**0.24**	**0.30**	0.018				
		32	**0.30**	0.52	0.66	0.77	0.070				
	3	8	0.02	0.05	0.07	0.10	0.005	1	0.35	0.32	n/a
		16	0.18	**0.32**	0.44	0.54	0.038				
		32	0.47	0.72	0.85	0.92	0.120				
	4	8	0.03	0.07	0.10	0.13	0.007	1	0.42	0.15	n/a
		16	0.05	0.09	0.13	0.17	.0095				
		32	0.17	**0.31**	0.42	0.52	0.036				
Gumbel	5	8	0.18	**0.31**	0.41	0.50	0.048	4	0.06	0.16	0.88
		16	**0.22**	**0.36**	0.48	0.57	0.058				
		32	**0.29**	0.47	0.59	0.69	0.080				
	6	8	0.13	0.18	**0.23**	**0.27**	0.050	3	0.30	0.20	0.60
		16	**0.33**	0.45	0.54	0.60	0.140				
		32	0.43	0.57	0.66	0.72	0.190				
	7	8	**0.30**	0.44	0.54	0.61	0.110	2	0.07	0.15	0.70
		16	**0.35**	0.50	0.60	0.68	0.130				
		32	0.46	0.63	0.73	0.80	0.180				
	8	8	0.02	0.06	0.11	0.17	.0021	2	0.22	0.18	1.5
		16	0.07	0.17	**0.30**	0.42	0.006				
		32	0.11	**0.27**	0.44	0.59	0.010				
	9	8	0.17	**0.28**	**0.36**	0.43	0.049	2	0.29	0.25	0.81
		16	0.43	0.62	0.74	0.82	0.140				
		32	0.56	0.76	0.86	0.92	0.200				
	10	8	0.02	0.06	0.10	0.15	.0018	1	0.29	0.12	1.5
		16	0.03	0.07	0.13	0.17	.0023				
		32	0.07	0.17	**0.30**	0.42	0.006				
Frank	11	8	0.05	0.09	0.13	0.17	0.010	3	0.03	0.14	1.5
		16	0.09	0.17	**0.24**	**0.30**	0.020				
		32	0.19	**0.32**	0.42	0.50	0.043				
	12	8	0.19	**0.32**	0.43	0.51	0.044	3	0.08	0.16	1.5
		16	**0.26**	0.43	0.54	0.63	0.065				
		32	**0.31**	0.49	0.61	0.69	0.080				
	13	8	0.07	0.13	0.18	**0.23**	.0146	3	0.13	0.21	1.5
		16	0.18	**0.31**	0.41	0.49	0.042				
		32	**0.33**	0.52	0.64	0.72	0.087				
	14	8	0.18	**0.31**	0.41	0.49	0.042	2	0.04	0.19	1.5
		16	**0.29**	0.47	0.59	0.67	0.075				
		32	0.42	0.62	0.74	0.82	0.120				
	15	8	0.03	0.06	0.08	0.11	0.006	2	0.31	0.13	1.5
		16	0.05	0.09	0.13	0.19	0.010				
		32	0.15	**0.27**	**0.36**	0.44	0.035				
	16	8	0.53	0.74	0.84	0.90	0.170	0	0.01	0.15	1.5
		16	0.53	0.75	0.85	0.91	0.175				
		32	0.55	0.75	0.86	0.91	0.178				

in our model may be suspect and lead to a poorer ability of correctly identifying an MTC.

We simulated patients to have exponentially distributed inter-arrival times with a mean of two weeks, and we divided all the follow-up times by 10 to achieve better numeric stability for our model. When a new patient is enrolled, an interim analysis is performed in which a single chain of $6,000$ samples, after a burn-in of $4,000$ samples, is drawn from the posterior distribution for each parameter. These posterior draws are then used to determine the dose and schedule assigned to the new patient as well as any dose and/or schedule reassignments for each currently enrolled patient still being followed. We then simulate for each a binary indicator of DLT using the method outlined in Section 3.5.1 depending upon the scenario examined. If a patient is simulated to have a DLT, the time of the DLT is drawn uniformly from the interval $[4 + 24(k-1), 4 + 24k]$ under their assigned schedule k, which also implies that all possible DLTs occur by $\omega = 116$ days. We did perform simulations of our design using our assumed model to simulate DLTs and came to similar final conclusions, and we have omitted those results for brevity.

We compared the performance of our approach that allows for patient reassignment (Design A) with the traditional approach that does not allow for patient reassignment (Design B). We evaluated the performance of both approaches by comparing the correct selection frequency at the end of the study, the mean proportion of patients assigned to each dose-schedule combination and the mean proportion of patients who experienced DLTs. We performed $1,000$ simulations in each scenario; our computer code is available upon request.

3.3.2 Simulation Results

Table 3.4 contains a summary of the performance of Design A (with reassignment) and Design B (without reassignment) in the 16 scenarios described in Table 3.3. For each design, this summary is a series of eight columns. The first four columns describe

the proportion of simulations in which the MTC selected at the end of the study was not found, had a DLT rate more than 10 points below the desired DLT rate η (column "L"), within 10 points of η (column "In"), or more than 10 points above η (column "H"). The next three columns have a similar interpretation related to the average percentage of dose-schedule assignments during the study. The eighth column, labeled "DLT" is the average of the proportion of observed DLTs among the 1,000 simulations.

Overall, we are able to identify acceptable dose-schedule combinations at the end of the study in a majority of simulations in the first 15 scenarios, whether or not reassignment is used, as well as terminate the study early in scenario 16. These results are not surprising, as the primary goal of reassignment is to optimize the assignments of patients enrolled in the study, rather than improve the final decision at the end of the study. Scenarios 4, 8, 9, and 10 have the lowest percentages of identifying the MTC at an acceptable combination, which is partially explained by the fact that these scenarios have only 1 or 2 acceptable combinations to choose from. These scenarios also have some of the largest values of MSE, indicating that the assumption of linearity in our model is suspect. Nonetheless, we emphasize that all 15 scenarios have DLT rates that come from models that are different from our assumed model, so that our approach works well even when the model is misspecified.

With regard to patient assignments during the study, we see that including reassignment leads to a higher proportion of patients assigned to acceptable combinations than without reassignment. For example, Design A assigned 58% of the patients to acceptable combinations in scenario 4, compared with only 29% for Design B, and the corresponding percentages in scenario 10 are 53% and 26%, respectively. Moreover, in all 15 scenarios, the average DLT rate when using reassignment is never more than the average observed DLT rate without reassignment and is always close to the desired DLT rate. Although Scenarios 7, 9 and 10 have 33%, 40% and 31% of patients,

Table 3.4: Simulation results for Design A (with reassignment) and B (without reassignment), prior standard deviation $\sigma = 2$ and Skeleton 2. For each design, there are four columns under "Selection" list the percentage of simulations in which the MTC was not identified, was identified at combinations with DLT rates below that desired, within that desired, and above that desired, respectively. Corresponding columns are listed under "Assignment" to describe the average percentage of patients assigned to each combination. DLT = mean proportion of patients who experienced DLT. The columns under "Reassignment" gives summary statistics on reassignment: Rp = average proportion of patients receiving at least one reassignment during the study, Rm (Rsd) = mean (standard deviation) of number of reassignments a patient received, R1 (Rn) = the average minimum (maximum) number of reassignments a patient received.

		Selection				Assignment				Reassignment				
Scen	Design	None	L	In	H	L	In	H	DLT	Rp	Rm	Rsd	R1	Rn
1	A	0	5	**95**	0	10	**90**	0	27	38	0.9	1	0	4.1
	B	0	4	**96**	0	12	**88**	0	27					
2	A	0	13	**83**	3	19	**63**	18	31	56	1.2	1.3	0	5.2
	B	0	12	**86**	3	29	**52**	18	31					
3	A	1	14	**71**	14	22	**53**	24	32	53	1	1.2	0	4.7
	B	0	25	**62**	13	29	**48**	23	33					
4	A	0	27	**53**	20	19	**58**	23	30	68	1.6	1.5	0	5.8
	B	0	32	**48**	20	33	**29**	38	31					
5	A	2	2	**78**	17	7	**62**	29	35	43	0.7	0.9	0	3.8
	B	2	2	**81**	16	5	**63**	31	36					
6	A	1	18	**56**	25	18	**52**	29	34	44	0.7	0.9	0	3.7
	B	1	12	**64**	23	21	**44**	35	34					
7	A	11	0	**69**	20	0	**58**	33	36	20	0.3	0.6	0	2.5
	B	9	0	**66**	24	0	**54**	39	37					
8	A	0	12	**56**	32	14	**58**	28	32	73	1.7	1.5	0	5.9
	B	0	15	**58**	27	25	**36**	38	32					
9	A	6	18	**49**	28	18	**38**	40	34	25	0.4	0.7	0	2.7
	B	3	14	**54**	29	17	**38**	44	36					
10	A	0	25	**53**	22	16	**53**	31	30	62	1.5	1.5	0	5.6
	B	0	26	**53**	21	30	**26**	44	30					
11	A	0	10	**79**	11	15	**64**	22	31	67	1.6	1.5	0	5.7
	B	0	12	**76**	13	21	**48**	31	31					
12	A	2	5	**75**	18	9	**59**	30	35	37	0.6	0.9	0	3.5
	B	2	4	**74**	20	9	**53**	36	36					
13	A	0	14	**67**	19	15	**63**	22	33	56	1.1	1.2	0	4.8
	B	0	13	**69**	17	24	**52**	24	33					
14	A	3	8	**61**	28	12	**57**	28	35	32	0.5	0.8	0	3.2
	B	3	7	**58**	32	13	**51**	35	36					
15	A	0	17	**74**	8	16	**64**	20	30	64	1.5	1.5	0	5.7
	B	0	26	**65**	9	26	**44**	30	30					
16	A	**89**	0	0	11	0	1	37	21	5	0.1	0.3	0	1.4
	B	**87**	0	0	13	0	0	40	22					

respectively, assigned to toxic combinations, most of the combinations in Scenarios 7 and 9 are overly toxic, and exposing a higher proportion of patients to toxic combinations in these scenarios seems unavoidable. Furthermore, in Scenario 10, all of the assignments to toxic combinations were to combinations with DLT rates in the range of $40\% - 42\%$, and in Scenario 9, 25% of patients were assigned to combinations with DLT rates between 40% and 45%. We also conducted simulations using the EWOC distance measure described in Section 3.2.5, but due to space limitations, we omit these results. We found that using the EWOC distance measure led to treating fewer patients at overly toxic combinations in some scenarios, but this apparent increase in safety came with a reduced ability of finding the MTC at the end of the trial.

The final five columns of Table 3.4 contain a summary of the number of reassignments per patient that occurred in each of the 16 scenarios. Rates of reassignment above 0.60 were seen in scenarios 4, 8, 10, 11, and 15. Although the explanation for the high rate of reassignment is not immediately obvious, a partial explanation is that acceptable combinations in these scenarios appear with longer schedules of the highest dose, with even longer schedules then becoming overly toxic. In contrast, Scenario 1 has a much lower rate of reassignment because that scenario had no overly toxic combinations. We also see that less than two reassignments occurred per patient on average in all the scenarios and just under six reassignments was the maximum number of reassignments per patient on average in all scenarios.

We also examined the sensitivity of our design to the maximum sample size by repeating our simulations using sample sizes of 50 and 70, the results of which are summarized in Table 3.5. The results show that decreasing the sample size from 60 to 50 may result in a nontrivial loss is selecting correct combinations in Scenarios 4, 10 and 15. However, increasing the sample size from 60 to 70 only produces a minor increase in the selection of correct combinations. Therefore, we selected a maximum number of 60 patients for this study.

Table 3.5: Simulation results for Design A with sample sizes of $N = 50$ and $N = 70$. For each design, columns "Selection" gives the percentage of identifying three categories of combinations as the MTC: unacceptable inefficacious combinations ("L") ; acceptable combinations ("In") and too toxic combinations ("H"). For each design, columns "Assignment" gives the mean proportion of patients assigned to the three categories, columns "DLT" give the mean proportion of patients who experienced DLTs.

	Design A with $N = 50$								Design A with $N = 70$							
	Selection				Assignment				Selection				Assignment			
Scen	None	L	In	H	L	In	H	DLT	None	L	In	H	L	In	H	DLT
1	0	6	**94**	0	11	**89**	0	27	0	5	**95**	0	9	**91**	0	27
2	0	17	**79**	4	21	**58**	21	31	0	11	**86**	2	17	**66**	17	31
3	0	16	**66**	17	24	**49**	26	33	0	12	**73**	15	20	**56**	23	32
4	0	32	**45**	24	20	**54**	26	30	0	25	**57**	18	18	**61**	21	30
5	2	3	**78**	18	8	**58**	32	35	1	2	**82**	15	7	**66**	26	34
6	1	19	**55**	25	18	**50**	31	35	1	12	**64**	22	16	**55**	27	34
7	9	0	**66**	25	0	**55**	38	37	9	0	**73**	18	0	**61**	32	35
8	0	16	**55**	30	16	**54**	30	33	0	12	**57**	31	14	**61**	25	32
9	4	17	**48**	31	20	**34**	43	36	3	16	**53**	28	19	**39**	40	35
10	0	28	**50**	22	18	**48**	34	30	0	26	**54**	20	17	**54**	29	30
11	0	14	**74**	12	16	**60**	24	31	0	10	**80**	11	15	**65**	20	31
12	3	5	**68**	25	10	**53**	35	36	2	4	**78**	17	9	**62**	28	35
13	1	15	**64**	20	15	**59**	25	33	0	13	**70**	16	14	**66**	20	32
14	3	9	**54**	34	12	**54**	31	36	1	7	**61**	31	11	**61**	26	35
15	0	22	**68**	10	18	**58**	23	30	0	18	**75**	7	15	**66**	18	30
16	**83**	0	0	17	0	1	42	24	89	0	0	11	0	1	34	19

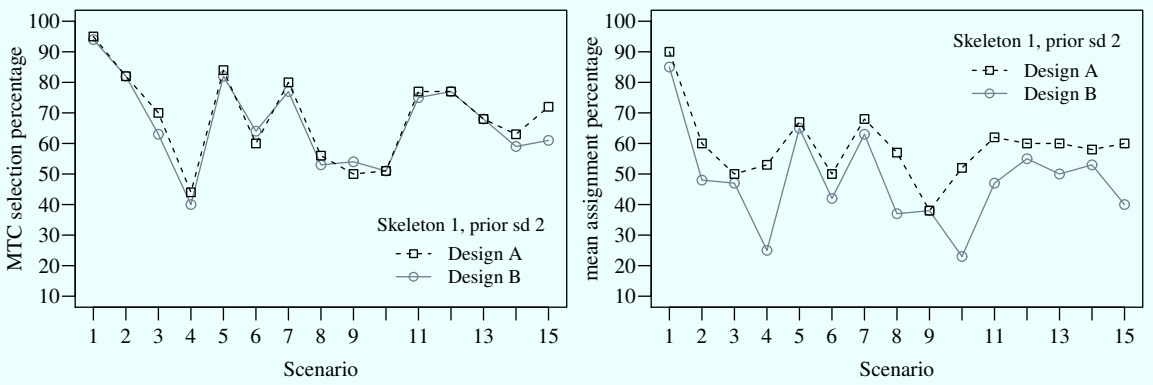

Figure 3.1: Comparison between Design A and B using Skeleton 1 and calibrated prior standard deviation 2.

Recall that we found that using a value of $\sigma = 2$ for the prior variances of the model parameters was insensitive to the skeleton used. To confirm this statement, we do not present the results when using Skeleton 1 in a tabular format like that of Table 3.4. Instead, Figure 3.1 contains a visual summary of the percentage of simulations in which the MTC was selected at an acceptable combination (left plot) and the percentage of patients assigned to acceptable combinations (right plot) in each of the 15 scenarios. As we found in Table 3.4, we are able to identify the MTC well whether or not reassignment is allowed, but that inclusion of reassignment greatly improves the treatment assignment of patients enrolled in the study.

3.4 Discussion

In our methods, we reassigned the dose and/or schedule of each enrolled patient only when a new patient was enrolled in order to optimize the treatment of each patient but maintain a feasible level of computation. As a result, with an average arrival of a new patient every 14 days relative to a follow-up of 116 days in our simulations, we considered approximately no more than eight reassignments with each new enrollment. In a study with much faster accrual, there would be many more possible reassignments to consider at each new enrollment and many of the patients could have several reassignments occurring during their treatment. Investigators may feel that

this level of possible reassignment impractical, requiring a different rule for determining when reassignment is possible. An interesting area of research is to compare our approach to other reassignment strategies to determine when fewer reassignments are allowed, yet do not lead to significantly worsened operating characteristics.

Although we restricted the estimate of the MTC at the end of the trial to only exist among dose-schedule regimens with dose constant within patient, our methods are flexible enough to regimens with dose variations within-patient that ultimately might provide a better MTC estimate since we are considering a richer set of candidate strategies. In addition, we could consider multiple MTC estimates each of which that are qualitatively different from each other. For example, we could consider two completing MTC estimates, one with a few administrations of a higher dose and another with many administrations of a lower dose. It is very possible that one schedule may be more effective or easier to administer and these facts are not yet part of our decision-making process. Certainly our methods can be generalized to incorporate more information and further the process of discovering new drugs and how to best administer them, and is an exciting avenue of research. Furthermore, our methods belong to the family of dynamic treatment regimes (Murphy, 2002) and use of additional patient information, such as a key biomarker for response or risk factors for lack of response, might further the cause to "personalize" the dose and schedule assigned to each patient. However, the major limiting factor for all these extensions is the small sample size used in most early-phase clinical trials.

We introduce our design in the setting of dose-schedule finding studies, although our methods could be easily applied to other settings. If the schedule were fixed while the dose varied, our design would be similar to the TITE-CRM in which the weight function would be determined by the functional form of the hazard function of a single administration in our model. In Phase I/II studies, one can easily adapt our method to the work of Yuan and Yin (2009) to model late-onset toxicity/efficacy

and introduce intra-patient dose changes. Similar modifications could be made to apply our approach to trials of combinations of two agents by specifying the hazard function of the toxicity or efficacy after a single administration to be a function of doses of both agents. However, the model that takes account of the joint distribution of toxicity and efficacy as a function of possibly multiple doses is not immediately obvious and should be carefully chosen. Lastly, in the motivating azacitidine trial, the highest dose and longest schedule was reached with no evidence of unacceptable toxicity and investigators decided to add more doses to the trial. The flexibility of our model would allow one to adaptively estimate the DLT rates of other doses and/or schedules to determine which, if any, might be added to the trial once it has begun.

3.5 Appendices

3.5.1 Algorithm for Simulating True DLT Rates

We let p_{d_j} denote the probability of DLT by ω after a single administration of dose $j = 1, 2, 3$, and we let n_j denote the number of administrations of dose j received. Then, for a treatment schedule of n_1, n_2, and n_3 administrations of doses 1, 2, and 3, respectively, regardless of their order, we denote the actual probability of DLT by time ω as $P_{true}(d_1, d_2, d_3, n_1, n_2, n_3)$. In scenarios 1-4, we assume that all administrations have independent effects (which is also the assumption used in our model), i.e.

$$P_{true}(d_1, d_2, d_3, n_1, n_2, n_3) = 1 - q_{d_1}^{n_1} q_{d_2}^{n_2} q_{d_3}^{n_3},$$

in which $q_{d_j} = 1 - p_{d_j}$. In scenarios 5-10, we assume that all administrations have correlated effects modeled via a Gumbel copula, i.e.

$$P_{true}(d_1, d_2, d_3, n_1, n_2, n_3) = 1 - \exp\left(- \left\{ \sum_{j=1}^{3} n_j [-\log(q_{d_j})]^{1/\xi} \right\}^{\xi} \right),$$

and is the same as Equation (3.4) when the correlation parameter $\xi = 1$. In scenarios 11-15, we assume that all administrations have correlated effects modeled via a Frank copula, i.e.

$$P_{true}(d_1, d_2, d_3, n_1, n_2, n_3) = 1 + \frac{1}{\xi} \log \left(1 + \frac{\prod_{j=1}^{3} [\exp(-\xi q_{d_j}) - 1]^{n_j}}{(e^{-\xi} - 1)^{\sum_j n_j - 1}} \right),$$

and is the same as Equation(3.4) when $\xi \to 0$.

The value of ξ used in each of scenarios 5-16 is shown in the last column of Table 3 of the manuscript and the actual values of p_{d_1}, p_{d_2}, and p_{d_3} used in each of the fifteen scenarios are shown in the column labeled "p_d" in Table 3 of the manuscript. Although Table 3 of the manuscript displays the DLT rates for each of the dose and schedule combinations under study, the values in the column labeled "p_d" can be used to compute the actual DLT rates for patients who receive a reassignment that does not fit one of these dose-schedule combinations. For example, suppose we have a patient in scenario 1 who has been assigned to five administrations of 8 mg/m^2 and five administrations of 16 mg/m^2. This patient has a probability of DLT by ω of $1 - (1 - 0.010)^5(1 - 0.013)^5 \approx 0.11$.

3.5.2 Calibration of Prior Variance

To achieve good operating characteristics, we first calibrated the prior standard deviations via simulation using a maximum number of 60 patients. We let σ_1, σ_2 and σ_3 have the same value σ to simplify the calibration process. The prior standard deviation that performs best among $\sigma = 1, 2, 5$ would be used in the study. We certainly could have examined more values of σ, but felt that choosing among these three values was sufficient and any small deviations in performance with other possible values for σ were outweighed by the increased amount of simulation time required.

The first row of Figure 3.2 shows how the prior variance impacts the proportion of patients assigned to acceptable dose-schedule combinations in Design A using either

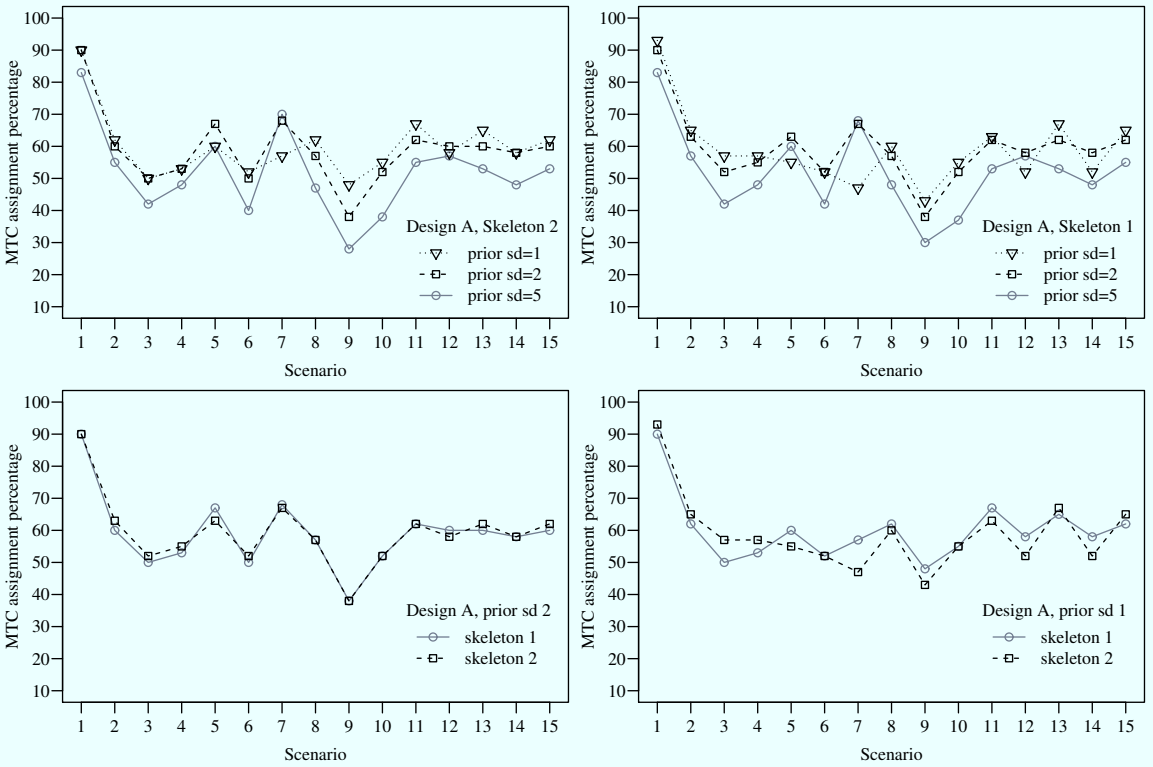

Figure 3.2: Impact of prior standard deviation and skeletons on the MTC assignment for Design A.

skeleton 2 (upper left plot) or skeleton 1 (upper right plot). For both skeletons, the design using $\sigma = 5$ performs worse than using $\sigma = 1$ or 2 in most of the 15 examined scenarios. However, from these two plots, it is not clear which among between $\sigma = 1$ and $\sigma = 2$ would be preferred. Therefore, the bottom two plots in Figure 3.2 attempt to assess the sensitivity of the results to the chosen skeleton when using $\sigma = 2$ (lower left plot) and $\sigma = 1$ (lower right plot). From these two plots, we see that there is greater variation in the results when using $\sigma = 1$ than $\sigma = 2$. We also performed similar analyses for the proportion of simulations in which acceptable dose-schedule combinations were selected at the end of the study, as well as repeating our calibration with Design B (no reassignment), and found little change in our conclusions. Therefore, we selected $\sigma = 2$ to be the prior standard deviation used in the study. And since our design is not sensitive to the skeletons when $\sigma = 2$, we have chosen to use Skeleton 2 in our study.

3.5.3 Illustration of Patient Reassignments

We selected the results from one of our simulations in Scenario 1 to visually display how patient reassignments occur during a trial. To that end, we have a series of tables that summarize patient assignments and reassignments after 10 patients have been enrolled (Table 1), after 20 patients have been enrolled (Table 2), after 30 patients have been enrolled (Table 3), etc., until the maximum sample size of $N = 60$ has been enrolled (Table 6).

Recall that there are 3 dose levels and 4 nested schedules under investigation. Each treatment course consists of 5 daily administrations followed by 28 days of rest. In Tables 1-6, we let (a*b) denote a treatment assignment that corresponds to b administrations at dose level a. Similarly, (a*b, c*d) denotes b administrations at dose level a followed by d administrations at dose level c. For instance, (2*5, 3*5) denotes that the patient received five administrations of dose level 2 for five consecutive days, had 28 days of rest, then received another five administrations of dose level 3.

With regard to the information in each table, ID refers to the identification number of the patient, "Time enrolled" is the study time when each patient is enrolled, "Assign" contains the most recent assignment of the patient, "Reassign" indicates the dose-schedule reassignment for each patient who received a reassignment, "DLT" contains a indicator (1=yes; 0=no) if the patient experienced a DLT, "Follow-up" is the time to DLT or the follow-up time if there has not been a DLT, and "Comp" indicates whether the patient completed their follow-up (Y=yes; N=no).

We see in Table 1 that the first patient received the lowest dose-schedule combination (dose 1, schedule 1), and because no DLT had been observed when patient 2 was enrolled, we initially assigned patient 2 to the combination of dose 2 and schedule 2. Similarly, we treated patient 3 at the combination of dose 3 and schedule 3 at enrollment. When the fourth patient was enrolled, not only was that patient assigned

Table 3.6: Summary of patient assignments and re-assignments when patient 11 is enrolled.

ID	Time Enrolled	Assign	Reassign	DLT	Follow-up	Comp
1	0.0	1*5	n/a	0	116.0	Y
2	4.5	2*10	2*5, 3*5	1	112.2	Y
3	15.7	3*15	3*20	0	116.0	Y
4	48.0	3*20	n/a	1	59.4	Y
5	83.3	3*20	3*15	0	42.0	N
6	90.7	3*20	3*5, 2*15	0	34.6	N
7	97.3	3*20	3*15	0	28.0	N
8	103.6	3*20	3*15	0	21.7	N
9	109.0	3*20	3*15	0	16.3	N
10	117.8	2*20	2*5, 3*15	0	7.5	N

to the current best regimen (dose 3 and schedule 4), but the assignments of the first three patients were re-examined for potential reassignment. As a result, the assignment for patient 2 was escalated to 2*5, 3*15 and the assignment for patient 3 was escalated from 3*15 to 3*20. Table 1 also demonstrates intra-patient de-escalation. For instance, after we observed two DLTs for patients 2 and 4, the initial assignments for patients 5-9 were changed to a dose-schedule combination with a lower DLT rate. Finally, when patient 11 was enrolled, we saw that the initial assignment for patient 10 was escalated from 2*20 to 2*5, 3*15 because we only saw two DLTs by that time.

Table 2 also demonstrates how patient reassignments can be altered as more patients are enrolled. For example, in Table 1, patient 6, who had been originally assigned to dose 3 and schedule 4 was reassigned to five administrations of dose 3 and 15 administrations of dose 2 by the time subject 11 was enrolled. Table 2 then shows that patient 6 was reassigned again with a reduction to 10 administrations of dose 2 instead of the original 15 administrations, due to the DLT experienced by patient 10. Similar conclusions can be reached from the information in Tables 3-6

Table 3.7: Summary of patient assignments and re-assignments when patient 21 is enrolled.

ID	Time Enrolled	Assign	Reassign	DLT	Follow-up	Comp
1	0.0	1*5	n/a	0	116.0	Y
2	4.5	2*5,3*5	n/a	1	112.2	Y
3	15.7	3*20	n/a	0	116.0	Y
4	48.0	3*20	n/a	1	59.4	Y
5	83.3	3*15	n/a	0	116.0	Y
6	90.7	3*5,2*15	3*5,2*10	0	116.0	Y
7	97.3	3*15	3*10, 1*5	0	116.0	Y
8	103.6	3*15	3*10, 1*5	0	116.0	Y
9	109.0	3*15	3*10, 1*5	0	116.0	Y
10	117.8	2*5,3*15	n/a	1	21.6	Y
11	125.3	3*15	3*5, 1*10, 3*5	0	116.0	Y
12	151.0	2*20	2*10, 3*10	0	116.0	Y
13	196.6	3*20	n/a	0	110.2	N
14	204.9	3*20	n/a	0	102.2	N
15	237.9	3*20	n/a	1	29.3	Y
16	239.2	3*20	n/a	0	67.9	N
17	247.4	3*20	n/a	0	59.7	N
18	278.6	3*20	n/a	0	28.5	N
19	282.7	3*20	n/a	0	24.4	N
20	294.2	3*20	n/a	0	12.9	N

Table 3.8: Summary of patient assignments and re-assignments when patient 31 is enrolled.

ID	Time Enrolled	Assign	Reassign	DLT	Follow-up	Comp
13	196.9	3*20	n/a	0	116.0	Y
14	204.9	3*20	n/a	0	116.0	Y
15	237.9	3*20	n/a	1	29.3	Y
16	239.2	3*20	n/a	0	116.0	Y
17	247.4	3*20	n/a	0	116.0	Y
18	278.6	3*20	n/a	1	75.9	Y
19	282.7	3*20	n/a	0	116.0	Y
20	294.2	3*20	n/a	0	116.0	Y
21	302.5	3*20	n/a	0	116.0	Y
22	307.1	3*20	n/a	1	71.8	Y
23	314.6	3*20	n/a	1	76.0	Y
24	339.2	3*20	3*15	0	116.0	Y
25	340.0	3*20	n/a	1	18.2	Y
26	377.7	3*20	3*5, 2*10	0	82.2	N
27	391.7	3*15	3*5, 2*15	1	27.8	Y
28	410.6	3*15	n/a	1	27.5	Y
29	458.4	2*20	n/a	0	1.5	N
30	458.9	2*20	n/a	0	1.0	N

Table 3.9: Summary of patient assignments and re-assignments when patient 41 is enrolled.

ID	Time Enrolled	Assign	Reassign	DLT	Follow-up	Comp
26	377.7	3*5, 2*10	n/a	0	116.0	Y
27	391.7	3*5, 2*15	n/a	1	27.8	Y
28	410.6	3*15	n/a	1	27.5	Y
29	458.4	2*20	2*5, 3*10	1	35.1	Y
30	458.9	2*20	2*5,3*5,1*5,2*5	0	116.0	Y
31	459.9	2*20	2*5,3*5,1*5,2*5	0	116.0	Y
32	468.0	2*20	2*17, 3*3	0	116.0	Y
33	500.6	2*20	2*10, 3*5, 2*5	0	91.1	N
34	520.2	2*20	2*10, 3*10	0	71.5	N
35	522.5	2*20	2*8, 3*2, 2*5,3*5	0	69.2	N
36	532.0	2*20	2*5,3*5,2*5,3*5	0	59.7	N
37	538.4	2*20	2*5, 3*10	0	53.3	N
38	540.6	2*20	2*5, 3*10	0	51.1	N
39	553.5	3*10	3*5,2*5,3*5	0	38.2	N
40	570.2	3*15	3*5, 2*15	0	21.5	N

Table 3.10: Summary of patient assignments and re-assignments when patient 51 is enrolled.

ID	Time Enrolled	Assign	Reassign	DLT	Follow-up	Comp
33	500.6	2*10, 3*5, 2*5	n/a	0	116.0	Y
34	520.2	2*10, 3*10	n/a	0	116.0	Y
35	522.5	2*8, 3*2, 2*5,3*5	n/a	0	116.0	Y
36	532.0	2*5,3*5,2*5,3*5	n/a	0	116.0	Y
37	538.4	2*5, 3*10	2*5,3*5,2*5,3*5	0	116.0	Y
38	540.6	2*5, 3*10	2*5,3*5,2*5,3*5	1	89.9	Y
39	553.5	3*5,2*5,3*5	3*5,2*10,3*5	0	116.0	Y
40	570.2	3*5, 2*15	3*5,2*10,3*5	0	116.0	Y
41	591.7	3*15	3*10,2*5,3*5	0	85.0	N
42	591.8	3*15	3*10,2*5,3*5	0	84.9	N
43	601.6	3*15	3*10,2*5,3*5	0	75.1	N
44	605.9	3*15	3*10,2*5,3*5	0	70.8	N
45	634.2	3*15	3*15	0	42.5	N
46	655.5	3*15	n/a	0	21.2	N
47	664.4	3*15	n/a	0	12.3	N
48	668.1	3*15	n/a	0	8.6	N
49	668.6	3*15	n/a	0	8.1	N
50	673.4	3*15	n/a	0	3.3	N

Table 3.11: Summary of patient assignments and re-assignments at the end of the study.

ID	Time Enrolled	Assign	Reassign	DLT	Follow-up	Comp
41	591.7	3*10,2*5,3*5	n/a	0	116.0	Y
42	591.8	3*10,2*5,3*5	n/a	0	116.0	Y
43	601.6	3*10,2*5,3*5	n/a	0	116.0	Y
44	605.9	3*10,2*5,3*5	n/a	0	116.0	Y
45	634.2	3*15	n/a	0	116.0	Y
46	655.5	3*15	3*20	0	116.0	Y
47	664.4	3*15	n/a	0	116.0	Y
48	668.1	3*15	n/a	0	116.0	Y
49	668.6	3*15	3*15,1*5	0	116.0	Y
50	673.4	3*15	3*15,1*5	0	116.0	Y
51	676.7	3*15	n/a	1	18.1	Y
52	712.7	3*15	3*20	1	28.0	Y
53	714.1	3*15	3*20	0	116.0	Y
54	728.3	3*15	3*20	1	113.5	Y
55	744.6	3*15	3*20	1	86.7	Y
56	761.0	3*20	n/a	0	116.0	Y
57	761.1	3*20	n/a	0	116.0	Y
58	766.6	3*20	n/a	0	116.0	Y
59	769.9	3*20	n/a	0	116.0	Y
60	773.7	3*20	n/a	0	116.0	Y

CHAPTER IV

A Bayesian Phase I/II Adaptive Design to Simultaneously Optimize Dose and Schedule Assignments by Modeling Toxicity and Efficacy as Time-to-event Outcomes

4.1 Introduction

Most published methodology of Phase I or Phase I/II oncology trials focuses on dose finding. In such studies, the investigators usually propose a dosing schedule prior to the onset of a study and aim to identify an optimal dose in terms of toxicity and/or efficacy under that schedule. The traditional 3+3 method or more recent model-based methods, for example, the Continual Reassessment Method (CRM) proposed by O'Quigley, Pepe, and Fisher (1990), are suitable for Phase I dose-finding studies. Alternatively, many Phase I/II dose-finding designs based on the joint outcomes of both toxicity and efficacy have been proposed. For example, see Braun (2002), Thall and Russell (1998), Thall and Cook (2004) and Yuan and Ying (2009).

It is obvious that such dose-finding studies might not perform well when the pre-determined dosing schedule is not chosen correctly and leads to the conclusion that all the doses under investigation are either too toxic or too inefficacious. To address this

issue, Braun et al. (2007) first proposed a dose- and schedule-finding design based on the schedule-finding design of Braun et al. (2004) by defining separate hazard functions for modeling time to a dose-limiting toxicity (DLT) for each dose. Liu and Braun (2009) also generalized the work of Braun et al. (2004) by incorporating smoother hazard functions. Although the cited work focuses on nested schedules, Li et al. (2008) explored a Phase I/II adaptive design for finding an optimal dose and schedule combination when schedules are not nested by using a Bayesian isotonic transformation and binary patient outcomes.

However, there are several design issues the current methods in dose- and schedule-finding studies fail to address. First, most published methods for dose-schedule finding were created with solely toxicity as the outcome while the extension to where toxicity and efficacy outcomes are considered jointly is not obvious. Second, even though Li et al. (2008) do model toxicity and efficacy outcomes jointly for dose-schedule finding, their design requires binary patient outcomes and patients have to be fully followed before a new cohort of patients can be enrolled. However, the efficacy outcome usually requires a relatively long period of follow-up, resulting in an undesirably long trial. In addition, a trial using such a design would be suspended to enrollment during fast accrual because incoming patients would have to wait until all the enrolled patients complete their follow-up. Third, intra-patient dose modifications during a trial are common in practice, often a result of unexpected outcomes that may benefit from a reduced dose. Also, in newer oncology trials involving molecularly targeted agents with low toxicity rates, the investigator might intentionally apply intra-patient dose escalation to those patients who do not achieve a certain efficacy endpoint after a first treatment cycle in order to have better clinical outcomes (Cutsem et al. 2012). However, there has not been any systematic approach to accommodate intra-patient dose modification in Phase I/II studies.

To address the above design issues, we first define separate hazard functions for

time to toxicity and time to response after a patient receives a single administration. We extend these hazards to accommodate multiple administrations using a non-mixture cure rate model and derive the marginal survival function for toxicity and efficacy outcomes. A copula model is then assumed for the joint survival function. We further define acceptable and optimal dose-schedule treatment regimes using the criteria based on both estimated toxicity and response rates. Patient accrual, data monitoring, early stopping and outcome-adaptive regimen assignment are done continually throughout the trial under a Bayesian formulation. We introduce two case studies that motivate our work in Section 4.2 for both nested and non-nested schedules. We describe our method and model in Section 4.3 and introduce the trial conduct in Section 4.4. We illustrate our method and algorithm in our motivating trials in Sections 4.5 and 4.6 for nested and non-nested schedules, respectively. We conclude with a discussion in Section 4.7.

4.2 Case Studies

4.2.1 Nested Schedule - Treating acute myelogenous leukemia (AML) and myelodysplastic syndrome (MDS)

Azacitidine is a DNA methyltransferase inhibitor with activity in myeloid disease. It was hypothesized by de Lima, et al. (2010) that low-dose azacitidine administered after transplant would reduce recurrence rates of AML and MDS. Instead of considering the long-term event-free survival, a short term efficacy endpoint in an early phase trial is complete response (CR). The investigators investigated combinations of five daily doses: 8, 16 and 32 mg/m^2 and 4 schedules: 1, 2, 3, or 4 cycles, each with five days of drug and 23 days of rest. A limited number of azacitidine cycles were administered because of concerns of sustained myelosuppression as well as possible induction of graft-versus-host disease(GVHD). Toxicity is defined as any of the following adverse events occurring within 116 days from the start of the first cycle:

1) National Cancer Institute Grade 3 or higher renal, hepatic, cardiac, pulmonary, or neurologic toxicity; 2) Grade 3/5 acute GVHD; 3) serious infection; 4) severe hematologic toxicity/graft failure. Response is defined as the event of CR, that is, $< 6\%$ bone marrow blasts and evidence of donor chimerism ($> 80\%$) by DNA microsatellite polymorphism analysis. Development of drug-related Grade 3 or 4 organ toxicity or severe infection led to azacitidine discontinuation. Azacitidine was also discontinued if platelet count dropped to $< 10,000/\text{mm}^3$, with 50% dose reduction if platelet count dropped to $< 20,000/\text{mm}^3$.

4.2.2 Non-nested Schedule - Treating metastatic colorectal cancer

Irinotecan was the first drug to improve survival beyond that achieved with standard first-line treatment of colorectal cancer. Combining irinotecan with oral capecitabine is an interesting alternative in view of the practicability of the treatment. However, there is no consensus on the dose and schedule of irinotecan in presence of standard treatment of oral capecitabine (Borner et al. 2008). Two non-nested schedules are suggested: weekly irinotecan on days $1, 8, 15, 22, 29$ or tri-weekly irinotecan on day 1 and days 22 every six weeks for six cycles. Doses for the weekly schedule are $35, 70$ and 90 mg/m^2 and doses for the tri-weekly schedule are $180, 240$ and 300 mg/m^2. Hence, there are a total of six dose-schedule treatment regimes under investigation. Toxicity is defined as any of the Grade 3/4 toxicity according to the National Cancer Institute. Response is defined as any of a complete response (CR), partial response (PR) or minor response (at least a 25% decrease in tumor size) reviewed by an independent radiology panel. There is an intra-patient dose reduction rule stating that doses should be reduced by 25% in the subsequent treatment cycles in case of Grade 3 non-hematological toxicity or by 50% in case of Grade 4 hematological toxicity.

4.3 Methods

4.3.1 Preliminary Notation

A typical dose-schedule-finding trial involving nested schedules aims to identify an optimal treatment regimen within a $J \times K$ matrix consisting of J doses and K nested schedules. Patients receive administrations under schedule $k, k = 1, \ldots, K$ at times $\mathbf{s}^{(k)} = (s_1, s_2, \ldots, s_{m_{(k)}})$ where $\mathbf{s}^{(1)} \subset \mathbf{s}^{(2)} \subset \ldots \subset \mathbf{s}^{(K)}$ and $m_{(1)} < m_{(2)} < \ldots < m_{(K)}$. In the azacitidine trial, there are $J = 3$ doses and $K = 4$ nested schedules. A single course of treatment, includes 5 administrations on first 5 consecutive days and the administration times are denoted by $\mathbf{s}^{(1)} = (0, 1, 2, 3, 4)$. Patients receive the second treatment course after 28 days of rest, hence $\mathbf{s}^{(2)} = (0, 1, 2, 3, 4, 28, 29, 30, 31, 32) = (\mathbf{s}^{(1)}, \mathbf{s}^{(1)} + 28)$ and so on. Let d_j denote the per-administration dose, $j = 1, \ldots, J$, then s^k and d_j denote the possible dose and schedule regimes to be initially assigned to a patient and $d_j \in (8, 16, 32)$ mg/m^2. However, due to possible intra-patient dose modification, the patient may not receive the same dose d_j for each administration. Hence, patient i may receive different doses $\mathbf{d}_i = (d_{i,1}, \ldots, d_{i,m_i})$ at administration times $\mathbf{s}_i = (s_{i,1}, \ldots, s_{i,m_i})$ where m_i is the number of administrations patient i has received.

Non-nested schedules, in contrast, are qualitatively different from each other and do not consist of the same treatment course. In the irinotecan trial, $K = 2$ non-nested schedules are under investigation. The administration times are denoted as $\mathbf{s}_0^1 = (0, 7, 14, 21, 28)$ for a single treatment cycle using the weekly schedule and $\mathbf{s}_0^2 = (0, 21)$ for a single treatment cycle using the tri-weekly schedule. Since the investigator plans six cycles, the administration times for the tri-weekly schedule are $\mathbf{s}^{(2)} = (0, 21, 42, 63, 84, 105, 126, 147, 168, 189, 210, 231) = (\mathbf{s}_0^2, \mathbf{s}_0^2+42, \mathbf{s}_0^2+84, \mathbf{s}_0^2+126, \mathbf{s}_0^2+168, \mathbf{s}_0^2+210)$. Similarly, for the weekly schedule $s^{(1)} = (\mathbf{s}_0^1, \mathbf{s}_0^1+42, \mathbf{s}_0^1+84, \mathbf{s}_0^1+126, \mathbf{s}_0^1+168, \mathbf{s}_0^1+210)$. Contrary to nested schedules, the set of doses under investigation for each schedule

could be different. Let d_j^1 denote the per-administration dose j for the weekly schedule and d_j^2 denote for the corresponding dose for the tri-weekly schedule, where $j = 1, \ldots, J$, the total number of doses considered for each schedule. In our example, J =3, $d_j^1 \in (35, 70, 90)$ mg/m^2 and $d_j^2 \in (180, 240, 300)$ mg/m^2. Again, even though the dose and schedule regimen $\mathbf{s}^k$ and $d_j^k, k = 1, \ldots, K$, might be assigned to patient i initially, she might experience dose modification for later treatment cycles. We still use $\mathbf{d}_i$ and $\mathbf{s}_i$ to denote the dose and schedule combination one actually receives.

Regardless of either nested or non-nested schedules, a patient will be followed for a maximum duration of ω_T and ω_E for evaluating whether the patient experiences a toxicity or response, respectively. The values ω_T and ω_E are often determined by the clinician and also known as the duration of the evaluation window for toxicity and efficacy.

4.3.2 Marginal Model for Time-to-Event After a Single Administration

Let c = T or R be an indicator for the toxicity outcome (T) or efficacy outcome (R). We expect the hazard function for either outcome after a single administration to increase with time to certain time and then attenuate as we want to mimic the biologic of the agent (Liu and Braun, 2009). We define the hazard function at time ν for event c as $g_c(\nu|\theta_c, \phi_c) = \theta_c f(\nu|\phi_c)$ where θ_c is a positive scale parameter and $f(\nu|\phi_c)$ is a probability density function with a certain restriction on the parameter ϕ_c such that the mode of $f(\cdot)$ exists. Let $F(\nu|\phi_c)$ denote the cumulative density function of $f(\cdot)$. The above model corresponds to the non-mixture cure rate model proposed by Chen, Ibrahim, and Sinha (1999) since the survival function $S(\nu|\theta_c, \phi_c) = \exp[-\theta_c F(\nu|\phi_c)]$ is not proper and $S(\infty|\theta_c, \phi_c) = \exp(-\theta_c) > 0$ is the cure probability, i.e., the proportion of the population who would never experience toxicity (c=T) or response (c=R). Therefore, $1 - \exp(-\theta_c)$ is the toxicity or response rate after a single administration of a dose with infinite follow-up.

Chen et al. (1999) also suggested modeling covariate effects through θ_c to maintain

a proportional hazards structure. To make the concept distinct, we denote $\theta_{i,l,c}$ as the cure parameter associated with dose $d_{i,l}$ for administration l to patient i. Since $\theta_{i,l,c} > 0$, we model it as a function of the dose $d_{i,l}$ for administration l through $\log(\theta_{i,l,c}) = \beta_{0,c} + exp(\beta_{1,c})d_{i,l}$, where the slope is positive because it is reasonable to assume that toxicity and response rates increase with dose. However, if the monotonicity assumption is not valid, one option is to use a second order polynomial. We choose $f(\cdot)$ to be the Weibull density function defined in Liu and Braun (2009): $f(\nu_{i,l}|\boldsymbol{\phi}_c) = \exp(-\gamma_c)\alpha_c\nu_{i,l}^{\alpha_c-1}\exp\left[-\nu_{i,l}^{\alpha_c}\exp(-\gamma_c)\right]$ where $\boldsymbol{\phi}_c = (\alpha_c, \gamma_c)$ and we require $\alpha_c > 1$ so the mode of the density exists. Therefore, the respective hazard and survival functions for patient i after administration l at dose $d_{i,j}$ are $g(\nu_{i,l}|\boldsymbol{\beta}_c, \boldsymbol{\phi}_c) = \theta_{i,l,c}f(\nu_{i,l}|\boldsymbol{\phi}_c)$ and $S(\nu_{i,l}|\boldsymbol{\beta}_c, \boldsymbol{\phi}_c) = \exp[-\theta_{i,l,c}F(\nu_{i,l}|\boldsymbol{\phi}_c)]$, in which $\boldsymbol{\beta}_c = (\beta_{0,c}, \beta_{1,c})$. More details about the model assumptions have been presented in Chapter III.

4.3.3 Marginal Model for Time-to-Event After Multiple Administrations

Assume patient i receives m_i administrations at administration times ($s_{i,1}$, $s_{i,2}$, ..., s_{i,m_i}) and the respective per-administration doses are ($d_{i,1}$, $d_{i,2}$, ..., d_{i,m_i}). By employing a competing risks approach, we assume that there is a latent event time $Y_{i,l}$ after each administration, and the patient time when patient i experiences an event is $\min(Y_{i,1} + s_{i,1}, \ldots, Y_{i,m_i} + s_{i,m_i})$. At patient time t_i, the survival function for Y_i is given by

$$\phi(t_i|\boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}_i, \mathbf{d}_i) = P(Y_i > t_i) = P(Y_{i,1} > y_{i,1}, \ldots, Y_{i,m_i} > y_{i,m_i})$$

where $y_{i,l} = t_i - s_{i,l}$.

However, we do not assume $Y_{i,1}, \ldots, Y_{i,m_i}$ are independent since they are from the same patient and we may want to consider within-patient correlation, which we did not take into account in Chapter III. We introduce a frailty variable $w_i > 0$, conditional on which $Y_{i,1}, \ldots, Y_{i,m_i}$ are independent. Specifically, after consider-

ing bivariate outcomes (c=T or R), the hazard and survival functions after administration l of multiple administrations are $g(\nu_{i,l}|\boldsymbol{\beta}_c, \boldsymbol{\phi}_c, w_i) = w_i\theta_{i,l,c}f(\nu_{i,l}|\boldsymbol{\phi}_c)$ and $S(\nu_{i,l}|\boldsymbol{\beta}_c, \boldsymbol{\phi}_c, w_i) = \exp[-w_i\theta_{i,l,c}F(\nu_{i,l}|\boldsymbol{\phi}_c)]$. While there are several choices for the distribution of w_i, the stable distribution is commonly used and expressed in the form

$$z(w|\lambda) = \frac{\lambda}{1-\lambda}w^{-1/(1-\lambda)}\int_0^1 J(\mu)\exp[-\frac{J(\mu)}{w^{\frac{\lambda}{1-\lambda}}}]d\mu,$$

where

$$J(\mu) = (\frac{\sin(\lambda\pi\mu)}{\sin(\pi\mu)})^{\frac{\lambda}{1-\lambda}}(\frac{\pi\mu(1-\lambda)}{\sin(\pi\mu)}).$$

Using the Laplace transform of w_i, considering bivariate outcomes (c =T, R) and following the derivation of Chen, Ibrahim, and Sinha (2002), one can derive the marginal survival function for the patient event time $Y_{i,c}$ as

$$\begin{aligned}\varphi_c(t_{c,i}|\boldsymbol{\beta}_c, \boldsymbol{\phi}_c, \lambda_c, \mathbf{s}_{c,i}, \mathbf{d}_{c,i}) &= \int_0^\infty P(Y_{i,1} > y_{i,1}, \ldots, Y_{i,m_i} > y_{i,m_i}|w_i)z(w_i|\lambda)\,\mathrm{d}w_i \\ &= \exp\left\{-\left[\sum_{l=1}^{m_{c,i}}\theta_{c,i,l}F\left(y_{c,i,l}|\boldsymbol{\phi}_c\right)\right]^{\lambda_c}\right\} \qquad (4.1)\end{aligned}$$

where $\log(\theta_{i,l,c}) = \beta_{0,c} + \exp(\beta_{1,c})d_{i,l}$. Note that the number of administrations $m_{i,c}$ could be different when considering different outcomes, since toxicity and response could occur at different times. $\lambda_c > 0$ is a correlation parameter that quantifies the amount of within-patient correlation and $\lambda_c = 1$ implies that there is no within-patient correlation.

The toxicity and response rate $\pi_c(j,k)$ for a dose-schedule combination (j, k) within the duration of the respective evaluation windows ω_c is then given by $\pi_c(j,k) = 1 - \varphi_c(\omega_c)$ where the parameters are omitted for brevity.

4.3.4 Joint Model, Likelihood and Posterior

In order to take into account the possible correlation between times to toxicity and efficacy, we construct the joint bivariate survival function by employing an Archimedean copula model, which specifies a well-defined relationship between marginal and joint distributions. We adopt the Clayton copula model (Clayton 1978; Yuan and Ying 2009) because it is simple and easy to interpret. Specifically, the bivariate survival function is given by

$$\begin{aligned} &S(t_{T,i}, t_{E,i}|\zeta, \boldsymbol{\beta}_T, \boldsymbol{\phi}_T, \lambda_T, \boldsymbol{\beta}_E, \boldsymbol{\phi}_E, \lambda_E, \mathbf{s}_i, \mathbf{d}_i) \\ &= \left[\varphi_T(t_{T,i}|\boldsymbol{\beta}_T, \boldsymbol{\phi}_T, \lambda_T, \mathbf{s}_{T,i}, \mathbf{d}_{T,i})^{-\zeta} + \varphi_E(t_{E,i}|\boldsymbol{\beta}_E, \boldsymbol{\phi}_E, \lambda_E, \mathbf{s}_{E,i}, \mathbf{d}_{E,i})^{-\zeta} - 1\right]^{-1/\zeta} \quad (4.2) \end{aligned}$$

Note that $\zeta > 0$ is a correlation parameter that defines the amount of correlation between times to toxicity and response. There is no correlation when $\zeta = 0$ since Kendall's tau that measures bivariate association is equal to $\zeta/(2+\zeta)$.

For toxicity, define *patient* time $Y_{T,i} = \min(Y_{T,i}, U_{T,i})$ and $\delta_{T,i} = I(Y_{T,i} \leq U_{T,i})$, where $U_{T,i}$ denotes the censoring time and $I(\cdot)$ is the indicator function. Hence, we observe toxicity for patient i if $C_i = 1$. Similarly, $Y_{E,i}$ and $\delta_{E,i}$ are defined for efficacy. Since interim analyses are performed whenever a new cohort is enrolled, by the time cohort $n+1$ is enrolled, we denote the number of patients currently in the study as n and for each enrolled patient, we observe $y_{T,i}$, $\delta_{T,i}$, $y_{E,i}$, $\delta_{E,i}$, $\mathbf{s}_i$ and $\mathbf{d}_i$, where $\mathbf{s}_i$ and $\mathbf{d}_i$, as we defined previously, are the respective time and dose for each administration patient i has received, $i = 1, \ldots, n$. Note $U_i = \min(W_{i,n+1}, \omega)$, where $W_{i,n+1}$ is the inter-patient time between patient i and $n+1$. If $n = N$, the maximum number of patients for the study, then define $W_{i,n+1}$ as the time between patient i and the end of the study. It is reasonable to assume random censoring since $W_{i,n+1}$ is usually independent of the DLT time and ω is a fixed value. Based on the above information and (4.2), the likelihood of the data $\mathcal{D}_i = (y_{T,i}, \delta_{T,i}, y_{E,i}, \delta_{E,i}, \mathbf{s}_i, \mathbf{d}_i)$ for patient i is

given by

$$\mathcal{L}(\boldsymbol{\tau}|\mathcal{D}_i) = L_1^{\delta_{T,i}\delta_{E,i}} L_2^{\delta_{T,i}(1-\delta_{E,i})} L_3^{(1-\delta_{T,i})\delta_{E,i}} L_4^{(1-\delta_{T,i})(1-\delta_{E,i})},$$

where,

$$\begin{aligned}
L_1 &= \partial^2 S(y_{T,i}, y_{E,i}|\boldsymbol{\tau}, \mathbf{s}_i, \mathbf{d}_i)/\partial y_{T,i}\partial y_{E,i},\\
L_2 &= -\partial S(y_{T,i}, y_{E,i}|\boldsymbol{\tau}, \mathbf{s}_i, \mathbf{d}_i)/\partial y_{T,i},\\
L_3 &= -\partial S(y_{T,i}, y_{E,i}|\boldsymbol{\tau}, \mathbf{s}_i, \mathbf{d}_i)/\partial y_{E,i},\\
L_4 &= S(y_{T,i}, y_{E,i}|\boldsymbol{\tau}, \mathbf{s}_i, \mathbf{d}_i).
\end{aligned}$$

Let $\boldsymbol{\tau} = (\zeta, \boldsymbol{\beta}_T, \boldsymbol{\phi}_T, \lambda_T, \boldsymbol{\beta}_E, \boldsymbol{\phi}_E, \lambda_E)$. After determining the prior distribution $p(\boldsymbol{\tau})$, the posterior distribution of $\boldsymbol{\tau}$ based on $\mathcal{D} = \{\mathcal{D}_i : i = 1, ..., n\}$ is

$$p(\boldsymbol{\tau}|\mathcal{D}_i) \propto p(\boldsymbol{\tau}) \prod_{i=1}^{n} \mathcal{L}(\boldsymbol{\tau}|\mathcal{D}_i)$$

We can compute posterior quantities via adaptive Markov Chain Monte Carlo (MCMC) methods (Rosenthal 2007). Those posterior quantities will be used to identify the dose-schedule assignment for a new cohort of patients.

4.3.5 Establishing Prior Distributions

For the parameters $\boldsymbol{\beta_c}$ and λ_c (c=T or R), we assign independent Gaussian distributions with prior mean and prior variance $(\mu_0, \sigma^2_{0,c})$ for $\beta_{0,c}$, $(\mu_{1,c}, \sigma^2_{1,c})$ for $\beta_{1,c}$ and $(\mu_{3,c}, \sigma^2_{2,c})$ for $\log(\lambda_c)$. The means are elicited from the investigators who are asked to provide the "skeleton" $\mathcal{P}$, the *a priori* estimates of the toxicity and response rates for all dose-schedule combinations, in which element (j, k), denoted P_{jk}, corresponds to the toxicity rate or response rate of dose j and schedule k treatment regimen. We then fit the linear regression model $\log(-\log[1-P_{jk,c}]) = \lambda_c \log(m_{k,c}) + \lambda_c b_{0,c} + \lambda_c \exp(b_{1,c}) d_j$

(d_j is replaced by d_j^k for non-nested schedules) derived from Equation (4.1) and use the ordinary least square estimates $\hat{b}_{0,c}$, $\hat{b}_{1,c}$ and $\hat{\lambda}_c$ as the respective values for $\mu_{0,c}$, $\mu_{1,c}$ and $\mu_{2,c}$.

For the parameters of the Weibull density in the hazard, $\boldsymbol{\phi_c} = (\alpha_c, \gamma_c)$, we adopt the same approach outlined in Chapter III and in Liu and Braun (2009). Specifically, we keep α_c fixed to maintain a parsimonious model and limit the number of parameters to estimate. In addition, we assign a Gaussian prior distribution for γ_c with mean $\mu_{\gamma,c}$ and variance $\sigma^2_{\gamma,c}$.

For the correlation parameter ζ for the bivariate time-to-event outcomes, we re-parameterize $\zeta = (\zeta_s)^2$ and assign a Gaussian distribution with mean μ_ζ and variance σ^2_ζ to ζ_s. We choose μ_ζ to be 0.

As in Chapter III, we recommend calibrating the prior variances through simulations using a few different skeletons and prior variances. The prior variance that is the most insensitive to skeletons and leads to the best operating characteristics will be used for a real trial.

4.4 Trial conduct

4.4.1 Decision Criteria

We adopt the decision criteria in terms of posterior probabilities given $\mathcal{D}$ at any interim analysis in the trial similar to Thall and Russell (1998) and many other Phase I/II trial designs. Given upper probability cutoffs p_T and p_E, we consider a dose-schedule combination, denoted by (j,k), to be a regimen with acceptably low toxicity rate if

$$P\{\pi_T(j,k) < \overline{\pi}_T | \mathcal{D}\} > p_T \tag{4.3}$$

and to be a regimen with acceptably high response rate if

$$P\{\pi_E(j,k) > \underline{\pi}_E|\mathcal{D}\} > p_E, \tag{4.4}$$

where $\overline{\pi}_T$ and $\underline{\pi}_E$ are a fixed upper limit on toxicity rate and lower limit on response rate respectively. A dose-schedule regimen is acceptable if it satisfies both (4.3) and (4.4). We denote the set of acceptable strategies based on $\mathcal{D}$ to be $\mathcal{A}(\mathcal{D})$. In order to have good operating characteristics for the design, the values of cut-off values p_T and p_E should be carefully calibrated via simulations prior to the onset of a trial.

The optimal regimen was defined to be the one in $\mathcal{A}(\mathcal{D})$ that maximizes the posterior response probability in (4.4), a rule proposed by Thall and Russell (1998). However, the above criteria might not work well when $\pi_E(j,k)$ is much above p_T for several dose-schedule regime because the probabilities in (4.4) would be very close to 1 and cause some numerical problems, since there would be more than one optimal treatment regime. Hence, we further propose the following additional criteria for accepting a dose-schedule regimen:

$$\hat{\pi}_T(j,k) < \overline{\pi}_T + \Delta_T \tag{4.5}$$

for toxicity and

$$\hat{\pi}_E(j,k) > \overline{\pi}_E - \Delta_E \tag{4.6}$$

for efficacy. $\hat{\pi}_T(j,k)$ and $\hat{\pi}_E(j,k)$ are posterior mean of the toxicity and response rates; Δ_T and Δ_E are cut-off values and usually assumed to be equal. We further re-define the set $\mathcal{A}(\mathcal{D})$ to be those strategies that satisfy (4.3) to (4.6).

4.4.2 Trial Conduct for Nested Schedules

We plan on enrolling a maximum of N patients in the trial in V cohorts with n_0 patients per cohort, and each patient will be followed for ω_T and ω_E days after enrollment for toxicity and efficacy respectively. The first cohort of patients are enrolled at study time $t = 0$ and assigned to the shortest schedule ($k = 1$) with the lowest dose ($j = 1$) or the dose-schedule regimen specified by the clinician. When cohort $i = 2, \ldots, V$ is to be enrolled in the study at study time t, we perform the following steps:

(a) Record the observed data $\mathcal{D}_{i'}$ for each enrolled patient $i' = 1, 2, \ldots, (i-1)n_0$;

(b) Compute the likelihood and draw samples from posterior distributions;

(c) Similar to Thall and Cook (2004), dose-schedule combination $(j, k) \in \mathcal{A}(\mathcal{D})$ if (j, k) satisfies (4.3) to (4.6), or if (j, k) is the lowest untried regimen above the starting regimen and it satisfies (4.3) and (4.5);

(d) If $\mathcal{A}(\mathcal{D}) \neq \phi$, then the next cohort is treated at the optimal regimen in $\mathcal{A}(\mathcal{D})$, subject to the constraint that no untried dose or schedule can be skipped when escalating;

(e) If $\mathcal{A}(\mathcal{D}) = \phi$, then consider the following conditions:

  (i) If (4.3) or (4.5) is not satisfied, then the trial is terminated and no dose-schedule regimen is selected;

  (ii) if it is the first time that (4.3) and (4.5) are satisfied but not for (4.4) or (4.6), then treat the next cohort at the regimen that maximizes the estimated response rate subject to the same no-skipping rule in (c); Otherwise if it is not the first time, then terminate the trial and no regimen is selected.

(f) If all N patients have been enrolled and the trial is not stopped early, then select the combination in $\mathcal{A}(\mathcal{D}_N)$ that maximizes the estimated response rate (we use its posterior mean) in (4.6).

4.4.3 Trial Conduct for Non-nested Schedules

In contrast with nested schedules, the ordering in terms of either toxicity or response rates of two dose-schedule regime with two different non-nested schedules are usually unknown. Hence, in order to have good operating characteristics, some aspects of the algorithm used for nested schedules need to be modified even though the rest remains the same. Firstly, the lowest dose-schedule combination is unknown and hence the regimen to be assigned to the first cohort has to be specified by the clinician. Secondly, we introduce adaptive randomization when assigning the a dose-schedule regimen to the next cohort rather than assigning the optimal regimen to the next cohort in the steps (d) and (e)(ii) with probability 1. The randomization is done as follows:

(i) Compute the estimated response rates $\hat{\pi}_E(j,k)$ for all dose-schedule regime in $\mathcal{A}(\mathcal{D})$;

(ii) For each schedule $k = 1, \ldots, K$, find the $j^{opt}(k) = \underset{j=1,2,\ldots,J}{\operatorname{argmax}}\ \hat{\pi}_E(j,k)$;

(iii) Treat the next cohort at the regimen $(j^{opt}(k), k)$ with the randomization probability

$$r(j^{opt}(k),k) = \frac{\hat{\pi}_E(j^{opt}(k),k)}{\sum_{k'=1}^{K} \hat{\pi}_E(j^{opt}(k'),k')}$$

where $\hat{\pi}_E(j^{opt}(k') = 0$ if $j^{opt}(k') = \phi$.

4.5 Application - The Azacitidine Trial

4.5.1 Simulation Design

In the motivating azacitidine trial, the clinicians specified $J = 3$ doses of interest: 8, 16 and 24 mg/m^2, and $K = 4$ nested schedules, for a total of 12 combinations. A course consists of five daily consecutive administrations followed by 28 days of rest as described in the example in Section 4.3.1, and schedule k consists of k consecutive courses. The maximum follow-up time for each patient is $\omega_T = \omega_E = 116$ days. A maximum of $N = 60$ patients will be enrolled with the cohort size 3. The algorithm was implemented with $p_T = 0.4, p_E = 0.1, \pi_T = 0.35$,$\pi_E = 0.35$ and $\Delta_T = \Delta_E = 0.05$ starting at the lowest dose and schedule (dose 1 and schedule 1). Investigators would like to determine which of the 12 combinations is the optimal treatment regimen in terms of both toxicity and efficacy.

We asked the investigator to provide the estimates of the toxicity and response rates for each of the 12 dose-schedule combinations. The skeleton specifies the middle combinations to be the optimal regimes. The actual values of the skeleton can be found in Table 4.2. We used the methods described in Section 4.3.5 to calculate the prior means that are shown in Table 4.1. We calibrated the prior variances of our model parameters through a process outlined in Chapter III. We assumed the variances for all the prior distributions are equal and tried three different prior standard deviations, 1, 2 and 4. We selected prior standard deviation(SD)=2 for the trial.

We examined our approach in eight different scenarios that are summarized in Table 4.2 in which the true toxicity and response rates are denoted by a pair (π_T, π_E). The true toxicity and response rates of every combination of dose and schedule were generated by the model used in our methods with the frailty following the positive stable distribution defined in Section 4.3.3. We simulated patients to have exponen-

Table 4.1: Means and standard deviations of the prior normal distributions for each parameter in the azacitidine trial.

Parameter	Mean	Sd
$\beta_{0,T}$	-5.02	2
$\beta_{1,T}$	-0.76	2
$\log(\lambda_T)$	-0.43	2
α_T	1.87	0
γ_T	1.27	2
$\beta_{0,E}$	-5.29	2
$\beta_{1,E}$	-0.55	2
$\log(\lambda_E)$	-0.73	2
α_E	1.88	0
γ_E	1.27	2
ζ_s	0	2

tially distributed inter-arrival times with a mean of two weeks, and we divided all the follow-up times by 10 to achieve better numerical stability for our model. When a new cohort is enrolled, an interim analysis is performed in which a single chain of $6,000$ samples, after a burn-in of $4,000$ samples, is drawn from the posterior distribution for each parameter. These posterior draws are then used to determine the dose and schedule assigned to the new patient. We evaluated the performance of our approach by comparing the correct selection frequency at the end of the study and the mean proportion of patients assigned to each dose-schedule combination. We performed 500 simulations in each scenario; our computer code is available upon request.

4.5.2 Simulation Results

Table 4.3 contains a summary of the performance of our design in the eight scenarios described in Table 4.2. Columns 3-6 describe the proportion of simulations in which each of the dose-schedule combination is selected as the optimal regimen at the end of the study. Column 7 describes the proportion of simulations that are stopped early. The next three columns have a similar interpretation related to the average

Table 4.2: The true toxicity and response rates and the skeleton used in the simulation study. The boldfaced values correspond to acceptable combinations.

		Schedule (π_T, π_E)			
Sc.	Dose	1	2	3	4
1	1	(0.54,0.41)	(0.71,0.55)	(0.80,0.63)	(0.86,0.69)
	2	(0.56,0.43)	(0.73,0.57)	(0.82,0.66)	(0.88,0.71)
	3	(0.59,0.45)	(0.75,0.59)	(0.84,0.68)	(0.89,0.73)
2	1	(0.03,0.02)	(0.05,0.04)	(0.07,0.05)	(0.09,0.06)
	2	(0.03,0.02)	(0.06,0.04)	(0.09,0.06)	(0.11,0.08)
	3	(0.04,0.03)	(0.08,0.05)	(0.11,0.07)	(0.14,0.09)
3	1	(0.06,0.18)	(0.11,0.26)	(0.16,0.32)	**(0.19,0.36)**
	2	(0.11,0.26)	**(0.19,0.36)**	**(0.27,0.44)**	**(0.33,0.49)**
	3	**(0.19,0.36)**	**(0.32,0.49)**	(0.43,0.58)	(0.51,0.64)
4	1	(0.05,0.10)	(0.07,0.18)	(0.1,0.24)	(0.11,0.27)
	2	(0.08,0.14)	(0.12,0.26)	(0.16,0.34)	**(0.18,0.39)**
	3	(0.12,0.21)	**(0.20,0.36)**	**(0.25,0.48)**	**(0.29,0.53)**
5	1	(0.14,0.28)	**(0.20,0.40)**	**(0.25,0.48)**	**(0.29,0.55)**
	2	(0.28,0.30)	(0.39,0.43)	(0.46,0.52)	(0.52,0.58)
	3	(0.50,0.33)	(0.65,0.46)	(0.74,0.55)	(0.80,0.62)
6	1	(0.13,0.14)	(0.21,0.20)	(0.28,0.24)	(0.33,0.27)
	2	(0.20,0.28)	**(0.32,0.39)**	(0.41,0.46)	(0.48,0.49)
	3	**(0.30,0.50)**	(0.46,0.65)	(0.58,0.74)	(0.65,0.77)
7	1	**(0.26,0.54)**	**(0.38,0.59)**	(0.46,0.62)	(0.51,0.64)
	2	**(0.34,0.55)**	(0.48,0.60)	(0.57,0.63)	(0.62,0.65)
	3	(0.44,0.55)	(0.60,0.60)	(0.69,0.63)	(0.74,0.65)
8	1	(0.05,0.07)	(0.10,0.13)	(0.13,0.18)	(0.16,0.21)
	2	(0.10,0.20)	(0.18,0.33)	**(0.24,0.44)**	**(0.29,0.51)**
	3	**(0.18,0.47)**	**(0.32,0.70)**	(0.42,0.82)	(0.48,0.88)
Skeleton	1	(0.13,0.18)	(0.21,0.28)	(0.28,0.33)	(0.32,0.35)
	2	(0.14,0.28)	(0.30,0.33)	(0.34,0.40)	(0.38,0.44)
	3	(0.30,0.33)	(0.35,0.44)	(0.40,0.48)	(0.45,0.53)

percentage of dose-schedule assignments during the study.

Overall, we are able to identify acceptable dose-schedule combinations at the end of the study in a majority of simulations and stop the trial early if there is no regimen that is both safe and efficacious. In Scenario 1, all the regimes are overly toxic because all of their toxicity rates are above the upper limit $\overline{\pi}_T$. We terminated the trial in 100% of the simulations. In Scenario 2, none of the regimes is efficacious enough because all of the response rates are below the lower limit $\underline{\pi}_E$. The trial was terminated in all the simulations. In Scenarios 3-8, we are able to identify the acceptable and the optimal regimes with high probability. For example, in Scenario 5, the combinations of dose 1 and one of schedules 2-4 are acceptable treatment regimes. But the combinations (dose 1, schedule 3) and (dose 1, schedule 4) have much higher response rates but still relatively low toxicity rates; hence they could be considered as the optimal combinations in this scenario. We select them 15% and 65% respectively at the end of trial. Scenario 6 might be the most difficult for dose-schedule finding among the scenarios we explored, because there are only two acceptable dose-schedule treatment regimes and their toxicity rates are very close to $\overline{\pi}_T$ but their response rates are not much higher than $\underline{\pi}_E$. The trial was stopped early in 10% of the simulations due to the difficulty in finding the correct treatment regimes. Even in this scenario, we selected the acceptable regimes in 50% of the simulations. In Scenario 7, the trial was stopped with the probability 13%. This is because most of the combinations are overly toxic and only two are acceptable ones.

In terms of patient assignment, we find that a high percentage of patients were treated at the acceptable dose-schedule regimes. In Scenario 1, we treated 20% of the total patients before the trial was stopped early due to safety and most of the enrolled patients were treated at the lowest dose and schedule combination. In Scenario 2, we enrolled 35% of the total patients before the trial was terminated due to futility and most of treated patients were on the highest dose-schedule regimen. In Scenarios

Table 4.3: Simulation results for the azacitidine trial with sample size of $N = 60$. Columns "% Sel" gives the percentage of identifying each dose-schedule combination as the optimal treatment regimen. Column " None" gives the percentage of simulations that were stopped early. Columns "% Pat" gives the mean proportion of patients assigned to each dose-schedule regimen. The boldfaced values correspond to acceptable combinations.

		% Sel					% Pat			
		Schedule					Schedule			
Scenario	Dose	1	2	3	4	None	1	2	3	4
1	1	0	0	0	0	**100**	11	2	1	3
	2	0	0	0	0		2	1	0	1
	3	0	0	0	0		0	0	0	0
2	1	0	0	0	0	**100**	0	0	0	6
	2	0	0	0	0		0	0	0	5
	3	0	0	0	0		0	0	0	24
3	1	0	0	1	**10**	1	0	1	1	**16**
	2	0	**4**	**17**	**33**		1	**2**	**6**	**31**
	3	**8**	**17**	6	2		**3**	**8**	9	22
4	1	0	0	0	0	1	0	0	0	6
	2	0	0	0	**5**		0	0	0	**10**
	3	0	**1**	**6**	**87**		0	**1**	**2**	**80**
5	1	2	**6**	**15**	**65**	2	2	**5**	**8**	**51**
	2	3	3	2	1		4	6	6	8
	3	1	0	0	0		3	2	3	3
6	1	0	3	5	11	10	1	3	4	17
	2	14	**14**	4	0		8	**10**	6	9
	3	**36**	2	0	0		**19**	9	4	6
7	1	**52**	**7**	1	0	13	**29**	**14**	6	9
	2	**19**	1	0	0		**14**	5	3	3
	3	6	0	0	0		5	2	1	1
8	1	0	0	0	1	0	0	0	0	8
	2	0	2	**6**	**12**		0	2	**3**	**17**
	3	**26**	**39**	9	5		**9**	**19**	14	27

3-8, we were able to treat many patients at the acceptable regimes: 66%, 93%, 64%, 29%, 57% and 48% respectively. At the same time, we did not expose many patients to overly toxic combinations. For example, if we consider the regimes with toxicity rate over $\overline{\pi}_T + 0.05$ as overly toxic, the percentage of overdosed patients are 20%, 0%, 22%, 0%, 25%, 34%, 31% and 27% in Scenarios 1-8 respectively.

We calibrated the variances of the prior distributions via simulations. Table 4.4 summarizes the simulation results using prior SD = 1 and 4. We note that the results are quite similar for prior SD =1 and 2. However, the results from using prior SD=4 differ greatly from others because we note more simulations were stopped early. Hence, we chose prior SD= 2 because our design selected the correct regimen slightly more often in Scenario 5 however quite similar in other scenarios.

We examined the sensitivity of our design to model misspecification in Table 4.5. Our model assumes a positive stable frailty, so we evaluated our method when the frailty actually follows an Inverse Gaussian distribution with mean 1 and variance adjusted to make the toxicity and response rates for each scenarios match Table 4.3. Overall, we find that our design's ability to identify and assign the correct dose-schedule regimes is not compromised, although we might assign slightly more patients to toxic treatment regimes.

We also examined the sensitivity of our design to the maximum sample size by repeating our simulations using sample sizes of 48 and 72, the results of which are summarized in Table 4.6. The results show that decreasing the sample size from 60 to 48 may result in a nontrivial loss is selecting and assigning correct combinations. However, increasing the sample size from 60 to 72 only produces a minor increase in the selection and assignment of correct combinations. Therefore, we selected a maximum number of 60 patients for this study.

Table 4.4: Simulation results for the azacitidine trial with sample size of $N = 60$ using different prior variances. Columns "% Sel" gives the percentage of identifying each dose-schedule combination as the optimal treatment regimen. Column " None" gives the percentage of simulations that were stopped early. Columns "% Pat" gives the mean proportion of patients assigned to each dose-schedule regimen. The boldfaced values correspond to acceptable combinations.

			% Sel					% Pat			
			Schedule					Schedule			
Prior Sd	Scenario	Dose	1	2	3	4	None	1	2	3	4
1	1	1	0.2	0	0	0	**99.8**	16	4	2	2
		2	0	0	0	0		2	1	0	1
		3	0	0	0	0		0	0	0	0
	2	1	0	0	0	0	**100**	0	0	0	6
		2	0	0	0	0		0	0	0	6
		3	0	0	0	0		0	0	0	25
	5	1	1	**9**	**18**	**59**	1	1	**4**	**9**	**50**
		2	2	6	3	1		3	7	7	9
		3	1	0	0	0		4	3	3	1
	6	1	0	1	5	14	8	0	2	4	19
		2	13	**15**	3	0		6	**10**	8	10
		3	**37**	3	0	0		**20**	10	5	4
	7	1	**59**	**16**	4	1	2	**24**	**20**	10	13
		2	**16**	1	0	0		**12**	6	3	4
		3	1	0	0	0		4	2	1	1
4	1	1	0.2	0	0	0	**99.8**	10	2	0	3
		2	0	0	0	0		1	1	0	1
		3	0	0	0	0		0	0	0	0
	2	1	0	0	0	0	**100**	0	0	0	6
		2	0	0	0	0		0	0	0	6
		3	0	0	0	0		0	1	1	21
	5	1	3	**7**	**12**	**63**	4	3	**6**	**7**	**48**
		2	3	2	2	1		4	5	5	8
		3	1	0	0	0		4	2	2	4
	6	1	0	5	3	9	12	3	5	4	14
		2	18	**13**	3	1		9	**10**	5	7
		3	**31**	4	0	0		**20**	9	4	7
	7	1	**37**	**6**	1	1	27	**26**	**10**	4	7
		2	**16**	0	0	0		**13**	4	2	3
		3	12	0	0	0		9	2	1	1

Table 4.5: Simulation results for the azacitidine trial under model mis-specification with sample size of $N = 60$. Columns "% Sel" gives the percentage of identifying each dose-schedule combination as the optimal treatment regimen. Column " None" gives the percentage of simulations that were stopped early. Columns "% Pat" gives the mean proportion of patients assigned to each dose-schedule regimen. The boldfaced values correspond to acceptable combinations.

		% Sel					% Pat			
		Schedule					Schedule			
Scenario	Dose	1	2	3	4	None	1	2	3	4
1	1	0.2	0	0	0	**99.8**	13	3	2	2
	2	0	0	0	0		2	1	0	0
	3	0	0	0	0		0	0	0	0
2	1	0	0	0	0	**100**	0	0	0	5
	2	0	0	0	0		0	0	0	6
	3	0	0	0	0		0	0	0	23
3	1	0	0	0	**6**	2	0	0	0	**12**
	2	0	**2**	**16**	**33**		0	**1**	**5**	**31**
	3	**6**	**21**	11	4		**3**	**9**	11	28
4	1	0	0	0	0	2	0	0	0	5
	2	0	0	0	**3**		0	0	0	**9**
	3	0	**0**	**5**	**89**		0	**1**	**2**	**82**
5	1	0	**6**	**10**	**78**	2	1	**4**	**7**	**56**
	2	1	2	1	0		4	6	6	6
	3	0	0	0	0		4	2	3	2
6	1	0	2	3	11	5	1	3	3	15
	2	11	**19**	5	1		6	**11**	9	10
	3	**40**	2	0	0		**23**	9	4	6
7	1	**54**	**12**	2	1	7	**28**	**16**	6	9
	2	**18**	1	0	0		**15**	5	3	4
	3	4	0	0	0		5	2	1	1
8	1	0	0	0	0	1	0	0	0	7
	2	0	0	**3**	**17**		0	1	**2**	**18**
	3	**20**	**38**	15	6		**7**	**20**	14	32

Table 4.6: Simulation results for the azacitidine trial under with sample size of $N = 48$ and 72. Columns "% Sel" gives the percentage of identifying each dose-schedule combination as the optimal treatment regimen. Column " None" gives the percentage of simulations that were stopped early. Columns "% Pat" gives the mean proportion of patients assigned to each dose-schedule regimen. The boldfaced values correspond to acceptable combinations.

			% Sel					% Pat			
			Schedule					Schedule			
Prior Sd	Scenario	Dose	1	2	3	4	None	1	2	3	4
48	1	1	0	0	0	0	**100**	12	2	2	2
		2	0	0	0	0		2	0	0	0
		3	0	0	0	0		0	0	0	0
	2	1	0	0	0	0.2	**99.8**	0	0	0	6
		2	0	0	0	0		0	0	0	6
		3	0	0	0	0		0	0	0	24
	5	1	2	**11**	**14**	**57**	2	2	**5**	**6**	**35**
		2	3	3	4	1		3	5	6	8
		3	2	0	0	0		3	2	2	3
	7	1	**49**	**12**	2	1	10	**21**	**11**	5	9
		2	**17**	1	0	0		**11**	5	2	4
		3	6	0	0	0		5	2	1	1
	8	1	0	0	0	0	1	0	0	0	7
		2	0	3	**6**	**13**		0	1	**3**	**14**
		3	**22**	**35**	12	7		**6**	**14**	10	25
72	1	1	0	0	0	0	**100**	12	2	1	3
		2	0	0	0	0		2	1	0	1
		3	0	0	0	0		0	0	0	0
	2	1	0	0	0	0	**100**	0	0	0	6
		2	0	0	0	0		0	0	0	5
		3	0	0	0	0		0	0	0	23
	5	1	1	**6**	**13**	**73**	1	2	**5**	**9**	**67**
		2	1	3	2	0		4	6	6	9
		3	0	0	0	0		3	3	3	3
	7	1	**58**	**8**	0	1	11	**40**	**17**	7	9
		2	**19**	0	0	0		**19**	5	3	3
		3	4	0	0	0		6	2	1	1
	8	1	0	0	0	0	0	0	0	0	7
		2	0	1	**5**	**17**		0	2	**3**	**21**
		3	**32**	**35**	8	2		**13**	**27**	16	30

4.6 Application - The Irinotecan Trial

4.6.1 Simulation Design

In the motivating irinotecan trial, clinicians specified $J = 3$ doses for each of the two non-nested schedules, for a total of 6 combinations. Specifically, the doses are $35, 70$ and 90 mg/m^2 for weekly schedule and $180, 240$ and 300 mg/m^2 for tri-weekly schedule. The maximum follow-up time for each patient is $\omega_T = \omega_E = 420$ days. A maximum of $N = 72$ patients will be enrolled in cohorts of three patients. The algorithm was implemented with $p_T = 0.4, p_E = 0.1, \pi_T = 0.35$ and $\pi_E = 0.35$. (dose 1, schedule 1) is assigned to the first cohort and (dose 1, schedule 2) is assigned to the second cohort, because the ordering of toxicity or response rate of the two dose-schedule regimes is unknown. Our algorithm begins with the third cohort. Investigators would like to determine which of the six combinations is the optimal treatment regimen for both toxicity and efficacy. The prior means as well prior variances of model parameters are shown in Table 4.7 and the skeleton can be found in Table 4.8 and Table 4.9. We tried three different prior SDs (1, 2 and 4) and selected prior SD = 1 for the trial.

We examined our approach in seven different scenarios that are summarized in Table 4.8. The true toxicity and response rates are also generated from our assumed model. We simulated patients to have exponentially distributed inter-arrival times with a mean of four weeks. The rest of the simulation designs are the same as those in Section 4.5.

4.6.2 Simulation Results

Overall, our design performs well in all of the seven scenarios according to the results summarized in Table 4.8. In Scenarios 1-2, the trial was stopped early in 100% and 99.6% of the simulations due to toxicity and futility respectively. We assigned 15% and 29% of the patients before the trial was terminated on average. In Scenarios 3-8,

Table 4.7: Means and standard deviations of the prior normal distributions for each parameter: non-nested

Parameter	Mean	Sd
$\beta_{0,T}$	-4.88	1
$\beta_{1,T}$	-1.131	1
$\log(\lambda_T)$	0.72	1
α_T	1.87	0
γ_T	1.26	1
$\beta_{0,E}$	-4.59	1
$\beta_{1,E}$	-1.134	1
$\log(\lambda_E)$	0.47	1
α_E	1.88	0
γ_E	1.26	1
ζ_s	0	1

we are able to select the acceptable regimes with probabilities $100\%, 100\%, 57\%, 80\%$ and 65%. For patient assignment, we note that most of patients were treated at acceptable dose-schedule regimes.

We examined the sensitivity of our design to the maximum sample size by repeating our simulations using sample sizes of 60, the results of which are summarized in Table 4.9. The results show that decreasing the sample size from 72 to 60 may result in a nontrivial loss is selecting and assigning correct combinations. Therefore, we selected a maximum number of 60 patients for this study. We also examined the sensitivity of our design for prior variance and model misspecification. Our conclusion is similar to that in the azacitidine trial and we have omitted these results for brevity.

4.7 Discussion

In our algorithm, we define the acceptable dose-schedule regimes to be the set $\mathcal{A}(D)$ that satisfies both toxicity and efficacy constraints and we further define the optimal regimen to be the regimen in $\mathcal{A}(D)$ with the highest estimated response rate. More generally, the optimal dose-schedule regimen could be based on maximizing

Table 4.8: Simulation results for the itinotecan trial with sample size of $N = 72$. Columns "% Sel" gives the percentage of identifying each dose-schedule combination as the optimal treatment regimen. Column " None" gives the percentage of simulations that were stopped early. Columns "% Pat" gives the mean proportion of patients assigned to each dose-schedule regimen. (π_T^0,π_E^0) denotes the skeleton used in the trial. The boldfaced values correspond to acceptable combinations.

		Weekly Schedule			Tri-weekly Schedule			
Sc.		Dose 1	Dose 2	Dose 3	Dose 1	Dose 2	Dose 3	None
	(π_T^0,π_E^0)	(.07,.15)	(.15,.40)	(.33,.50)	(.05,.15)	(.19,.40)	(.32,.50)	
1	(π_T,π_E)	(.55,.27)	(.67,.36)	(.74,.43)	(.51,.28)	(.72,.46)	(.9,.67)	
	% Sel	0	0	0	0	0	0	**100**
	% Pat	10	1	0	15	1	0	
2	(π_T,π_E)	(.07,.08)	(.09,.09)	(.11,.11)	(.06,.05)	(.1,.07)	(.18,.09)	
	% Sel	0	0	0	0	0	0.4	**99.6**
	% Pat	4	4	15	4	4	8	
3	(π_T,π_E)	(.07,.27)	(.09,.36)	**(.11,.43)**	(.06,.28)	(.1,.46)	**(.18,.67)**	
	% Sel	0	0	**10**	0	0	**90**	0
	% Pat	4	4	**36**	9	5	**41**	
4	(π_T,π_E)	(.11,.5)	(.15,.56)	**(.18,.59)**	(.1,.28)	(.17,.35)	**(.29,.43)**	
	% Sel	0	0	**98**	0	0	**2**	0
	% Pat	5	5	**46**	9	8	**27**	
5	(π_T,π_E)	(.25,.51)	(.4,.63)	(.51,.71)	(.49,.52)	(.83,.74)	(.99,.91)	
	% Sel	**57**	10.4	0	1	0	0	31
	% Pat	**47**	12	1	16	1	0	
6	(π_T,π_E)	**(.37,.51)**	(.45,.63)	(.5,.71)	**(.26,.52)**	**(.38,.74)**	(.53,.91)	
	% Sel	**10**	5	0	**43**	**27**	4	11
	% Pat	**21**	7	4	**34**	**19**	8	
7	(π_T,π_E)	(.2,.29)	**(.26,.4)**	**(.31,.47)**	(.21,.32)	**(.33,.52)**	(.5,.76)	
	% Sel	3	**9**	**28**	6	**37**	15	3
	% Pat	12	**12**	**24**	17	**19**	15	

Table 4.9: Simulation results for the itinotecan trial with sample size of $N = 60$. Columns "% Sel" gives the percentage of identifying each dose-schedule combination as the optimal treatment regimen. Column " None" gives the percentage of simulations that were stopped early. Columns "% Pat" gives the mean proportion of patients assigned to each dose-schedule regimen. π_T^0,π_E^0 denotes the skeleton used in the trial. The boldfaced values correspond to acceptable combinations.

		Weekly Schedule			Tri-weekly Schedule			
Sc.		Dose 1	Dose 2	Dose 3	Dose 1	Dose 2	Dose 3	None
	(π_T^0,π_E^0)	(.07,.15)	(.15,.40)	(.33,.50)	(.05,.15)	(.19,.40)	(.32,.50)	
1	(π_T,π_E)	(.55,.27)	(.67,.36)	(.74,.43)	(.51,.28)	(.72,.46)	(.9,.67)	
	% Sel	0	0	0	0	0	0	**100**
	% Pat	12	0	0	17	2	0	
2	(π_T,π_E)	(.07,.08)	(.09,.09)	(.11,.11)	(.06,.05)	(.1,.07)	(.18,.09)	
	% Sel	0	0	0.4	0	0	0.4	**99.2**
	% Pat	5	5	18	10	5	10	
3	(π_T,π_E)	(.07,.27)	(.09,.36)	**(.11,.43)**	(.06,.28)	(.1,.46)	**(.18,.67)**	
	% Sel	0	0	**17**	0	1	**83**	0
	% Pat	5	5	**35**	10	6	**38**	
4	(π_T,π_E)	(.11,.5)	(.15,.56)	**(.18,.59)**	(.1,.28)	(.17,.35)	**(.29,.43)**	
	% Sel	0	1	**95**	0	0	**4**	0
	% Pat	6	6	**41**	11	9	**26**	
5	(π_T,π_E)	**(.25,.51)**	(.4,.63)	(.51,.71)	(.49,.52)	(.83,.74)	(.99,.91)	
	% Sel	**52**	13	1	2	0	0	33
	% Pat	**44**	12	2	20	2	0	
6	(π_T,π_E)	**(.37,.51)**	(.45,.63)	(.5,.71)	**(.26,.52)**	**(.38,.74)**	(.53,.91)	
	% Sel	**7**	2	1	**42**	**31**	4	13
	% Pat	**22**	8	4	**32**	**18**	9	
7	(π_T,π_E)	(.2,.29)	**(.26,.4)**	**(.31,.47)**	(.21,.32)	**(.33,.52)**	(.5,.76)	
	% Sel	5	**11**	**31**	5	**30**	15	3
	% Pat	14	**13**	**22**	18	**18**	14	

the posterior estimate of a utility function that explicitly incorporates the trade-off between toxicity and efficacy. Obviously, the utility function should decrease with the toxicity rate and increase with the response rate. For instance, a possible choice is the odds ratio between the response and toxicity rate. However, the odds ratio might not work for all drugs, since the critical aspect of using a utility function is that it should reflect clinical knowledge and background. Hence, it is important to interact with clinicians regarding the details on toxicity and efficacy trade-off.

Another possible extension of our current method is to implement dose and/or schedule re-assignment to our design. While dose-schedule regimen is assigned to the next cohort adaptively during the trial, a patient, once being assigned, is assumed to receive the same initially assigned regimen (except when the dose is reduced due to the dose reduction rule) even though it might turn out to be a sub-optimal regimen later in the trial. However, the benefit of introducing re-assignment in a Phase I/II trial is not directly clear. Furthermore, the potential benefit might really depend on what utility function and re-assignment rule are used in the trial.

We used a relatively parsimonious model in our design because the sample size is usually small in early phase trials and the overall model fit is not the primary goal for a dose and/or schedule finding study. However, if the investigator is concerned about the model assumptions, we could certainly adopt a more complex model. For example, one could model the correlation parameter λ_c as a function of dose to make the model more flexible. A strong assumption made in our method is that the toxicity and response rates increase with dose, which is reasonable for azacitidine, irinotecan and many other cytotoxic agents. However, for some newer cytostatic agents, the response rates might not necessarily increase with dose. In such situations, one might use a second order polynomial to model the relationship between $\log(\theta_c)$ and dose. Another assumption in our model is that we can observe the exact event time for both toxicity and efficacy outcomes. However, in practice, one might only observe

them in a certain interval, for instance, it might be sensible to evaluate the efficacy outcomes every month instead of continually. Our model could be easily extended to accommodate such interval censored data, in which the intervals are pre-specified by the clinician.

Another interesting area of research is to develop methods to accommodate patient dropout. Patients might have such severe side effects that they might drop out from the study, or worse, they might die. In such situations, if the response has not yet occurred, the response would most likely be informatively censored. However, if a response has occurred without toxicity, patients would still receive treatment and be under follow-up. Hence, toxicity could possibly censor efficacy outcome, but not vise versa. Possible future work is to adopt a semi-competing risks model to solve the problem of patient death or drop-out.

CHAPTER V

Summary

In this dissertation we relax the assumption of a fixed prior variance in the traditional CRM for a Phase I dose-finding design and propose three systematic approaches to adaptively calibrate the prior variance continually throughout the trial. Our approaches have the ability to perform better than the traditional CRM using a constant prior variance as well as methods that calibrate the prior variance only at the beginning of the trial.

We have contributed to the existing literature on the design of Phase I dose- and schedule-finding studies by first expanding the schedule-finding model used by Liu and Braun (2009) to incorporate variations in dose and have then incorporated an algorithm for assessing whether or not currently enrolled patients should have their assigned dose and/or schedule changed to one with an estimated DLT rate closer to that desired. We modeled the per-administration dose as a covariate through the cure parameter of the hazard function for a single administration in the framework of the cure rate model, which is particularly attractive as a significant number of patients would not have DLTs from a new agent. While the functional form of the hazard function of a single administration can be chosen to reflect the background of a specific study, another major advantage of our method lies in the flexibility to model any sequence of dose and administration combinations. Furthermore, we have proposed

an algorithm to optimize the intra-patient dose-schedule reassignment in addition to the conventional method for inter-patient optimization. Simulations indicate that our design identifies correct dose and schedule combinations as well as the traditional method that does not allow for intra-patient doses-schedule reassignments, but with a larger number of patients assigned to those combinations.

Finally, we generalize the Phase I dose-schedule-finding design to allow for modeling both toxicity and efficacy in Phase I/II dose-schedule-finding clinical trials. We propose marginal time-to-toxicity and time-to-response models based on the cure rate model for multiple administrations with the per-administration dose as a covariate through the cure parameter of the hazard function for each single administration. The bivariate model is then constructed using a copula function. Furthermore, we describe both toxicity and efficacy acceptability criteria for selecting the optimal dose-schedule regimen as well as for stopping the trial early in case of a overly toxic or inefficacious trial.

We have discussed the future work for each project in the corresponding discussion section for each chapter. However, one important next step common to all three projects is to write, validate, document and publish robust, efficient and open-source routines to implement each method, using uniform syntax and a parsimonious set of control parameters. This will generate a simulation platform that can be used by any member of the statistical community to compare the operating characteristics of various trial designs or to modify our source code to tailor to their specific studies. These routines will be maintained in a single `R` package and posted on `CRAN`. Furthermore, we are also planning to develop user-friendly software or a web-based application to allow better bench-to-bedside translational statistical tools.

BIBLIOGRAPHY

BIBLIOGRAPHY

Babb, J., Rogatko, A., and Zacks, S. (1998), "Cancer Phase I Clinical Trials: Efficient Dose Escalation with Overdose Control," *Statistics in Medicine*, 17, 1103-1120.

Borner, M. M., Bernhard, J., Dietrich, D., Popescu, R., Wernli, M., Saletti, P., Rauch, D., Herrmann, R., Koeberle, D., Honegger, H., brauchli, P., Lanz, D. and Roth, A. D. (2005), "A Randomized Phase II Trial of Capecitabine and Two Different Schedules of Irinotecan in First-line Treatment of Metastatic Colorectal Cancer: Efficay, Quality-of-life and Toxicity," *Annals of Oncology*, 16, 282-288.

Braun, T. M. (2002), "The Bivariate Continual Reassessment Method: Extending the CRM to Phase I Trials of Two Competing Outcomes," *Controlled Clinical Trials*, 23, 240-256.

Braun, T.M. (2005), "Generalizing the TITE-CRM to Adapt for Early- and Late-onset Toxicities," *Statistics in Medicine*, 20, 153-169.

Braun, T. M., Thall, P. F., Nguyen, H., and de Lima, M. (2007), "Simultaneously Optimizing Dose and Schedule of a New Cytotoxic Agent," *Clinical Trials*, 4, 113-124.

Braun, T. M., Yuan, Z., Thall, P. F., Nguyen, H., and Lima, M. (2005), "Determining a Maximum-Tolerated Schedule of a Cytotoxic Agent," *Biometrics*, 61, 335-343.

Chen, M., Ibrahim, J., and Sinha, D. (1999), "A New Bayesian Model for Survival Data with a Surviving Fraction," *Journal of American Statistical Association*, 94, 909-919.

Chen, M., Ibrahim, J., and Sinha, D. (2002), "Bayesian Inference for Multivariate Survival Data with a Cure Fraction," *Journal of Multivariate Analysis*, 80, 101-126.

Cheung, Y. K. (2005), "Coherence Principles in Dose-finding studies," *Biometrika*, 92, 863-873.

Cheung, Y. K., and Chappell R. (2000), "Sequential Designs for Phase I Clinical Trials with Late-onset Toxicity," *Biometrics*, 56, 1177-1182.

Cheung, Y. K., and Chappell R. (2002), "A Simple Technique to Evaluate Model Sensitivity in the Continual Reassessment Method," *Biometrics*, 58, 671-674.

Chevret, S. (1993), "The Continual Reassessment Method in Cancer Phase I Clinical Trials: A Simulation Study," *Statistics in Medicine*, 30, 1093-1108.

Cutsem, V. E., Tejpar, S., Vanbeckevoort, D., Peeters, M., Humblet, Y., Gelderblom, H., Vermorken, J. B., Viret, F., Glimelius, B., Gallerani, E., Hendlisz, A., Cats, A., Moehler, M., Sagaert, X., Vlassak, S., Schlichting, M., and Ciardiello, F. (2012), "Intrapatient Cetuximab Dose Escalation in Metastatic Colorectal Cancer According to the Grade of Early Skin Reactions: The randomized EVEREST Study," *Journal of Clinical Oncology*, 30, 28612868.

Faries, D. (1994), "Practical Modifications of the Continual Reassessment Method for Phase I Cancer Clinical Trials." *Journal of Biopharmaceutical Statistics*, 4, 147-164.

Hougaard, P. (2000), *Analysis of Multivariate Survival Data*, New York: Springer.

Iasonos, A., and O'Quigley, J. (2011), "Continual Reassessment and Related Designs in Dose-finding Studies," *Statistics in Medicine*, 30, 2057-2061.

Jeffreys, H. (1961), *Theory of Probability*, London: Oxford University Press.

Korn, E. L., Midthune, D., Chen, T. T., Rubinstein, L. V., Christain, M. C., and Simon, R. A. (1994), "Comparison of Two Phase I Trial Designs," *Statistics in Medicine*, 13, 1799-1806.

Lee, S. M., Cheung, Y. K. (2009), "Model Calibration in the Continual Reassessment Method," *Clinical Trials*, 6, 227-238.

Lee, S. M., Cheung Y. K. (2011), "Calibration of Prior Variance in the Bayesian Continual Reassessment Method," *Statistics in Medicine*, 30, 2081-2089.

Li, Y., Bekele, B. N., Ji, Y., and Cook, J. D. (2008), "Dose-Schedule Finding in Phase I/II Clinical Trials using a Bayesian Isotonic Transformation," *Statistics in Medicine*, 27, 4895-4913.

Liu, C., and Braun, T. M. (2009), "Parametric Non-Mixture Cure Models for Schedule Finding of Therapeutic Agents," *The Journal of the Royal Statistical Society, Series C (Applied Statistics)* , 58, 225-236.

de Lima, M., Giralt, S., Thall, P. F., Silva, L., Jones,R. B., Komanduri K., Braun, T. M., Nguyen, H. Q., Champlin, R., and Garcia-Manero, G. (2010), "Maintenance Therapy With Low-Dose Azacitidine After Allogeneic Hematopoietic Stem Cell Transplantation for Recurrent Acute Myelogenous Leukemia or Myelodysplastic Syndrome: A Dose and Schedule Finding Study," *Cancer* , 116, 5420-5431.

Moller, S. (1995), "An Extension of the Continual Reassessment Methods using a Preliminary Up-and-down Design in a Dose-finding Study in Cancer Patients, in order to Investigate a Greater Range of Doses," *Statistics in Medicine*, 14, 911-922.

Murphy, S. A. (2002), "Optimal Dynamic Treatment Regimes," *Journal of the Royal Statistical Society, Series B*, 65, 331-336.

Neuenschwander, B., Branson, M., Gsponer, T. (2008) Critical Aspects of the Bayesian Approach to Phase I Cancer Trials. *Statistics in Medicine*, 27, 2420-2439.

O'Quigley, J., Pepe, M., and Fisher, L. (1990), "Continual Reassessment Method: A Practical Design for Phase I Clinical Trials in Cancer," *Biometrics* 46, 33-48.

O'Quigley, J., and Reiner, E. (1998), "A Stopping Rule for the Continual Reassessment Method," *Biometrika*, 85, 741-748.

O'Quigley, J., and Shen, L. Z. (1996), "Continual Reassessment Method: A Likelihood Approach," *Biometrics*, 52, 673-684.

Paoletti, X., and Kramar, A. (2009), "A comparison of model choices for the continual reassessment method in phase I cancer trials," *Statistics in Medicine*, 28, 3012-3028.

R Development Core Team (2008), *R: A Language and Environment for Statistical Computing*, Vienna, Austria: R Foundation for Statistical Computing.

Raftery, A. E., Madigan, D., and Hoeting, J. A. (1997), "Bayesian Model Averaging for Linear Regression Models," *Journal of the American Statistical Association*, 92, 179-191.

Rosenberger, W., and Haines, L. (2009), "Competing Designs for Phase I Clinical Trials: A Review," *Statistics in Medicine*, 28, 3012-3028.

Rosenthal, J. (2007), "AMCMC: An R Interface for Adaptive MCMC," *Computational Statistics and Data Analysis*, 51, 5467-5470.

Shu, J., and O'Quigley, J. (2008), "Dose-escalation designs in oncology: ADEPT and the CRM," *Statistics in Medicine*, 27, 5345-5353.

Storer, B. (1989), "Consistency of Continual Reassessment Method under Model Misspecification," *Biometrics*, 45, 925-937.

Thall, P. F. and Cook J. D. (2004), "Dose-Finding Based on Efficacy-Toxicity Trade-Offs," *Biometrics*, 60, 684-693.

Thall, P. F. and Russell K. E. (1998), "A Strategy for Dose-Finding and Safety Monitoring Based on Efficacy and Adverse Outcomes in Phase I/II Clinical Trials," *Biometrics*, 54, 251-264.

Tsiatis, A. (1975), "A Nonidentiability Aspect of the Problem of Competing Risks," *Proceedings of the National Academy of Sciences in USA*, 72, 20-22.

Tsodikov, A., Ibrahim J., and Yakovlev, A. (2003), "Estimating Cure Rates from Survival Data: an Alternative to Two-component Mixture Models," *Journal of American Statistical Association*, 98, 1063-1078.

Yin, G., and Yuan, Y. (2009), "Bayesian Model Averaging Continual Reassessment Method in Phase I Clinical Trials," *Journal of the American Statistical Association*, 104, 954-972.

Yuan, Y., and Yin, G. (2009), "Bayesian Dose Finding by Jointly Modelling Toxicity and Efficacy as Time-to-event Outcomes," *The Journal of the Royal Statistical Society, Series C (Applied Statistics)*, 58, 719-736.

Yuan, Y., and Yin, G. (2011), "Bayesian Hybrid Dose-finding Design in Phase I Oncology Clinical Trials," *Statistics in Medicine*, 30, 2098-2108.

Zhang, J., Braun, T. M., and Taylor, J. M.G. (2012), "Adaptive Prior Variance Calibration in the Bayesian Continual Reassessment Method," *Statistics in Medicine*, doi: 10.1002/sim.5621