

8-2011

BAYESIAN PHASE I DOSE FINDING IN CANCER TRIALS

Lin Yang

Follow this and additional works at: https://digitalcommons.library.tmc.edu/utgsbs_dissertations



Part of the [Biostatistics Commons](#), [Clinical Trials Commons](#), [Statistical Models Commons](#), and the [Survival Analysis Commons](#)

Recommended Citation

Yang, Lin, "BAYESIAN PHASE I DOSE FINDING IN CANCER TRIALS" (2011). *The University of Texas MD Anderson Cancer Center UTHealth Graduate School of Biomedical Sciences Dissertations and Theses (Open Access)*. 136.

https://digitalcommons.library.tmc.edu/utgsbs_dissertations/136

This Dissertation (PhD) is brought to you for free and open access by the The University of Texas MD Anderson Cancer Center UTHealth Graduate School of Biomedical Sciences at DigitalCommons@TMC. It has been accepted for inclusion in The University of Texas MD Anderson Cancer Center UTHealth Graduate School of Biomedical Sciences Dissertations and Theses (Open Access) by an authorized administrator of DigitalCommons@TMC. For more information, please contact digitalcommons@library.tmc.edu.

BAYESIAN PHASE I DOSE FINDING IN CANCER TRIALS

A

DISSERTATION

Presented to the Faculty of

The University of Texas

Health Science Center at Houston

and

The University of Texas

M. D. Anderson Cancer Center

Graduate School of Biomedical Sciences

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

by

Lin Yang, MS

Houston, Texas

August, 2011

BAYESIAN PHASE I DOSE FINDING IN CANCER TRIALS

by

Lin Yang, MS

APPROVED:

Supervisory Professor: Donald A. Berry, Ph.D.

B. Nebiyu Bekele, Ph.D.

Yisheng Li, Ph.D.

Ying Yuan, Ph.D.

Funda Meric-Bernstam, MD

APPROVED:

Dean, The University of Texas

Graduate School of Biomedical Sciences at Houston

To my parents, Xingfan and Zeqing, for their unconditional love,
inspiration and sacrifice;

To my dear husband, Meng, without whom I could never have
reached this far;

To all people who have made the impossible possible in fighting
cancer.

Acknowledgements

I offer my sincerest gratitude to my supervisors, Dr. Donald Berry and Dr. Nebiyu Beleke for their knowledge and patient support while allowing me the room to work in my own way. I attribute my Ph.D. degree to their encouragement and dedication.

Thanks to my supervisory committee for their kind contributions.

One simply could not wish for a better coordinator than Lydia Davis, who has provided administrative support during my pursuit of this degree. I am indebted to LeeAnn Chastain for assisting me in proofreading the dissertation and for further tutoring me in scientific writing.

In my daily work I have been blessed with a friendly and cheerful group of faculty and fellow students, from whom I have learned much in many ways. Dr. Peter Mueller referred me for the internship and offered much advice throughout the study. Dr. Gary Rosner served as the chair for my PhD candidacy examination and provided invaluable insights on the proposal. Dr. Yisheng Li and Dr. Yuan Ji shared with me their expertise in various aspects of statistics. Violeta Hennessey has been a good companion on an otherwise lonely and uncertain road.

Abstract

This dissertation explores phase I dose-finding designs in cancer trials from three perspectives: the alternative Bayesian dose-escalation rules, a design based on a time-to-dose-limiting toxicity (DLT) model, and a design based on a discrete-time multi-state (DTMS) model.

We list alternative Bayesian dose-escalation rules and perform a simulation study for the intra-rule and inter-rule comparisons based on two statistical models to identify the most appropriate rule under certain scenarios. We provide evidence that all the Bayesian rules outperform the traditional “3+3” design in the allocation of patients and selection of the maximum tolerated dose.

The design based on a time-to-DLT model uses patients’ DLT information over multiple treatment cycles in estimating the probability of DLT at the end of treatment cycle 1. Dose-escalation decisions are made whenever a cycle-1 DLT occurs, or two months after the previous check point. Compared to the design based on a logistic regression model, the new design shows more safety benefits for trials in which more late-onset toxicities are expected. As a trade-off, the new design requires more patients on average.

The design based on a discrete-time multi-state (DTMS) model has three important attributes: (1) Toxicities are categorized over a

distribution of severity levels, (2) Early toxicity may inform dose escalation, and (3) No suspension is required between accrual cohorts. The proposed model accounts for the difference in the importance of the toxicity severity levels and for transitions between toxicity levels. We compare the operating characteristics of the proposed design with those from a similar design based on a fully-evaluated model that directly models the maximum observed toxicity level within the patients' entire assessment window. We describe settings in which, under comparable power, the proposed design shortens the trial. The proposed design offers more benefit compared to the alternative design as patient accrual becomes slower.

Contents

List of Figures	x
List of Tables	xiii
1 Introduction	1
1.1 Motif	1
1.2 Background and significance	4
1.2.1 Phase I cancer trials	4
1.2.2 Bayesian adaptive design	5
1.2.3 Traditional designs for phase I dose-finding	7
1.2.3.1 “3+3” algorithm	7
1.2.3.2 Continual reassessment method (CRM)	9
1.2.3.3 Escalation with overdose control (EWOC)	10
2 Bayesian Dose-Assignment Rules for Phase I Dose Finding in Cancer Trials	12
2.1 Two probability models	14
2.1.1 Bayesian beta/binomial model	14
2.1.2 Logistic regression model	15

2.2	Five dose-assignment rules	16
2.2.1	<i>CRM_r</i> by O’Quigley <i>et al.</i>	16
2.2.2	<i>MaxTI</i> by Neuenschwander <i>et al.</i>	16
2.2.3	<i>ODC</i> by Bekele <i>et al.</i>	17
2.2.4	<i>Loss</i> by Conaway <i>et al.</i>	19
2.2.5	Modified “3+3” algorithm	19
2.3	Universal overdose control criteria	20
2.4	Simulation study	21
2.4.1	Scenarios	21
2.4.2	Prior specification	23
2.4.3	Operating characteristics based on modified “3+3” algorithm	24
2.4.4	Operating characteristics for <i>CRM_r</i>	25
2.4.5	Operating characteristics for <i>MaxTI</i>	25
2.4.6	Operating characteristics for <i>ODC</i>	29
2.4.7	Operating characteristics for <i>Loss</i>	32
2.4.8	Comparison between dose-assignment rules	35
2.4.8.1	Comparison based on the beta/binomial model	35
2.4.8.2	Comparison based on the logistic regression model	38
2.5	Discussion	45
3	A Phase I Dose-Finding Design Based on a Time-to-DLT Model	53
3.1	Introduction	53
3.2	Time-to-DLT model	55
3.3	Study Design	59
3.3.1	Decision rule	59
3.3.2	Trial conduct	60
3.4	Simulation	61

CONTENTS

3.4.1	Scenarios	61
3.4.2	Operating characteristics	64
3.5	Discussion	75
4	Adaptive Dose Finding in Phase I Cancer Trials Based on a Discrete-Time Multi-State Model	77
4.1	Introduction	78
4.2	Bayesian discrete-time multi-state model	80
4.3	Dose-assignment rule	85
4.4	An alternative fully-evaluated model as a comparison	87
4.5	Simulation comparing the proposed design with the alternative . .	88
4.5.1	Scenarios	88
4.5.2	Sample trial conduct based on the fully-evaluated design and the <i>DTMS</i> design	90
4.5.3	Simulation results	91
4.6	Sensitivity test	107
4.7	Discussion	108
5	Summary	109
6	Bibliography	112

List of Figures

2.1	Dose-toxicity scenarios.	22
2.2	95% CI of prior probability of toxicity	24
2.3	Operating characteristics based on the modified “3+3” algorithm	26
2.4	Operating characteristics for <i>CRM_r</i> based on beta/binomial model	27
2.5	Operating characteristics for <i>CRM_r</i> based on the logistic regression model	28
2.6	Operating characteristics for <i>MaxTI</i> based on beta/binomial model	30
2.7	Operating characteristics for <i>MaxTI</i> based on the logistic regres- sion model	31
2.8	Operating characteristics for <i>ODC</i> based on the beta/binomial model	33
2.9	Operating characteristics for <i>ODC</i> based on the logistic regression model	34
2.10	Operating characteristics for <i>Loss</i> based on the beta/binomial model	36
2.11	Operating characteristics for <i>Loss</i> based on the logistic regression model	37
2.12	OC on 5 rules based on the beta/binomial model: Scenario-MTD0	39
2.13	OC on 5 rules based on the beta/binomial model: Scenario-MTD1	40
2.14	OC on 5 rules based on the beta/binomial model: Scenario-MTD2	41
2.15	OC on 5 rules based on the beta/binomial model: Scenario-MTD3	42

LIST OF FIGURES

2.16	OC on 5 rules based on the beta/binomial model: Scenario-MTD4	43
2.17	OC on 5 rules based on the beta/binomial model: Scenario-MTD5	44
2.18	OC on 5 rules based on the logistic regression model: Scenario-MTD0	46
2.19	OC on 5 rules based on the logistic regression model: Scenario-MTD1	47
2.20	OC on 5 rules based on the logistic regression model: Scenario-MTD2	48
2.21	OC on 5 rules based on the logistic regression model: Scenario-MTD3	49
2.22	OC on 5 rules based on the logistic regression model: Scenario-MTD4	50
2.23	OC on 5 rules based on the logistic regression model: Scenario-MTD5	51
3.1	Prior distribution of the end-of-cycle-1 DLT rate	58
3.2	Trial conduct	62
3.3	Simulated trial conduct	63
3.4	Dose-toxicity plots	65
3.5	OC for 3 doses with 6-cycle follow-up	69
3.6	OC for 3 doses with 3-cycle follow-up	70
3.7	OC for 3 doses with 1-cycle follow-up	71
3.8	OC for 5 doses with 6-cycle follow-up	72
3.9	OC for 5 doses with 3-cycle follow-up	73
3.10	OC for 5 doses with 1-cycle follow-up	74
4.1	True value of ATS determined by τ	89
4.2	Sample trial conduct based on the alternative designs	92
4.3	Operating characteristic with $R \sim \text{exp}(1)$ for scenarios with constant toxicity hazard	95
4.4	Operating characteristic with $R \sim \text{exp}(1)$ for scenarios with decreasing toxicity hazard	96
4.5	Operating characteristic with $R \sim \text{exp}(1)$ for scenarios with increasing toxicity hazard	97

LIST OF FIGURES

4.6	Operating characteristic with $R \sim \text{exp}(1)$ for scenarios with scalloping toxicity hazard	98
4.7	Distribution of the posterior mean ATS for scenario group 1 with $R \sim \text{exp}(1)$	99
4.8	Distribution of the posterior mean ATS for scenario group 2 with $R \sim \text{exp}(1)$	100
4.9	Operating characteristic with $R \sim \text{exp}(0.5)$ for scenarios with constant toxicity hazard	101
4.10	Operating characteristic with $R \sim \text{exp}(0.5)$ for scenarios with decreasing toxicity hazard	102
4.11	Operating characteristic with $R \sim \text{exp}(0.5)$ for scenarios with increasing toxicity hazard	103
4.12	Operating characteristic with $R \sim \text{exp}(0.5)$ for scenarios with scalloping toxicity hazard	104
4.13	Distribution of the posterior mean ATS for scenario group 1 with $R \sim \text{exp}(0.5)$	105
4.14	Distribution of the posterior mean ATS for scenario group 2 with $R \sim \text{exp}(0.5)$	106

List of Tables

1.1	Brief comparison of Bayesian and frequentist methods in randomized trials (Spiegelhalter et al., 1999. Table used with permission from BMJ Publishing Group Ltd.)	8
2.1	Prior specification for $\theta_d \sim \text{beta}(\alpha_d, \beta_d)$ for the beta/binomial model	23
2.2	Bivariate normal prior specification on $(\alpha, \log \beta)$ for the logistic regression model	23
3.1	Prior quantiles for the end-of-cycle-1 DLT rate	57
3.2	Parameters for bivariate normal distribution	57
3.3	Scenarios	66
3.4	Key design parameters of the time-to-DLT model and the logistic regression model	67
4.1	Phase I dose-finding methods and design factors	80
4.2	One patient's contribution to the probability model through $p_{j,k,k'}$. . .	82
4.3	Structure of the transition probability matrix	89
4.4	Operating characteristics of the sensitivity test	107

1

Introduction

1.1 Motif

Phase I oncology trials determine the highest dose that does not result in a prohibitive toxicity, which is known as the maximum tolerated dose (MTD). This determination is accomplished by designating the doses given to consecutive patient cohorts based on the toxicity responses from patients in the trial who have already received the treatment. As the initial opportunity to observe patients' responses to a new drug or treatment combination, a phase I trial is a fundamental component of the most important steps in drug development.

Current methods to determine the MTD in phase I cancer trials are inefficient and may unnecessarily prolong the duration of the trial. We address multiple issues in phase I clinical trials. These problems include that toxicity is typically treated as a binary outcome and accrual is regularly suspended. We develop methods that stem from an acknowledgment that toxicity is actually a set of heterogeneous adverse outcomes that can range from mild to severe. In addition, we develop approaches that do not require suspension of accrual. Thus our approach

wastes less information and can be completed more rapidly.

In this dissertation we explore phase I oncology trials from three perspectives: In Chapter 2 we investigate a set of Bayesian dose-assignment rules applied to two popular probability models, and recommend appropriate rules under certain scenarios. In Chapter 3 we describe a phase I dose-finding design based on a time-to-event model with flexible patient cohort depending on the occurrence of dose-limiting toxicity (DLT) within the first cycle of treatment, the development of which I participated in during an internship with Novartis Oncology. In Chapter 4 we propose an adaptive phase I dose-finding method based on a discrete-time multi-state model that incorporates the toxicity level and allows patients' partial toxicity information to inform decision about dose-escalation.

Traditional up-and-down algorithms are the most commonly used methods for phase I dose-finding trial designs. Among them is the well known "3+3" design, for which decisions are made according to the performance of the current cohort of three patients, in addition to the performance of the previous cohort at the same dose. Model-based dose-assignment rules developed over the past decades have also gained attention in cancer trials. A common dose-escalation rule used in conjunction with the continual reassessment method (CRM), is based on point estimates for the probability of toxicity at each dose. Using that rule, for example, the next cohort is treated at the dose with a mean posterior probability of toxicity closest to the prespecified target probability, e.g., 30%. Other model-based methods include partitioning the posterior probability of toxicity into sub-intervals: Dose-escalation is then based on the toxicity region of the current dose (Bekele *et al.*, 2010) or on augmenting the likelihood of achieving a target toxicity interval without causing prohibitive levels of toxicity (Neuenschwander *et al.*, 2008). Conaway *et al.* (2004) proposed an asymmetric loss function that penalizes the toxicity probability according to its relative relationship to the

target. Each of the above dose-assignment rules has merit in certain scenarios. To compare the performance among the various rules, we investigate each rule's operating characteristics by changing a set of key decision parameters (described in detail in Chapter 2) and evaluate which method performs the best under each decision parameter configuration.

Traditional phase I dose-finding designs usually determine the maximum tolerated dose (MTD) based on the frequency of DLTs observed within the first cycle of treatment (e.g., 21 days). These designs ignore toxicity beyond the first cycle. To incorporate late-onset toxicity into the dose-escalation decisions, we propose a design based on a Bayesian time-to-DLT model. While dose-escalation decisions are based on the estimates of the DLT rate at the end-of-cycle 1, the model allows toxicity beyond cycle 1 to inform the end-of-cycle-1 DLT rate. The next dose-escalation point is decided by the occurrence of the cycle 1 DLT. The operating characteristics of the proposed design are compared with those from a design based on a Bayesian logistic regression model that uses the frequency of the cycle 1 DLT for the decisions of dose-escalation (see Chapter 3).

Standard phase I trials dichotomize toxicity based on whether the toxicity is dose limiting. For example, if grade 4 fatigue is a dose-limiting toxicity, while fatigue of grade 1, 2, and 3 are considered not dose limiting, then grade 1 fatigue and grade 3 fatigue are treated identically. These approaches are designed to address early toxicity of a sufficiently high level. However, due to the ethical requirement to limit the exposure of patients in the trial to toxic doses, the above methods discard useful information and lead to other problems. First, if a drug results in delayed toxicity and the trial's accrual rate is high, many patients may be treated at overly toxic doses before clinicians are aware of any toxicities. Second, different categories of toxicity are not equal importance and do not occur independently. If several categories of toxicity are likely to occur, a

model describing the progression within category is needed. To better characterize toxicity severity levels and their transitions, we propose a phase I dose-finding design based on a discrete-time multi-state model (see Chapter 3). The new design considers patients' partial toxicity information when guiding dose-escalation decisions and does not require accrual suspension between patient cohorts, which shortens the trial duration. An alternative design by Bekele *et al.* (2009) based on a fully-evaluated model is also introduced for comparison. The alternative design directly models the maximum observed toxicity level within the patient assessment window and therefore requires accrual suspension between patient cohorts (see Chapter 4).

1.2 Background and significance

1.2.1 Phase I cancer trials

This section explains concepts that are fundamental to phase I cancer trials, including trial conduct, trial features, ethical concerns, and the definition of toxicity, etc. Readers who are familiar with these concepts may want to skip this section.

The components of a typical phase I design are the prespecified doses to be evaluated, a starting dose, cohort sizes, methods for designating the appropriate doses for consecutive cohorts, a maximum sample size, stopping rules, and a criterion for choosing the maximum tolerated dose (MTD) at the trial's conclusion. A common definition of *the maximum tolerated dose (MTD)* is a dose associated with a clinically permissible degree of toxicity, which if surpassed could result in prohibitive degree of toxicity (Chevret, 2006). Trial conduct in phase I generally involves three basic processes: designating the starting dose for the first group

of patients; designating doses for successive patient cohorts based on the toxicity data accumulated from patients in the trial who already received the treatment; and monitoring the trial for early stopping if there is evidence of excessive drug toxicity or enough evidence to identify the MTD.

According to Storer (1989), the challenges of designing and analyzing a phase I trial include the use of a small sample of heterogeneous patients who are evaluated over a relatively long period. This approach is taken partially because of the ethical requirement to use a conservative approach when escalating doses, combined with a relatively subjective assessment of what does and does not constitute toxicity, and the potential need to assess response in patients who drop out of the study early for reasons other than toxicity.

The definition of toxicity, which varies substantially for different trials, may be generally considered as detrimental secondary effects of treatment that render the treatment unfeasible (Thall & Lee, 2003). However, in oncology trials, some level of toxicity is expected. Unlike patients in phase I trials for nonlethal conditions or diseases, patients with cancer often choose to enroll in a phase I trial because their disease has progressed to a terminal stage and they have few or no other treatment options. Such patients accept the risk of severe toxicity along with the possibility that treatment may prolong their lives. Oncology clinicians also accept the risk of severe toxicity, even presuming that better clinical results can be achieved when patients receive higher doses of anti-tumor drugs. Thus, a common goal is to treat patients at the highest tolerable dose.

1.2.2 Bayesian adaptive design

This section provides an essential introduction to Bayesian adaptive designs that will help readers better understand subsequent sections. Readers who are familiar

with the Bayesian paradigm and the basic concept of the adaptive design may want to skip this section.

An adaptive design prospectively allows the investigators to modify the design parameters, while patients are still enrolling into the trial by periodically examining the accumulating data. Modifications include early stopping, extending accrual, adding or dropping treatment arms, assigning patients to arms that are performing better, etc.

In addition to the goal of learning as efficiently and as rapidly as possible, the aim of adaptive designs is also to treat patients as effectively as possible. A phase I cancer trial should be designed to limit the treatment of patients at both ends of the spectrum: at low, non-therapeutic doses and at severely toxic doses. A good adaptive design navigates between the two ends of the spectrum (Berry, 2004). Furthermore, a Bayesian adaptive design provides a straightforward and flexible way to determine dose selection.

The Bayesian approach uses probability to characterize all uncertainty regarding parameters, which distinguishes it from the frequentist approach. The term “Bayesian” refers to mathematical work first proposed by Thomas Bayes (circa 1702 - 1761), who is best known for a theorem that takes his name. The essence of Bayes’ theorem is that, given some data and some hypothesis, the posterior probability that the hypothesis is true is proportional to the likelihood multiplied by the prior probability. The likelihood characterizes the data, whereas the prior specifies the belief in the hypothesis before the data were observed.

A Bayesian analysis explicitly states that:

- The prior distribution carries historical knowledge of the credibility of different values of a treatment effect,

- The likelihood is the credibility of different values of the treatment effect based on data from the trial, and
- The posterior distribution is a combination of the prior distribution and the likelihood, constructing an assessment of the treatment effect.

The above process focuses on how information gleaned in a trial should influence researchers' assessment of the treatment effect. Bayes' rule enables the distribution of the parameter under investigation to be updated as new observations accumulate. The Bayesian approach to clinical trials involves ongoing revisions of the assessment of the treatment effect during the trial as more information is gathered. This allows for greater flexibility in the trial and is a discriminating characteristic between Bayesian and frequentist approaches to clinical trials (Berry, 2006).

Fundamental differences in Bayesian and frequentist methods in randomized trials are listed in Table 1.1 (Spiegelhalter *et al.*, 1999).

1.2.3 Traditional designs for phase I dose-finding

This section introduces the traditional phase I dose-finding design and provides the basis for the development of a new design and model in later chapters.

Traditional methods in phase I dose finding include the “3+3” algorithm, the continual reassessment method (CRM) by O’Quigley, *et al.* (1990) and escalation with overdose control (EWOC) by Babb, *et al.* (1998).

1.2.3.1 “3+3” algorithm

The standard methodology in the classical “3+3” design is as follows:

- Treat 3 subjects at the starting dose level D_i

Table 1.1: Brief comparison of Bayesian and frequentist methods in randomized trials (Spiegelhalter et al., 1999. Table used with permission from BMJ Publishing Group Ltd.)

Issue	Frequentist methods		Bayesian methods
Prior information other than that in the study being analyzed	Informally used in design		Used formally by specifying a prior probability distribution
Interpretation of the parameter of interest	A fixed state of nature		An unknown quantity which can have a probability distribution
Basic question	how likely is the data, given a particular value of the parameter?		How likely is a particular value of the parameter given the data
Presentation of results	Likelihood functions, P values, confidence intervals		Plots of posterior distributions of the parameter, calculation of specific posterior probabilities of interest, and use of the posterior distribution in formal decision analysis
Interim analyses	P values and estimates adjusted for the number of analyses		Inference not affected by the number or timing of interim analyses
Interim predictions	Conditional power analyses		Predictive probability of getting a firm conclusion
Dealing with subsets in trials	Adjusted P values (for example, Bonferroni)		Subset effects shrunk towards zero by a skeptical prior

- If 0/3 subjects experience dose-limiting toxicity (DLT), escalate to dose D_{i+1}
- If 1/3 or more subjects experience DLT, treat 3 more subjects at dose level D_i
- If 1/6 subjects experience DLT, escalate to dose D_{i+1}
- If 2/6 or more subjects experience DLT, $MTD = D_{i-1}$

Dose escalation stops when one-third of the subjects experience DLT at a given dose level. The MTD is then the next lower dose level.

Storer (1989) lists four single-stage designs and two two-stage designs based on the “up and down” scheme, and compares them with several simple alternatives with regard to the conservativeness of the design and with regard to the point and interval estimates of the MTD.

This approach is easy; however, it is not adjustable. The method is “designed” to estimate the MTD as a 33rd percentile, thus it is inappropriate if some other toxicity probability is targeted. The requirement that 3 subjects be exposed to the experimental treatment at each dose level leads to the inefficient use of patients. In addition, the design is conservative, it usually underestimates the MTD. Chapter 2 will describe a modified “3+3” method in comparison with other Bayesian decision rules.

1.2.3.2 Continual reassessment method (CRM)

The continual reassessment method (CRM) (O’Quigley *et al.*, 1990) is a Bayesian dose-finding algorithm that is based on a parametric model and a fixed target toxicity rate. The method continually updates our knowledge about the

1.2 Background and significance

therapy and treats patients at a level that is closest to the target level, as indicated by the current accumulated data. The model under the CRM assumes that toxicity probability is conditional on the dose and requires a set of fixed probabilities that are associated with the dose levels. The one-parameter power model for the CRM is

$$p_{\theta}(d) = c_d^{\theta}, \quad d = 1, \dots, D, \quad \theta > 0$$

where the values of c_d are monotonically increasing, prespecified (skeleton) probabilities.

The dose with the estimated mean probability of toxicity that is closest to the target level is designated for each patient cohort. A fixed sample size n is decided upon at the beginning of the study and the estimated MTD is the dose that would be allocated to patient $n+1$ were they to be included in the trial. The CRM represents one of the earliest uses of Bayesian methods in clinical trials to determine the MTD of a drug. It is flexible in allowing variations in the number of patients who are give each dose, and can integrate the dose-limiting toxicity (DLT) rates that are expected for different therapies. Extensions of the CRM include the time-to-event continual reassessment method (TITE-CRM) by Cheung *et al.* (2000), which allows patients to enter a trial without suspending accrual and thus decreasing trial duration. We will adopt the dose-assignment rule from the CRM as one alternative Bayesian decision rule to compare with other methods in Chapter 2.

1.2.3.3 Escalation with overdose control (EWOC)

Escalation with overdose control (EWOC) was first proposed by Babb *et al.* (1998). This method allows the clinician to prespecify the dose levels that will be available in a phase I trial such that the anticipated proportion of patients

who experience toxicity is set at a specified value α , which is called the feasibility bound. This is accomplished by computing the posterior cumulative distribution (CDF) of the MTD at the time of each dose assignment. For the k th dose assignment, the posterior CDF of the MTD is given by

$$p_k(\gamma) = \text{Prob}\{MTD \leq \gamma | \mathcal{D}_k\},$$

where \mathcal{D}_k denotes the data at the time of treatment for the k th patient. $p_k(\gamma)$ is the conditional probability that γ is an overdose given the data currently available. Based on this, EWOC selects for the k th patient the dose level d_k such that

$$p_k(d_k) = \alpha.$$

EWOC is designed to consider all the data that has accumulated when designating a dose for a new patient cohort. This method also controls the probability of toxicity. Neuenschwander *et al.* (2008) proposed a similar dose-assignment rule by categorizing the toxicity interval into sub-intervals and selecting the dose by maximizing the toxicity in the target interval while controlling the toxicity within the excessive and unacceptable interval. We will describe this rule as an alternative Bayesian decision rule in Chapter 2.

2

Bayesian Dose-Assignment Rules for Phase I Dose Finding in Cancer Trials

This chapter investigates four published Bayesian dose-assignment rules for phase I dose finding in cancer trials. Each rule involves decisions based on the posterior probability of toxicity. A modified “3+3” design is also introduced for comparison as it remains widely used in phase I dose finding trials. There are no new methodologies or dose-assignment rules developed in this chapter. Our purpose is to use simulation to demonstrate the appropriateness of the rules under certain scenarios. The metric we use to evaluate the performance of each rule is the percentage of trials in which the true maximum tolerated dose (MTD) is successfully declared, and the average proportion of patients assigned to each dose. The five dose-assignment rules are:

-
- *CRM_r* - the rule in the CRM design by O’Quigley *et al.* (1990) that assigns the next cohort to the dose with the posterior mean probability of toxicity closest to the target;
 - *MaxT_I* - the rule by Neuenschwander *et al.* (2008) that assigns the next cohort to the dose with the maximum posterior probability of toxicity within a target interval;
 - *ODC* - the overdose control rule by Bekele *et al.* (2010) that is based on the probability that the current dose is overly toxic;
 - *Loss* - the rule by Conaway *et al.* (2004) that defines the closeness of an estimated probability of toxicity to the target based on a loss function; and
 - the modified “3+3” algorithm.

In addition to their specific dose-assignment rules, all trials in our study are conducted as follows: The first cohort of patients are treated at the lowest dose and their clinical responses to treatment are fully observed. The probability model (or toxicity information) provides the current estimates of toxicity probability at each of the available doses. The next cohort of patients is then treated at the dose determined by the corresponding dose-assignment rule. Proceeding by sequentially updating the probability of toxicity for the available doses, the trial stops at a fixed sample size (if it is not stopped early) and provides an estimate of the MTD.

We begin this chapter by introducing two probability models on which the four Bayesian decision rules are based. Then we provide the five dose-assignment rules and universal over-dose control criteria. We proceed with the simulation studies and close this chapter with a discussion.

2.1 Two probability models

2.1.1 Bayesian beta/binomial model

The Bayesian beta/binomial model assumes that the individual probability of toxicity for each dose follows a binomial distribution. Let $\mathbf{y}_i(d)$ be the toxicity status for patient i on dose d , where $Y_i(d) = 0$ indicates no toxicity within the patient's assessment window (e.g., treatment cycle 1), and $Y_i(d) = 1$ indicates toxicity within the patient's assessment window. Define θ_d as the probability of experiencing toxicity within the patient's assessment window for dose d . Let $n(d)$ be the total number of patients on dose d . The general form of the contribution to the likelihood attributed to each patient is given by the following equation:

$$L(\mathbf{y}_i(d)) = \theta_d^{\mathbf{I}(y_i(d)=1)} [1 - \theta_d]^{\mathbf{I}(y_i(d)=0)}$$

where θ_d is *a priori* distributed as

$$\theta_d \sim \text{beta}(\alpha_d, \beta_d).$$

Define $n_1(d)$ as the number of patients experiencing toxicity within the assessment window and $n_1(d) = \sum_{i=1}^{n(d)} I(Y_i = 1)$. Since the binomial and beta distributions are conjugate, the posterior distribution of $\theta_d \mid \text{data}$ is

$$\theta_d \mid \text{data} \sim \text{beta}(\alpha_d + n_1(d), \dots, \beta_d + (n(d) - n_1(d))).$$

The above beta/binomial model provides the posterior distribution of the toxicity probability for each dose based on the current data. There remains a problem that the probability of toxicity may not be monotone across doses. In order to keep the monotone relationship of toxicity across doses, we apply an isotonic transformation (similar to that of Li *et al.* 2008) on the posterior probability of

toxicity. The *pooled adjacent violators algorithm* (PAVA; Robertson et al. 1988) is applied to borrow information across doses. To carry out the Bayesian isotonic transformation, we first draw posterior sample vectors of $\theta \mid \text{data}$ and then apply PAVA to $\theta \mid \text{data}$. Denote θ^{ord}_d as the order-restricted probability of toxicity for dose d after PAVA. The algorithm guarantees that $\theta^{ord}_{d_j} \geq \theta^{ord}_{d_i}$ for $d_j > d_i$.

2.1.2 Logistic regression model

The two-parameter logistic regression model parameterizes the probability of toxicity using the logistic transformation,

$$\text{Logit}\{\theta_d\} = \alpha + \beta \log(d/d^*), \quad \beta > 0 \quad (2.1)$$

where (α, β) are model parameters, d^* is a reference dose, and α is interpreted as the odds of toxicity at d^* . The ability to combine data from patients treated at different dose levels is a valuable feature of the regression model. This effectively enables us to “borrow” statistical strength, which can result in more accurate predictions of the future outcomes for patients who are given specific doses for the drug. Compared to the one-parameter CRM model introduced in Chapter 1, the logistic model imparts greater flexibility and more accurately depicts the dose-toxicity curve that characterizes the treatment under evaluation.

Assuming that the prior for (α, β) follows a bivariate normal distribution, the prior mean and covariance matrix can be calibrated by pre-specifying several percentiles (e.g., 2.5%, 50%, 97.5%) of toxicity probability for each dose based on empirical evidence. More details of prior tuning are described by Neuenschwander *et al.* (2008).

2.2 Five dose-assignment rules

Let $\theta^* = 0.33$ be the target toxicity rate. The four Bayesian dose-assignment rules summarize the relationship between $\theta_d \mid data$ and θ^* in different ways. The fifth “rule”, the modified “3+3” algorithm, which is not based on a probability model.

2.2.1 *CRM* by O’Quigley *et al.*

CRM, which is based on the dose-escalation concept of the CRM design proposed (O’Quigley *et al.*, 1990) assumes a monotonic dose-toxicity relationship. The “target” toxicity level was first introduced in this method for a life-threatening illness, when there is no immediate possibility of assessing the relative advantages of the treatment against its toxicities. In addition to the notion of sequentially updating the dose-toxicity relationship as new observations become available, patients are always treated at the dose that has a current probability of toxicity closest to the target. At the end of the trial, the estimated MTD is the dose that would be assigned to the next patient cohort were it be introduced in the trial.

2.2.2 *MaxTI* by Neuenschwander *et al.*

Neuenschwander *et al.* (2008) incorporated uncertainty about θ_d into their approach to phase I trial design. They were motivated by potential problems at the beginning of the trial because of the *CRM* approach of using point estimates for the probability of toxicity and ignoring the uncertainty about θ_d . Motivated by this concern, Neuenschwander *et al.* (2008), let θ_1 , θ_2 and θ_3 be the cutoff values of the toxicity probability, where $0 < \theta_1 < \theta_2 < \theta_3 < 1$. *MaxTI* classifies the probability of toxicity into four categories:

<i>Under-dosing:</i>	$\theta_d \in (0, \theta_1]$
<i>Targeted toxicity:</i>	$\theta_d \in (\theta_1, \theta_2]$
<i>Excessive toxicity:</i>	$\theta_d \in (\theta_2, \theta_3]$
<i>Unacceptable toxicity:</i>	$\theta_d \in (\theta_3, 1]$

Summarizing θ_d by toxicity intervals, a dose is designated according to the likelihood of achieving the targeted toxicity interval in addition to limiting the likelihood of prohibitive toxicity.

The rationale of controlling the overly toxic doses under *MaxTI* is comparable to that of EWOC (described in Chapter 1), where the overdose criteria is defined by a feasibility bound α as the posterior probability of exceeding the MTD. *MaxTI* is flexible in selecting appropriate toxicity intervals for the specific treatment. The upper bound for the target toxicity interval θ_2 is often set to be 0.33 to allow a maximum of 1 toxicity to occur out of 3 dose administrations estimating the MTD. The intervals for *excessive* and *unacceptable* toxicity are sometimes combined for overdose control (e.g., $\leq 70\%$).

2.2.3 ODC by Bekele *et al.*

Unlike other Bayesian dose-assignment rules, in which decisions are made upon evaluation of all available doses, decisions for *ODC* are made upon evaluation of the current dose's posterior probability of toxicity above the target (Bekele *et al.* 2010). Define $\xi_d = Pr(\theta_d > \theta^* | data)$. The method partitions the unit interval into three subintervals by the cutoff values $0 < \underline{\xi} < \bar{\xi} < 1$ and defines that toxicity at dose d is *inconsequential* if $\xi_d < \underline{\xi}$, *permissible* if $\underline{\xi} \leq \xi_d \leq \bar{\xi}$, and *prohibitive* if $\xi_d > \bar{\xi}$. The selection of cutoff values reflects the investigator's preference for a more aggressive or more guarded treatment approach. Information obtained from computer simulations should also be used to direct the selection of cutoff

2.2 Five dose-assignment rules

values so that desirable operating characteristics are achieved. Reasonable cutoff values are $0.05 \leq \underline{\xi} \leq 0.30$ and $0.70 \leq \bar{\xi} \leq 0.90$ (Bekele *et al.* 2008).

Given the above definitions, we describe below the dose-assignment rule.

- Treat the first cohort of patients at the lowest dose and fully observe the toxicity information of each cohort of patients.
- If the current dose evinces inconsequential toxicity, then escalate to the next higher dose that does not evince prohibitive toxicity. Stay at the current dose if the next higher dose evinces prohibitive toxicity.
- If the current dose evinces permissible toxicity, then stay at the current dose.
- If the current dose i evinces prohibitive toxicity and $i > 1$, then de-escalate to the highest dose less than i that is not prohibitively toxic. If $i = 1$ or all doses evince prohibitive toxicity, then stop the trial early and declare failure due to prohibitive toxicity of the drug.
- At the end of the trial, select a dose with either inconsequential or permissible toxicity having a posterior mean of θ_d closest to the target θ^* .

Other than being flexible in selecting appropriate cutoff values for overdose control, *ODC* is straightforward in making decisions based on the current dose's toxicity information with the intrinsic mechanism of "borrowing strength" across doses through the probability model. It relies more on the current dose rather than treating all doses equally (e.g., tried and untried doses). This may avoid assigning more patients to a higher untried dose that is later found to be too toxic. We will perform a simulation study to explore the impact of different values of ξ_d on the operating characteristics.

2.2.4 Loss by Conaway *et al.*

Conaway *et al.* (2004) proposed a loss function for defining how close an estimated probability of toxicity is to the target. The loss associated with an estimated $\hat{\theta}_d$ and a target probability θ^* is given by the following equation,

$$Loss(\hat{\theta}_d, \theta^*) = \begin{cases} \gamma(\theta^* - \hat{\theta}_d), & \text{if } \theta^* > \hat{\theta}_d \\ (1 - \gamma)(\hat{\theta}_d - \theta^*), & \text{if } \theta^* \leq \hat{\theta}_d. \end{cases} \quad (2.2)$$

Let $R(d)$ be the risk associated with the loss and $R(d) = E(L(\hat{\theta}_d, \theta^*)|data)$. The next cohort of patients is always treated at the dose with the lowest $R(d)$. At the end of the trial, the estimated MTD is the dose that has an estimated mean probability of toxicity that is closest to the target.

The operating characteristics of *loss* is largely decided by the value of γ in equation 2.2. A value of $\gamma = 0.5$ corresponds to a symmetric loss. An asymmetric loss penalizes toxicity greater than the target level differently than that below the target level. For each dose, the ratio of the penalty with $\{\theta : \theta < \theta^*\}$ and $\{\theta : \theta > \theta^*\}$ is $\frac{\gamma}{1-\gamma}$. A smaller γ assigns less penalty to the probability of toxicity below the target and thus makes it conservative in selecting the next dose. On the other hand, a larger γ assigns less penalty to the probability of toxicity above the target, which makes it easier to escalate doses. We will perform a simulation study to explore the different choices of γ and will recommend an appropriate γ that is a compromise between the two notions.

2.2.5 Modified “3+3” algorithm

The traditional “3+3” algorithm treats a maximum of 6 patients at any dose level. To keep the trial sample size the same across the different methods, we modify the algorithm to allow more than 6 patients to be treated at one dose:

2.3 Universal overdose control criteria

- Treat 3 subjects at the starting dose d_i .
- If $\leq 1/6$ subjects experience toxicity on current dose d_i , escalate to dose d_{i+1} .
- If between $(1/6, 1/3)$ subjects experience DLT on d_i , treat 3 more subjects at the same dose level d_i .
- If $1/3$ subjects experience DLT on d_i , treat 3 more subjects at d_i if 3 subjects have been treated on d_i , otherwise de-escalate to dose d_{i-1} .
- If $> 1/3$ subjects experience DLT on d_i , de-escalate to dose d_{i-1} or stop the trial early.

At the end of the trial, the estimated MTD is the highest dose that has a toxicity rate below 0.33.

This “3+3” approach represents a guarded clinical approach and therefore usually underestimates the MTD. Also, it estimates the MTD only at the 33rd percentile, and does not allow for adjustments when the target MTD is not at that level.

2.3 Universal overdose control criteria

In addition to each individual rule, we impose two universal overdose control criteria for the Bayesian dose-assignment rules: (i) Skipping doses when escalating is not allowed, even though the model may recommend so. (ii) For the rules without built-in overdose control criteria, the selected dose’s probability of toxicity above the target must be no larger than a threshold (i.e., $\xi_d \leq 0.8$). *MaxTI* has its own built-in criteria on ξ_d and may not be in compliance with the second criteria defined above (ii).

2.4 Simulation study

2.4.1 Scenarios

The simulation study aims to learn how the above dose-assignment rules perform under certain dose-toxicity scenarios. There are 5 dose levels in each scenario. The sample size for each trial is fixed at 30. Trials may be stopped early for futility using the overdose control criteria. Each cohort of 3 patients is fully observed before assigning patients to the next cohort. We generate 6 representative scenarios within the dose-toxicity domain, as shown in Figure 2.1. The probability of toxicity is above the target for all doses in scenario MTD-0; therefore, the overdose control criteria should stop the trial early in this scenario. For scenario MTD-1, the lowest dose is the true MTD and the probability of toxicity for all other doses is significantly above the target. For scenario MTD-2, dose 2 is the true MTD and the probability of toxicity for dose 1 and dose 3 show significant separation from that of dose 2. For scenario MTD-3, dose 3 is the true MTD and the probability of toxicity for dose 2 and dose 4 are very close to that of dose 3. For scenario MTD-4, dose 4 is the true MTD and has very low probability of toxicity; whereas dose 5 is very toxic. For scenario MTD-5, the highest dose is the true MTD.

We first elicit priors for the two model parameters, and proceed with investigating the individual rules. When the rule parameters influence the rules, such as in *MaxTI*, *ODC* and *Loss*, we change the parameters and compare the performances to capture the impact of different parameter values on the operating characteristics. The five rules are then lined up for comparison. The simulations are divided into two parts by the two probability models. The modified “3+3” algorithm, which is not based on any probability model, is introduced first.




Figure 2.1: Dose-toxicity scenarios. The dashed lines mark the target toxicity rate of 0.33. The names of the scenarios reflect the true MTD.

2.4.2 Prior specification

For simulations based on the beta/binomial model, we derive the prior probability of toxicity by specifying α_d and β_d for each dose. The prior mean probability of toxicity for each dose is listed in Table 2.1. Each prior carries information for one patient.

Table 2.1: Prior specification for $\theta_d \sim \text{beta}(\alpha_d, \beta_d)$ for the beta/binomial model

	d_1	d_2	d_3	d_4	d_5
α_d	0.128	0.159	0.172	0.210	0.241
β_d	0.871	0.841	0.827	0.790	0.759
$E(\theta_d)$	0.128	0.159	0.172	0.210	0.241

For simulations based on the logistic regression model, we assume that the priors for α and $\log\beta$ are noninformative and follow a bivariate normal distribution, as defined in Table 2.2. The 95% confidence interval, mean and median for the prior probability of toxicity are shown in Figure 2.2.

Table 2.2: Bivariate normal prior specification on $(\alpha, \log\beta)$ for the logistic regression model

$E(\alpha), E(\log(\beta))$	Standard deviation	Correlation
(-3,-0.6)	(3.16,0.76)	0.4




Figure 2.2: The 95% confidence interval, mean and median of the prior probability of toxicity for the Bayesian logistic regression model

2.4.3 Operating characteristics based on modified “3+3” algorithm

Operating characteristics for the modified “3+3” algorithm are shown in Figure 2.3. The top bars represent the average proportion of patients assigned to each dose. The bottom bars represent the MTD selection percentage for each dose. The unshaded space of the top bars marks the average proportion of patients not assigned to any dose (out of 30) due to early stopping. The unshaded space of the bottom bars marks the percentage of trials in which no dose is declared as the MTD due to early stopping.

For scenario MTD-1 where dose 1 is the true MTD, 71.5% of the trials stopped

early and failed to declare the MTD. The average total number of patients per trial for this scenario is 14.63. Overall, the algorithm performs very conservatively in dose escalation and in MTD declaration for trials for which the target toxicity rate is 0.33. Among the 6 scenarios, scenario MTD-4 yields the best performance.

2.4.4 Operating characteristics for CRM_r

Figures 2.4 and 2.5 show the operating characteristics for the CRM_r based on the beta/binomial model and the logistic regression model, respectively, with the particular priors specified in Section 2.4.2. Compared with the modified “3+3” algorithm, the CRM_r performs less conservatively. For scenario MTD-0, in which all doses are too toxic, CRM_r performs aggressively by assigning more patients to doses with higher toxicity based on the beta/binomial model (Figure 2.4). This occurs primarily because the point estimation of the posterior probability ignores uncertainty and is problematic at the beginning of the trial. For scenario MTD-4, in which the toxicity of the true MTD is very low, CRM_r performs aggressively based on the logistic regression model (Figure 2.5). The performances of CRM_r are appropriate for the other scenarios. A lower value for the overdose control criteria (e.g., $\xi_d \leq 0.6$) will make the performance less aggressive for scenarios MTD-0 and MTD-4. However, this will make the performances too conservative for the other scenarios.

2.4.5 Operating characteristics for $MaxTI$

In investigating $MaxTI$, we combine the *excessive* and *unacceptable* toxicity intervals for overdose control. Therefore only θ_1 and θ_2 need to be specified. The cutoff points for *under-dosing* and *target* toxicity intervals are ($\theta_1 = 0.167$, $\theta_2 = 0.333$).




Figure 2.3: Operating characteristics based on the modified “3+3” algorithm. The x-axis represents 6 scenarios. The top bars represent the average proportion of patients on each dose. The bottom bars represent the MTD selection percentage for each dose.




Figure 2.4: Operating characteristics for CRM based on beta/binomial model. The x-axis represents 6 scenarios. The top bars represent the average proportion of patients on each dose. The bottom bars represent the MTD selection percentage for each dose.




Figure 2.5: Operating characteristics for *CRM* based on logistic regression model. The x-axis represents 6 scenarios. The top bars represent the average proportion of patients on each dose. The bottom bars represent the MTD selection percentage for each dose.

The values for θ_1 and θ_2 are selected in alignment with the modified “3+3” algorithm for comparison purpose. The selected dose’s probability of toxicity beyond the target interval is controlled such that $\xi \leq (0.5, 0.7, 0.9)$. Note that the built-in overdose control criteria in *MaxTI* are not in compliance with the universal overdose control criteria (ii) defined in Section 2.3. Figures 2.6 and 2.7 show the operating characteristics based on the beta/binomial model and the logistic regression model, respectively, with the particular priors specified in Section 2.4.2. As shown in the figures, the higher upper-bound of ξ results in a more aggressive performance. For scenario MTD-4, in which the true MTD (dose 4) is very safe and the adjacent increased dose is very toxic, the beta/binomial model yields the ideal performance by assigning most of the patients to the true MTD and successfully declaring the MTD at the end of the trial. For scenario MTD-5, in which the highest dose is the true MTD, the Bayesian logistic regression model yields the ideal performance. *MaxTI* performs appropriately in patient allocation and dose selection overall. The choice of toxicity intervals and overdose control should be guided by the specific trial characteristics and simulation results.

2.4.6 Operating characteristics for *ODC*

Decisions based on *ODC* will escalate the dose if the probability of overdose for the current dose is below $\underline{\xi}$, and de-escalate if it is above $\bar{\xi}$. We fix $\underline{\xi}$ at 0.2 and choose $\bar{\xi} = (0.7, 0.8, 0.9)$. The operating characteristics are shown in Figure 2.8 for the beta/binomial model and in Figure 2.9 for the Bayesian logistic regression model, with the particular priors specified in Section 2.4.2. Overall, a higher value of $\bar{\xi}$ yields a more aggressive performance in patient allocation because higher ξ makes the algorithm unlikely to reject higher doses. As a special case based on the logistic regression model (Figure 2.9), in scenario MTD-4, in which

2.4 Simulation study




Figure 2.6: Operating characteristics for $MaxTI$ with different overdose control criteria based on the beta/binomial model. In each scenario, the cutoff values for θ_1 and θ_2 are fixed at $(0.167, 0.333)$. ξ values for overdose control are $(0.5, 0.7, 0.9)$. The top bars represent the average proportion of patients on each dose. The bottom bars represent the MTD selection percentage for each dose.

2.4 Simulation study




Figure 2.7: Operating characteristics for $MaxTI$ with different overdose control criteria based on the logistic regression model. In each scenario, the cutoff values for θ_1 and θ_2 are fixed at $(0.167, 0.333)$. ξ values for overdose control are $(0.5, 0.7, 0.9)$. The top bars represent the average proportion of patients on each dose. The bottom bars represent the MTD selection percentage for each dose.

the true MTD (dose 4) is very safe and the next higher dose is very toxic, the next lower dose (dose 3) is more often selected as the MTD when $\bar{\xi}$ increases. The individual trial simulation shows that dose 3 is more likely to be selected as the MTD when more patients are assigned to dose 5. This is because having more patients assigned to dose 5 allows the logistic regression model to borrow more strength from dose 5. As a consequence, the probability of toxicity for each dose is overestimated due to the high occurrence of toxicity on dose 5. Using overdose control, the posterior mean toxicity of dose 3 is therefore more likely to be the one closest to the target. *ODC* performs appropriately in patient allocation and dose selection overall.

2.4.7 Operating characteristics for *Loss*

To investigate the impact of γ on the design performance, we select γ to be (0.4, 0.5, 0.7). The operating characteristics are shown in Figure 2.10 for the beta/binomial model and in Figure 2.11 for the logistic regression model, with the particular priors specified in Section 2.4.2. Overall, the rule performs consistently in that more information from patients information on higher doses results in a lower percentage of likelihood that the higher doses are declared as the MTD. For scenarios MTD-3, MTD-4 and MTD-5 based on the beta/binomial model (Figure 2.10), where $\gamma = 0.4$, no patients are assigned to dose 5 due to a higher penalty on overdose in the loss function. For trials that declare dose 5 as the MTD in the above scenarios, toxicity information on dose 5 is based entirely on the prior and the toxicity information for other doses. Therefore, for $\gamma < 0.5$, the trial is at risk of declaring a higher dose as the MTD due to a lack of information. The MTD selection is significantly improved with a slightly higher proportion of patients assigned to dose 5 when $\gamma = 0.7$. Overall, $\gamma = 0.7$ yields the best performance




Figure 2.8: Operating characteristics for *ODC* with different $\bar{\xi}$ based on the beta/binomial model. In each scenario, $\bar{\xi} = (0.7, 0.8, 0.9)$ from left to right. The top bars represent the average proportion of patients on each dose. The bottom bars represent the MTD selection percentage on each dose.




Figure 2.9: Operating characteristics for *ODC* with different $\bar{\xi}$ based on the logistic regression model. In each scenario, $\bar{\xi} = (0.7, 0.8, 0.9)$ from left to right. The top bars represent the average proportion of patients on each dose. The bottom bars represent the MTD selection percentage on each dose.

among the three choices of γ .

This section describes an intra-rule analysis based on two probability models to give readers a perspective on each rule’s performance. Some decision rules involve specification of the parameters that summarize the posterior distribution of toxicity probability. The selection of decision parameters should be guided by trial characteristics as well as simulation outcomes to fully investigate possible scenarios. The next section describes an inter-rule comparison to better identify the appropriate rules under certain scenarios.

2.4.8 Comparison between dose-assignment rules

For comparison purposes, we select within each decision rule the parameter that yields the best overall operating characteristics, based on the particular priors specified in Section 2.4.2. For *Loss*, $\gamma = 0.7$. For *MaxTI*, $\xi = 0.7$. All the other Bayesian rules are restricted by the overdose control criteria with $\xi \leq 0.8$. The operating characteristics of the five rules for each scenario based on the beta/binomial model are described in Section 2.4.8.1; those based on the Bayesian logistic regression model are described in Section 2.4.8.2. In the figures that accompany the descriptions, each pair of bars represents the average proportion of patients assigned to each dose and the MTD selection percentage for each dose.

2.4.8.1 Comparison based on the beta/binomial model

This section summarizes the inter-rule performance based on the beta/binomial model for the 6 scenarios illustrated in Figures 2.12 through 2.17. Compared to the model-based Bayesian decision rules, the modified “3+3” algorithm performs very conservatively in patient allocation and MTD selection. For scenarios in which the lower dose is the true MTD (e.g., MTD-0 to MTD-2), *MaxTI* slightly




Figure 2.10: Operating characteristics for $Loss$ with different γ based on the beta/binomial model. In each scenario, $\gamma = (0.4, 0.5, 0.7)$ from left to right. The top bars represent the average proportion of patients on each dose. The bottom bars represent the MTD selection percentage on each dose.




Figure 2.11: Operating characteristics for *Loss* with different γ based on the logistic regression model. In each scenario, $\gamma = (0.4, 0.5, 0.7)$ from left to right. The top bars represent the average proportion of patients on each dose. The bottom bars represent the MTD selection percentage on each dose.

outperforms the other decision rules, with similar correct MTD selection percentages and conservative patient allocation. For scenarios in which the middle dose is the true MTD (e.g., MTD-3), *CRMr*, *ODC* and *Loss* achieve similar MTD selection percentages, with *ODC* assigning more patients to the true MTD. *MaxTI* achieves slightly less aggressive MTD selection by assigning more patients to higher doses. Overall, *ODC* and *MaxTI* are recommended for scenarios in which the middle dose is the true MTD. For scenarios where the higher dose is the true MTD (e.g., MTD-4 and MTD-5), the four Bayesian decision rules achieve very similar performances for scenario MTD-4. For scenario MTD-5, *CRMr* achieves a slightly higher MTD selection percentage and allocates more patients to dose 5. *ODC* and *Loss* are more conservative in patient allocation, with MTD selection percentages that are similar to that of *CRMr*.

2.4.8.2 Comparison based on the logistic regression model

This section summarizes the inter-rule performance based on the Bayesian logistic regression model for the 6 scenarios illustrated in Figures 2.18 through 2.23. The modified “3+3” algorithm continues to perform very conservatively compared to the Bayesian decision rules. For scenarios in which the lower dose is the true MTD (e.g., MTD-0 to MTD-2), *MaxTI* performs conservatively. *CRMr* and *ODC* outperform the other decision rules with similar operating characteristics. For scenarios in which the middle dose is the true MTD (e.g., MTD-3), there is no significant difference among the four Bayesian decision rules. For *ODC*, less information on the higher doses results in a higher MTD selection percentage for the doses. On the other hand, more toxicity information on the higher doses when using *Loss* results in a lower MTD selection percentage for the higher doses. For scenarios in which the higher dose is the true MTD (e.g., MTD-4 and MTD-5), *Loss* outperforms the other decision rules on MTD selection for scenario MTD-4,




Figure 2.12: Operating characteristics for the 5 rules based on the beta/binomial model for scenario-MTD0. *CRM* is the CRM rule; *MaxTI* is the rule by Neuen-schwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.13: Operating characteristics for the 5 rules based on the beta/binomial model for scenario-MTD1. *CRM* is the CRM rule; *MaxTI* is the rule by Neuen-schwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.14: Operating characteristics for the 5 rules based on the beta/binomial model for scenario-MTD2. *CRM* is the CRM rule; *MaxTI* is the rule by Neuen-schwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.15: Operating characteristics for the 5 rules based on the beta/binomial model for scenario-MTD3. *CRM* is the CRM rule; *MaxTI* is the rule by Neuen-schwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.16: Operating characteristics for the 5 rules based on the beta/binomial model for scenario-MTD4. *CRM* is the CRM rule; *MaxTI* is the rule by Neuen-schwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.17: Operating characteristics for the 5 rules based on the beta/binomial model for scenario-MTD5. *CRM* is the CRM rule; *MaxTI* is the rule by Neuen-schwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.

with similar patient allocation. All the Bayesian decision rules perform similarly in scenario MTD-5.

In this section we compared the five decision rules for dose finding in cancer trials based on two probability models. The modified “3+3” algorithm performs very conservatively in patient allocation and MTD selection. The toxicity boundaries for *MaxTI* is chosen to be (0.167, 0.33) to align with those in the “3+3” algorithm. Simulation results show that with the same toxicity boundaries, the *MaxTI* significantly outperforms the modified “3+3” algorithm in patient allocation and MTD selection. All the Bayesian decision rules perform reasonably, with slight differences in operating characteristics. We identified the appropriate decision rules to use under specific scenarios.

Compared to the beta/binomial model, the logistic regression model shows more effect in borrowing strength across doses by performing conservatively in scenarios in which a lower dose is the MTD, and aggressively in scenarios in which a higher dose is the MTD.

2.5 Discussion

This chapter has described an investigation of five dose-assignment rules used in phase I cancer trials. Other than the modified “3+3” algorithm, all Bayesian decisions are based on the beta/binomial model and the Bayesian logistic regression model. The Bayesian decision rules summarize the posterior distribution of toxicity in different ways such that the selected dose’s probability of toxicity is close enough to the target with appropriate overdose control. Both intra-rule and inter-rule comparisons were performed based on the 6 scenarios.

Based on the particular priors specified in Section 2.4.2, the modified “3+3” algorithm performs very conservatively and thus may not be desirable in cancer




Figure 2.18: Operating characteristics for the 5 rules based on the logistic regression model for scenario-MTD0. *CRM* is the CRM rule; *MaxTI* is the rule by Neuenschwander *et al.* (2008); *ODC* is the rule by Bekele *et al.* (2010); *Loss* is the rule by Conaway *et al.* (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.19: Operating characteristics for the 5 rules based on the logistic regression model for scenario-MTD1. *CRM* is the CRM rule; *MaxT* is the rule by Neuenschwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.20: Operating characteristics for the 5 rules based on the logistic regression model for scenario-MTD2. *CRM* is the CRM rule; *MaxTI* is the rule by Neuenschwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.21: Operating characteristics for the 5 rules based on the logistic regression model for scenario-MTD3. *CRM* is the CRM rule; *MaxTI* is the rule by Neuenschwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.22: Operating characteristics for the 5 rules based on the logistic regression model for scenario-MTD4. *CRM* is the CRM rule; *MaxTl* is the rule by Neuenschwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.




Figure 2.23: Operating characteristics for the 5 rules based on the logistic regression model for scenario-MTD5. *CRM* is the CRM rule; *MaxT* is the rule by Neuenschwander *et al.* . (2008); *ODC* is the rule by Bekele *et al.* . (2010); *Loss* is the rule by Conaway *et al.* . (2004); 3+3 is the modified “3+3” algorithm.

trials in which a higher toxicity rate can be tolerated. $MaxTI$, ODC and $Loss$ are flexible in attaining the desired operating characteristics, depending on how conservatively or aggressively investigators expect the trial to perform. In contrast to the point estimation of toxicity for CRM_r , the other Bayesian decision rules incorporate uncertainty on the probability of toxicity. Information obtained from historical trials and computer simulations should also be used to direct the selection of the appropriate dose-assignment rule so that desirable operating characteristics are achieved.

The next chapter introduces a new design for phase I dose finding in which the time-to-DLT data is used to allow for the late-onset toxicity to inform dose-escalation decisions. A time-to-DLT model is proposed for use along with the $MaxTI$ decision rule introduced in this chapter.

3

A Phase I Dose-Finding Design Based on a Time-to-DLT Model

3.1 Introduction

Chapter 2 investigated four Bayesian dose-assignment rules for phase I dose-finding designs in cancer trials, in which the estimation of toxicity is based on a beta/binomial model or a logistic regression model. Both of these models account for the occurrence of dose-limiting toxicity (DLT) up to a certain cycle (e.g., number of DLT occurrences at the end of cycle 1). This chapter introduces a new Bayesian adaptive design based on a time-to-DLT model, which uses the available information on all the current patients regarding time to DLT (or time to censoring) in estimating the probability of DLT at the end of cycle 1. The *MaxTI* introduced in Chapter 2 is applied in this new design as the dose-assignment rule.

The content of this chapter represents a collaborative work with Novartis Oncology under the supervision of Beat Neuenschwander, Lilla Di Scala and Ilona

Pylvaenäinen in supporting a series of trials of an investigational drug for lung cancer and breast cancer. Prior to this study, Thomas Gsponer developed simulations based on the same time-to-DLT model and a simplified dose-escalation rule. His work served as a foundation and contributed substantially to the study, which investigates the feasibility of the combination of an investigational drug with various chemotherapies based on the safety profile of the combination. The various chemotherapies are maintained at a fixed dose throughout the trial, so that dose-escalation decisions are made with respect to the investigational drug.

The study aims to understand the operating characteristics of the new design and to determine the toxicity scenarios for which the design is more appropriate than a logistic regression model. The key issues addressed by the study:

- Provide evidence of the operating characteristics of the new design under certain scenarios;
- Determine the limitations of the design and how it performs in extreme scenarios;
- Measure the performance of the time-to-DLT model as compared to the Bayesian logistic regression model to investigate:
 - Whether the time-to-DLT model is more reactive to safety than the competing model;
 - Whether the time-to-DLT model provides more information;
 - How the model picks up late-onset toxicity and the impact of this information on the results;
- Provide evidence regarding the impact of late-onset toxicity on the estimation of the DLT at the end of cycle 1.

Section 3.2 reviews the time-to-DLT model on which the decision rules are based. Section 3.3 lists the dose-escalation rule for the trial conduct. Section 3.4 provides the operating characteristics of the proposed design compared to the competing design based on a logistic regression model. Section 3.5 closes this chapter with a discussion.

3.2 Time-to-DLT model

Assume that for a fixed dose d , the time to DLT follows a Weibull distribution with the density given by

$$f(t) = \gamma \lambda t^{\gamma-1} e^{-\lambda t^\gamma}$$

with $\gamma > 0$ as the shape parameter and $\lambda > 0$ as the scale parameter. Time t is normalized by the cycle length, which is a technical choice to facilitate prior tuning and aid in the interpretation. The measure of interest will be the hazard function at time t , $h(t)$, which, in this context, corresponds to the probability that a patient experiences a DLT at time t , given that he/she did not experience a DLT prior to time t . The hazard function is given by

$$h(t) = \gamma \lambda t^{\gamma-1}.$$

We commonly assume that the hazard increases as the dose increases. Define the scale parameter $\log \lambda(d) = \log(\alpha) + \beta * \log(\frac{d}{d^*})$ with $\beta > 0$, the log-hazard at time t for dose d is then given by

$$\log(h(t, d)) = \log(\alpha) + \beta * \log(\frac{d}{d^*}) + \log(\gamma) + (\gamma - 1)\log(t), \quad \beta > 0$$

with d^* being the reference dose.

Parameter γ determines the shape of the hazard function (monotone, increasing or decreasing), α together with γ determine the hazard at the end of cycle 1 on the reference dose, and β determines the hazard ratio at any time between two different doses.

Statistical analysis of the model is achieved by a Bayesian approach as this allows for the proper quantification and combination of uncertainties of both the model parameters and the derived quantities of interest, such as the risk of DLT at various dose levels. The likelihood for this model is $h(t, d)S(t, d)$, where $S(t, d)$ is the survival function at time t for dose d .

The posterior distribution of the model parameters is then translated to the probability of DLT within the first treatment cycle (0, 1), on which the dose-escalation decision will be based:

$$P(t_{DLT} < 1|d) = 1 - S(1, d).$$

The posterior distribution of the cycle-1 DLT rate will be summarized using the following intervals:

- $[0, 0.20)$ *underdosing* toxicity;
- $[0.20, 0.35)$ *target* toxicity;
- $[0.35, 0.60)$ *excessive* toxicity;
- $[0.60, 1]$ *unacceptable* toxicity.

The choice of the above cut-off points does not depend on the model but rather needs to be discussed and selected on a trial-by-trial basis.

The priors

An uninformative bivariate normal prior for the vector $[\log(\alpha), \log(\beta)]$, and an

3.2 Time-to-DLT model

uninformative normal prior for $\log(\gamma)$, are used in the simulation study. The procedure for deriving the prior values of the model parameters is based on the prior assumptions of the probability of DLT within the first treatment cycle, as shown in Table 3.1. Hereafter this quantity will be referred to as the end-of-cycle-1 DLT rate.

Table 3.1: Prior quantiles for the end-of-cycle-1 DLT rate

	2.5%QNT	50%QNT	97.5%QNT
2.5mg/day	0.01	0.10	0.67
5mg/day	0.02	0.17	0.80
10mg/day	0.03	0.26	0.97

The parameters derived for the bivariate normal distribution of $[\log(\alpha), \log(\beta)]$ are given in Table 3.2. The parameters for the normal distribution for $\log(\gamma)$ are selected to be $Normal(0, 0.20)$. Figure 3.1 shows the prior probability of cycle-1 DLT in each toxicity interval for each dose as a consequence of the priors of the model parameters. The dose-escalation decision is based on each dose's probability in the target toxicity interval (overdose control criteria are introduced in Section 3.3.1).

Table 3.2: Parameters for bivariate normal distribution

$E(\log(\alpha))$	$E(\log(\beta))$	std $\log(\alpha)$	std $\log(\beta)$
-1.74	-0.54	1.17	0.75

Posterior sampling

This study implements a direct sampling mechanism for determining the posterior cycle-1 DLT rate (Gelman *et al.* 2004). The method approximates the distribution of a continuous parameter as a discrete distribution on a grid of points. For low-dimensional likelihood functions, this method can significantly




Figure 3.1: Prior distribution of the end-of-cycle-1 DLT rate for each dose as a consequence of the priors of the model parameters. The four toxicity intervals are tox1: *underdosing* $[0, 0.20)$, tox2: *targeted* $[0.20, 0.35)$, tox3: *excessive* $[0.35, 0.60)$ and tox4: *unacceptable* $[0.60, 1.00]$.

improve computational time as compared to Markov chain simulations. A brief description of this method is provided hereafter.

For the simplest discrete approximation, compute the target density, $p(\theta|y)$, at a set of evenly spaced values $\theta_1, \dots, \theta_N$ that cover a broad range of the parameter space for θ , then approximate the continuous $p(\theta|y)$ by the discrete density at $\theta_1, \dots, \theta_N$ with probabilities $\frac{p(\theta_i|y)}{\sum_{i=1}^N p(\theta_i|y)}$. Once the grid of density values is computed, a random draw from $p(\theta|y)$ is obtained by drawing a random sample from the uniform distribution on $[0,1]$, then transformed by the inverse cdf. method to obtain a sample from the discrete approximation (Gelman *et al.* 2004).

3.3 Study Design

3.3.1 Decision rule

The primary end point of this study is the end-of-cycle-1 DLT rate, expressed for each dose level in terms of the cycle-1 DLT rate that falls within the predefined toxicity intervals. The trials covered by this simulation study could use up to three doses in a daily regimen of the investigational drug, namely 2.5 mg/day, 5 mg/day and 10 mg/day in a 21-day cycle. In addition, for the purpose of evaluating model performance, the simulations will also investigate five doses in a daily regimen, with the addition of 1mg/day and 7.5mg/day. In order to better categorize the operating characteristics, a dummy dose (0mg) will be imposed in cases in which the first dose is found to be too toxic. The model will be able to recommend the dummy dose, thereby forcing the trial to stop immediately.

The study design uses the posterior end-of-cycle-1 DLT rate per dose level, summarized in the toxicity intervals given in Figure 3.1, to identify the feasible dose level at which the model can safely assign the succeeding patient cohort. A dose level is defined as feasible with respect to the rate of DLTs if it satisfies the following properties:

1. It maximizes the probability of the end-of-cycle-1 DLT rate within the *target* toxicity interval; and
2. It corresponds to a less than 25% probability of the end-of-cycle-1 DTL rate falling within the combination of the *excessive* and *unacceptable* toxicity intervals.

Condition (2) will be referred to as the overdose control criteria, i.e., criteria that controls the risk of *excessive* and *unacceptable* toxicity such that it remains below a specified threshold.

3.3.2 Trial conduct

Figure 3.2 illustrates the conduct of the design. At every decision point the model identifies the dose level that is currently the most feasible, by evaluating the end-of-cycle-1 DLT rate within the predefined toxicity interval. This dose level is then adjusted by the safety-related issue (i.e., no skipping a dose when escalating). The next dose level is then recommended for the next cohort of patients. This procedure is repeated until one of the prespecified stopping rules is satisfied. The trials are conducted as follows:

- Simulation starts at the time of the first patient’s first visit (FPFV). Patients are recruited at a specific rate, which is exponentially distributed with a mean of 1/7 per day. Patient accrual is suspended after the 6th patient has been enrolled. The model is first evaluated when all of the first six patients have finished cycle 1 (or experienced DLT within cycle 1). Recruitment continues thereafter without accrual suspension. Dose-escalation decisions will be made when a cycle-1 DLT happens, or two months after the last decision time point.
- If the model recommends a dose that is more than one level higher than the current dose, only the next higher dose will be assigned to the next patient cohort.
- The trial will be stopped for successfully declaring the MTD if all of the following criteria are satisfied:
 - The corresponding dose is found to be feasible according to the criteria given in Section 3.3.1;
 - At least 6 patients on this dose have finished cycle 1 (or experienced DLT);

- At least 12 patients in the trial have finished cycle 1 (or experienced DLT).
- The trial will be stopped for failure to declare the MTD if any of the following events occur before the MTD is declared:
 - The dummy dose (0mg) is recommended by the model;
 - The maximum number of patients in the trial is reached;
 - The maximum trial length is reached.

A graphical presentation of a simulated trial based on the particular priors specified in Section 3.2 is displayed in Figure 3.3. The middle plot shows the patient accrual from the first patient’s first visit and the dose assignment. The columns on the right side list the time to DLT (or time to censoring) for each patient. The plot at the bottom of the figure marks the dose-escalation points and the total number of patients assigned to each dose up to that point. The upper panel of bar graphs shows the probability of toxicity within the predefined toxicity intervals for each dose. The decision of dose escalation is based on these bar graphs. The true MTD in this simulated trial is 10mg. The trial successfully declared 10mg as the MTD after 238 days, with a total of 21 patients enrolled in the trial. Twelve patients out of the 21 enrolled in the trial were assigned to the true MTD, 3 of which experienced cycle 1 DLT.

3.4 Simulation

3.4.1 Scenarios

The underlying time-to-DLT model from which the events are simulated is a Weibull two-parameter model. The parameters of the Weibull time-to-DLT model




Figure 3.2: Trial conduct using the proposed phase-I dose-finding design based on the time-to-DLT model.




Figure 3.3: Simulated trial conduct. The bottom plot marks decision time points and total number of patients up to that point. The middle plot shows the patient accrual. The columns on the right indicate each patient's time to DLT (or censoring) and whether the DLT is a cycle-1 DLT. The upper panel of bar graphs shows the probability of toxicity within the predefined toxicity intervals for each dose.

for the scenarios are defined by fixing the probability of DLT on the reference dose at the end of cycle 1 (pd^*) and at the end of cycle 2 (pe^*). The probabilities pd^* and pe^* are used to determine α and γ :

$$\alpha = -\log(1 - pd^*),$$

and

$$\gamma = \log[-\log(1 - pe^*)] - \log(\alpha)/\log(2).$$

The parameters of the simulation scenarios are shown in Table 3.3. Within each of the three groups, the parameters are selected such that each scenario has the same end-of-cycle-1 DLT rate on the reference dose (5mg) with increasing hazard (IH), constant hazard (CH) and decreasing hazard (DH) over cycles. Parameter β is similar across the scenarios. Plots of the probability of DLT over the treatment cycles for all scenarios are shown in Figure 3.4. For each scenario, the true MTD's probability of DLT at the end of cycle 1 falls within the targeted toxicity interval. The curvature of each plot is determined by the hazard rate.

3.4.2 Operating characteristics

In order to better investigate the performance of the time-to-DLT model, a competing model is selected for comparison: the Bayesian logistic regression model introduced in Chapter 2, which uses binary toxicity data and requires full observation of each patient cohort. That model bases dose-escalation decisions on the posterior estimation of the cycle-1 DLT rate, following the same prior probabilities of toxicity that were outlined in Section 3.2 and the same rules that were outlined in Section 3.3. A metric we use to evaluate model performance is the percentage of trials in which the true MTD is successfully declared, and the average number of patients assigned to each dose. Table 3.4 illustrates the primary

3.4 Simulation




Figure 3.4: Dose-toxicity plots for three groups of scenarios. IH indicates increasing hazard, CH indicates constant hazard, and DH indicates decreasing hazard. The title of each plot indicates the true MTD, e.g., SC5-CH means that the true MTD is 5mg with constant toxicity hazard over the treatment cycles.

Table 3.3: Scenarios defined by the time-to-DLT model.

	Hazard	p_d^*	p_e^*	β	α	γ
Group 1 MTD 2.5mg/day	Increasing	0.47	0.83	1.0	0.63	1.5
	Constant	0.47	0.72	1.0	0.63	1.0
	Decreasing	0.47	0.59	1.0	0.63	0.5
Group 2 MTD 5mg/day	Increasing	0.28	0.61	1.1	0.33	1.5
	Constant	0.28	0.48	1.1	0.33	1.0
	Decreasing	0.28	0.37	1.1	0.33	0.5
Group 3 MTD 10mg/day	Increasing	0.12	0.30	1.1	0.13	1.5
	Constant	0.12	0.23	1.1	0.13	1.0
	Decreasing	0.12	0.17	1.1	0.13	0.5

input parameters for the models and trial conduct. The maximum number of patients enrolled in each trial is limited to 100 for simulation purposes only. The simulated trials rarely reach this upper limit for all scenarios.

Figures 3.5 to 3.7 illustrate the operating characteristics of the designs based on the time-to-DLT model vs. the logistic regression model. Doses of 2.5mg, 5mg and 10mg are evaluated. Patients are followed for 6 cycles, 3 cycles and 1 cycle, respectively, for the design based on the time-to-DLT model. Patients are always followed for 1 cycle for the logistic regression design. Trials always start at the lowest dose, which is 2.5mg in these scenarios. The top sections of the bars represent the average number of patients assigned to each dose. The bottom sections of the bars represent the MTD selection percentage for each dose. Each bar group represents the results from one group of scenarios in Table 3.3, with

3.4 Simulation

	Logistic model	Time-to-DLT model
How to define scenario	$P(\text{cycle-1 DLT})$	Pd^*, Pe^*, β
Max. follow-up	1 cycle	6,3,2,1 (cycles)
Min. patients on MTD	6	6
Min. patients on trial for success	12	12
Cohort size	6	Not fixed
Over-dose Criteria	25%	25%
Decision time point	After each cohort	First 6 patients finished cycle 1; Every cycle 1 DLT or 2 months from last point;

Table 3.4: Key design parameters based on the time-to-DLT model and the logistic regression model.

the first 3 bars (from the left) representing the results based on the time-to-DLT model for scenarios with increasing hazard (IH), constant hazard (CH) and decreasing hazard (DH). The 4th bars (from the left) are the results based on the logistic regression (LR) model.

There is no significant difference in the operating characteristics based on the time-to-DLT model for all the scenarios with 6-cycle follow-up or 3-cycle follow-up. For scenarios with increasing toxicity hazard, the trials based on the time-to-DLT model perform more conservatively, with longer follow-up (shown in Figure 3.5 and Figure 3.6). This is because more late-onset toxicities beyond cycle 1 contribute to a higher estimation of the DLT rate at the end of cycle 1. For scenarios with decreasing toxicity hazard and with longer follow-up, the trials based on the time-to-DLT model perform less conservatively because fewer late-onset toxicities contribute to a lower estimation of the DLT rate at the end of cycle 1. For the trials with 1-cycle follow-up (in Figure 3.7), the MTD selection percentage shows a different trend compared with those for trials with longer follow-up: the trials based on the time-to-DLT model perform less conservatively

for scenarios with increasing toxicity hazard because fewer early-onset toxicities contribute to a lower estimation of the DLT rate at the end of cycle 1. By contrast, the trials based on the time-to-DLT model perform more conservatively for scenarios with decreasing toxicity hazard because more early-onset toxicities contribute to a higher estimation of the DLT rate at the end of cycle 1.

Based on these prespecified parameters for trial conduct (e.g., accrual rate, decision points, etc.), the results indicate that a 3-cycle follow-up is appropriate for the design based on the time-to-DLT model. Thus, the estimation of the DLT rate at the end of cycle 1 combines more toxicity information from multiple cycles. The new design shows more safety benefits for trials in which more late-onset toxicities are expected. As a trade-off, the new design requires more patients on average.

Figures 3.8 to 3.10 illustrate the operating characteristics of the time-to-DLT model and the logistic regression model for which the evaluated doses are 1mg, 2.5mg, 5mg, 7.5mg and 10mg, using the scenarios that were defined in Table 3.3. Although the lowest dose for model input is 1mg, the starting dose continues to be 2.5mg. There is no significant difference in the operating characteristics based on the time-to-DLT model with 6-cycle follow-up or 3-cycle follow-up. The true toxicity hazard shape continues to contribute to the difference in patient allocation and MTD selection based on the time-to-DLT model. With the addition of a lower dose of 1mg, trials are stopped less for failure for the scenario group in which the true MTD is 2.5mg. Overall, when compared to the logistic regression model, the time-to-DLT model yields higher MTD selection percentages and requires more patients on average.




Figure 3.5: Operating characteristics of the design based on the time-to-DLT model with 6-cycle follow-up, and of the logistic regression model for 3 groups of MTD scenarios. The top sections of the bars represent the average number of patients on each dose; the bottom sections of the bars represent the MTD selection percentage for each dose. Within each group, the first 3 bars (from the left) are the results based on the time-to-DLT model with 3 different hazard shape scenarios; the 4th bars (from the left) are the results based on the logistic regression model.




Figure 3.6: Operating characteristics of the design based on the time-to-DLT model with 3-cycle follow-up, and of the logistic regression model for 3 groups of MTD scenarios. The top sections of the bars represent the average number of patients on each dose; the bottom sections of the bars represent the MTD selection percentage for each dose. Within each group, the first 3 bars (from the left) are the results based on the time-to-DLT model with 3 different hazard shape scenarios; the 4th bars (from the left) are the results based on the logistic regression model.




Figure 3.7: Operating characteristics of the design based on the time-to-DLT model with 1-cycle follow-up, and of the logistic regression model for 3 groups of MTD scenarios. The top sections of the bars represent the average number of patients on each dose; the bottom sections of the bars represent the MTD selection percentage for each dose. Within each group, the first 3 bars (from the left) are the results based on the time-to-DLT model with 3 different hazard shape scenarios; the 4th bars (from the left) are the results based on the logistic regression model.




Figure 3.8: Operating characteristics of the design based on the time-to-DLT model with 6-cycle follow-up, and of the logistic regression model for 3 groups of MTD scenarios. The top sections of the bars represent the average number of patients on each dose; the bottom sections of the bars represent the MTD selection percentage for each dose. Within each group, the first 3 bars (from the left) are the results based on the time-to-DLT model with 3 different hazard shape scenarios; the 4th bars (from the left) are the results based on the logistic regression model.




Figure 3.9: Operating characteristics of the design based on the time-to-DLT model with 3-cycle follow-up, and of the logistic regression model for 3 groups of MTD scenarios. The top sections of the bars represent the average number of patients on each dose; the bottom sections of the bars represent the MTD selection percentage for each dose. Within each group, the first 3 bars (from the left) are the results based on the time-to-DLT model with 3 different hazard shape scenarios; the 4th bars (from the left) are the results based on the logistic regression model.




Figure 3.10: Operating characteristics of the design based on the time-to-DLT model with 1-cycle follow-up, and of the logistic regression model for 3 groups of MTD scenarios. The top sections of the bars represent the average number of patients on each dose; the bottom sections of the bars represent the MTD selection percentage for each dose. Within each group, the first 3 bars (from the left) are the results based on the time-to-DLT model with 3 different hazard shape scenarios; the 4th bars (from the left) are the results based on the logistic regression model.

3.5 Discussion

This chapter introduced a phase I dose-finding design based on a time-to-DLT model. Instead of using the binary DLT data within treatment cycle 1, the time-to-DLT model associates time with toxicity information in estimating the probability of DLT at the end of cycle 1. Dose-escalation decisions are made whenever a cycle-1 DLT occurs, or 2 months after the previous decision-making point. The new model allows a DLT beyond cycle 1 to inform the dose-escalation decision. The time-to-DLT model distinguishes scenarios in which the same end-of-cycle-1 DLT rate is achieved by different hazard shapes, a differentiation that is not captured by the logistic regression model. As a trade-off, a trial based on the time-to-DLT model requires a slightly higher patient accrual to declare the MTD.

The time-associated toxicity information in this design is dichotomized based on its dose-limiting characteristics. However, toxicities in cancer trials are usually categorized into severity levels, which are not of equal importance and do not occur independently. The next chapter introduces a new design for phase I dose finding in which the toxicity is categorized by severity levels and transitions over time. The new design allows partial toxicity information to contribute to dose-escalation decisions. A new discrete-time multi-state model is proposed for use along with the *ODC* decision rule introduced in Chapter 1.

Unlike the traditional logistic regression model which can only estimate the DLT rate at the end of cycle 1 based on patients' binary toxicity data within cycle 1, the time-to-DLT model introduced in this chapter can also estimate the DLT rate for cycles that occur after cycle 1. To better characterize toxicity information for dose-escalation decisions, we can combine the toxicity information across multiple cycles to inform dose-escalation decisions. One approach is to

assign a toxicity score to correspond with each treatment cycle such that it reflects the relative importance of the DLT within that cycle. The posterior probabilities of DLT for the multiple cycles are then integrated by the toxicity scores into an average toxicity score (ATS). Dose-escalation decisions are then based on the ATS. The ATS approach is introduced in the next chapter as a manner of integrating the probability of each toxicity level as the maximum for dose-escalation decisions.

4

Adaptive Dose Finding in Phase I Cancer Trials Based on a Discrete-Time Multi-State Model

Chapter 2 investigated four Bayesian dose-assignment rules for phase I dose-finding designs in cancer trials in which decisions are made based on the posterior probability of toxicity. One of the rules, *MaxTI*, was used in the dose-finding design in Chapter 3, which applied a time-to-DLT model to estimate the dose-limiting toxicity (DLT) rate at the end of cycle 1. None of the four dose-assignment rules or the time-to-DLT model considers toxicity severity levels and inter-level transitions. This chapter introduces a phase I dose-finding design in which (i) toxicity is categorized into multiple severity levels, and (ii) patients' toxicity states are measured periodically according to the treatment cycles. As the model allows patients' early toxicity states to inform dose escalation, there is no requirement of accrual suspension between patient cohorts.

This chapter views the problem in a different way and is not a special case of any other chapters. We start the chapter with a description of the discrete-time multi-state model, followed by the dose-assignment rule. We then introduce a fully-evaluated model proposed by Bekele *et al.* (2010) as a comparison. We perform simulations to explore the operating characteristics of the two designs, and close the chapter with a brief discussion.

4.1 Introduction

Standard phase I trials dichotomize toxicity based on whether the toxicity is dose limiting. For example, if grade 4 fatigue is a dose-limiting toxicity, while fatigue of grade 1, 2, and 3 are considered not dose limiting, then grade 1 fatigue and grade 3 fatigue are treated identically. These approaches are designed to address early toxicity of a sufficiently high level. However, due to the ethical requirement to limit the number of patients in the trial who experience toxicity, the above methods discard useful information and lead to other problems. First, if a drug causes delayed toxicity and the trial's accrual rate is high, many patients may be treated at overly toxic doses before clinicians are aware that any patients have experienced toxicity. Second, different categories of toxicity are not of equal importance and do not occur independently. If several categories of toxicity are likely to occur, a model describing the transitions from one toxicity grade to another is needed.

De Moor *et al.* (1996) proposed a modification of the standard continual reassessment method (CRM) in which ordinal toxicity grade information is incorporated into the estimation of the maximum tolerated dose (MTD) using a cumulative logit model. The model considers that the increasingly severe toxicity grades associate with one another in a nesting manner, and confines the

dose-toxicity curves of each toxicity grade to like forms. Yuan *et al.* (2007) proposed a Bayesian quasi-likelihood approach based on CRM for multiple toxicity grades. Bekele and Thall (2004) developed a Bayesian method for dose finding based on a vector of correlated, ordinal-valued toxicities with severity levels that varied with the dose. Bekele et al. (2010) considered possible outcomes of toxicity with differing degrees of severity, defined appropriate scores for the degree of severity, and characterized the overall toxicity using an average toxicity score (ATS) parameter. These methods go beyond the conventional approaches and allow toxicity levels to inform dose-escalation decisions. In addition to controls for early-onset toxicity, clinical trials must address controls for late-onset toxicity.

To better characterize toxicity severity levels and inter-level transitions, we propose a phase I dose-finding design based on a discrete-time multi-state (DTMS) model. We use the term “discrete-time” to address multiple decision time points. We use “multi-state” to address multiple toxicity severity levels. The methodology is motivated by our desire to address several limitations arising from the use of established methods. First, this method accounts for the fact that toxicities can be clinically characterized over a distribution of severity levels. In addition, the method allows early lower severity levels to inform dose-escalation decisions. Last, the method allows for the enrollment of the next patient or small cohort of patients into the trial before the previous patient or cohort has been completely followed. This is possible because the Bayesian model allows us to account for uncertainty while applying sequentially adaptive decision rules.

Table 4.1 lists the key factors that highlight the proposed design compared with a fully-evaluated design by Bekele et al. (2010) and other phase I dose-finding designs that were introduced in Chapter 2. This chapter focuses on the first two designs that categorize toxicity into multiple severity levels.

4.2 Bayesian discrete-time multi-state model

Table 4.1: Phase I dose-finding methods and design factors

Design Factors	DTMS Design	Fully-Evaluated Design	3+3	Designs CRMr/MaxTI/ODC/Loss
Multi-level toxicity?	Yes	Yes	No	No
Accrual suspended between cohorts?	No	Yes	Yes	Yes
Cycles during which toxicity assessed	Multiple	1	1	1
Cohort size specified	No	No	Yes	No
Target toxicity rate	Flexible	Flexible	0.33	Flexible

4.2 Bayesian discrete-time multi-state model

Multi-state models are extensions of competing risks models, which incorporate an initial state and several mutually-exclusive absorbing states. In our design problem, toxicity is characterized by multiple severity levels, but not classified as initial or final states. We classify our model as a multi-state model that considers the transitions between multiple toxicity severity levels. In practice, the multi-state model is often considered to be a Markov model. A general interpretation of the Markov trait is that our predictions depend on history only through the present time. Applied to our multi-state model, this trait indicates that, given a patient's history and the present toxicity level, the designation of the next toxicity level, as far as how much and when, depends on only the present toxicity level. In addition to the Markov properties stated above, we propose a discrete-time multi-state model, which assumes a specific transition probability matrix for each time interval.

We define a sequence of time points such that $0 = t_0 < t_1 < \dots < t_J < \infty$, where $(t_{j-1}, t_j]$ is expressed as the time interval j and $(t_0, t_J]$ is the patients' entire assessment window with J intervals. Let Y_j take on one of K possible values $(1, \dots, K)$, where $Y_j = k$ means that the k th level of toxicity is observed during period j . Define the transition probability, i.e., the probability of transitioning

4.2 Bayesian discrete-time multi-state model

from a toxicity level k during the time interval $j - 1$ to a level k' during the time interval j , as

$$p_{j,k,k'} = Pr(Y_j = k' \mid Y_{j-1} = k).$$

For the purpose of modeling toxicity, we assume that

- 1) If $k = K$, then $p_{j,k,k'} = Pr(Y_j = K \mid Y_{j-1} = K) = 1$, which implies that the highest toxicity level K is irreversible. In cancer trials, toxicity level K usually indicates treatment cessation.
- 2) If $k < K$, then $p_{j,k,k'} = Pr(Y_j = k' \mid Y_{j-1} = k) > 0$, which implies that other than the highest level, a toxicity level can change bi-directionally.

For example, we categorize toxicity into 5 severity levels ($K = 5$) where level 1 indicates no toxicity or minimum toxicity, and level 5 indicates treatment cessation. All patients start at toxicity level 1 prior to entry into the trial. A patient receives a treatment on day 1 and is followed for 6 weeks ($J = 6$). Suppose he/she develops a level 3 toxicity during the first week. The toxicity increases to level 4 during the second week and drops back to level 3 during the third week. The toxicity then drops to level 2 during the fourth week and stays at level 2 for the remaining weeks. The accumulating data appear as $(Y_1 = 3, Y_2 = 4, Y_3 = 3, Y_4 = Y_5 = Y_6 = 2)$. Table 4.2 illustrates the observed toxicity level for each week and its contribution to the likelihood through $p_{j,k,k'}$. It is also possible that due to the pathological nature of a certain drug or treatment, patients are more likely to develop late onset toxicity, e.g., $(Y_1 = 1, Y_2 = 1, Y_3 = 2, Y_4 = 3, Y_5 = 4, Y_6 = 4)$. For trials in which patients are scheduled to receive a treatment periodically, e.g., once every week for 6 weeks, it is more likely that the observed toxicity level will jump up and down over time, e.g., $(Y_1 = 3, Y_2 = 2, Y_3 = 3, Y_4 = 2, Y_5 = 3, Y_6 = 2)$.

Let $y_{i,j} \in [1, \dots, K]$, ($i = 1, \dots, n$), denote the toxicity state of the i th patient during time period j . Let $\mathbf{y}_i = (y_{i,1}, \dots, y_{i,J})$ represent the sequence of toxicity

4.2 Bayesian discrete-time multi-state model

Table 4.2: One patient's contribution to the probability model through $p_{j,k,k'}$.

Week	0	1	2	3	4	5	6
Toxicity level	1	3	4	3	2	2	2
Contribution		$p_{1,1,3}$	$p_{2,3,4}$	$p_{3,4,3}$	$p_{4,3,2}$	$p_{5,2,2}$	$p_{6,2,2}$

states observed for the i th patient over the $[1, \dots, J]$ periods during which data are available for this patient. We use a discrete-time multi-state model and assume that transitions from a given toxicity level are independent of the history of the process prior to entry into that level. The general form of the contribution of each patient to the likelihood is as follows:

$$L(\mathbf{y}_i) = \prod_{j=1}^J \prod_{k=1}^K \prod_{k'=1}^K p_{j,k,k'}^{\mathbf{I}(y_{i,j-1}=k)\mathbf{I}(y_{i,j}=k')}$$

Let vector $\mathbf{P}_{j,k} = (p_{j,k,1}, p_{j,k,2}, \dots, p_{j,k,K})$. A *priori*, $\mathbf{P}_{j,k,k'}$ is distributed as

$$\mathbf{P}_{j,k} \sim \text{Dirichlet}(\alpha_{j,k,1}, \alpha_{j,k,2}, \dots, \alpha_{j,k,K}).$$

A Dirichlet distribution is the multivariate generalization of the beta distribution, and conjugate prior of the multinomial distribution in Bayesian statistics. In our study, its probability density function returns to the form in which the probabilities of K competing events are $p_{j,k,k'}$, given that each event has been observed $\alpha_{j,k,k'} - 1$ times.

Define $n_{j,k,k'}$ as the number of patients during time interval j who were initially in toxicity level k and transitioned to level k' with $n_{j,k,k'} = \sum_{i=1}^n I(Y_{i,j-1} = k)I(Y_{i,j} = k')$, where n is the total number of patients currently in the trial. The model assumes that

$$(n_{j,k,1}, n_{j,k,2}, \dots, n_{j,k,K}) \sim \text{Multinomial}(\sum_{k'=1}^K n_{j,k,k'}, P_{j,k}).$$

4.2 Bayesian discrete-time multi-state model

Since the multinomial and Dirichlet distributions are conjugate, the posterior distribution of $\mathbf{P}_{j,k}$ is

$$\mathbf{P}_{j,k} \mid \text{Data} \sim \text{Dirichlet}(\alpha_{j,k,1} + n_{j,k,1}, \alpha_{j,k,2} + n_{j,k,2}, \dots, \alpha_{j,k,K} + n_{j,k,K}).$$

Denote P_j as a matrix with transition probabilities $p_{j,k,k'}$. In our study we define toxicity level 1 as no toxicity or minimum toxicity and assume that all patients start with toxicity level 1, so that P_1 is a $1 \times K$ vector with transition probabilities $p_{1,1,k'}$, and P_j is a $K \times K$ matrix for $j > 1$ with transition probabilities $p_{j,k,k'}$.

The posterior distribution of the transition matrices P_j are updated as new data accumulates. These transition probabilities serve as the basis, upon which decisions are made.

The above multinomial/Dirichlet model combines the prior information and the current observed data into P_j . In addition to the model described above, we borrow an idea from Bekele *et al.* (2010), to define an average toxicity score (ATS). We describe this quantity in more detail below. There remains a need to integrate these matrices into a quantity that reflects the toxicity severity for the patient's entire assessment window. We address this need using an idea borrowed from Bekele *et al.* (2010), who defined appropriate scores for these toxicity levels and characterized the overall toxicity using an average toxicity score (ATS) parameter.

Define π_k as the probability of observing toxicity level k as the maximum toxicity level within the patient assessment window. To calculate π_k from P_j , we first define Q_k as the probability of observing less than or equal to level k toxicity within the assessment window. Q_k is expressed as

$$Q_k = P(\max\{Y_j\} \leq k \mid \text{data}) = \left\{ \prod_{j=1}^J P_j[1 : k] \right\} \{(1, \dots, 1)_k\}^T,$$

4.2 Bayesian discrete-time multi-state model

where $\{(1, \dots, 1)_k\}^T$ is a k -vector of ones, $P_j[1 : k]$ represents the first $1 \times k$ sub-vector from P_j for $j = 1$, and $P_j[1 : k]$ represents the upper-left $(k \times k)$ sub-matrix from P_j for $j > 1$. $\prod_{j=1}^J P_j[1 : k]$ is a $1 \times k$ vector, which represents the probability of observing level $(1, \dots, k)$ toxicity at the end of patient assessment window without observing greater than level k toxicity during the entire assessment window. Q_k is the sum of the elements of $\prod_{j=1}^J P_j[1 : k]$ by multiplying the k -vector of ones. We can then express π_k as

$$\pi_k = \begin{cases} Q_k & \text{if } k = 1 \\ Q_k - Q_{k-1} & \text{if } k \geq 2 \end{cases}$$

The above process combines transition matrices from all time intervals into π_k , which reflects the overall probability that toxicity level k is the maximum. By applying an approach that is similar to that of Bekele and Thall (2004), we can integrate the influence of each toxicity severity level into the decision making. Bekele and Thall (2004) depicted the clinical importance of each toxicity severity level by assigning weights to each level. Following this approach, we obtain from the clinicians a score for each toxicity severity level. Let w_k be the score obtained for toxicity severity level k , and denote $\mathbf{w} = (w_1, \dots, w_K)$, where $0 = w_1 < w_2 < \dots < w_K$. Increasing severity of toxicity is represented by increasing score values.

To measure the overall toxicity for each dose, we define the average toxicity score (ATS) at each dose as

$$ATS = \sum_{k=1}^K w_k \cdot \pi_k \tag{4.1}$$

The toxicity score w_k represents the investigator's expert opinion and is limited to positive values that are contingent upon the monotonicity constraint. In our study, we scale the toxicity score between 0 and 1.

The above process combines the transition probability matrices so as to represent the toxicity severity for the entire patient assessment window. An underlying assumption in dose-finding trials is that toxicity increases with an increasing dose; however, the *ATS* may not be isotonically ordered across doses. To address that issue, we take an isotonic regression approach similar to that of Ji *et al.* (2007a), and apply the *pooled adjacent violators algorithm* (PAVA, robertson et al. 1988) to borrow information across doses. To carry out the Bayesian isotonic transformation, we obtain posterior samples of the *ATS* for each dose from equation 4.1 and apply the PAVA to the posterior sample vectors of the *ATS*. Denote ATS_d as the order-restricted *ATS* for dose d . The algorithm guarantees that $ATS_{d_j} \geq ATS_{d_i}$ for dose levels $d_j > d_i$. De Leeuw *et al.* (2009) provided a general framework for isotonic regression, as well as implementation software. Hereafter, we will use *ATS* as an abbreviation of the order-restricted *ATS*.

The above integration and isotonic regression combine all the transition probabilities into an order-restricted average toxicity score between 0 and 1. The *ATS* can then be treated as the probability of overall toxicity and serve in that manner for the dose-assignment rule.

4.3 Dose-assignment rule

The *maximum tolerated dose (MTD)* is defined as the highest dose with ATS_d closest to and no larger than a target ATS^* . In this section we introduce a dose-escalation rule based on the posterior probability that ATS_d is greater than the target ATS^* . This rule is similar to the overdose control rule *ODC* described in Chapter 2.

The dose-escalation decision-making approach we propose borrows an idea from Bekele *et al.* (2010), who based dose-escalation decisions on the posterior

probabilities $\xi_d = Pr(ATS_d > ATS^*|data)$. We first partition the unit interval into three subintervals by the cutoff values $0 < \underline{\xi} < \bar{\xi} < 1$ and define that toxicity at dose d is *inconsequential* if $\xi_d < \underline{\xi}$, *permissible* if $\underline{\xi} \leq \xi_d \leq \bar{\xi}$, and *prohibitive* if $\xi_d > \bar{\xi}$. The selection of the cutoff values reflects the investigator's preference for a more aggressive or more guarded treatment approach. Information obtained from computer simulation should also be used to achieve desirable operating characteristics.

Given the above definitions, we describe a dose-finding rule:

- Treat the first cohort of patients at the lowest dose.
- Whenever a new cohort of patients is recruited to the trial, reassess for toxicity with the currently available data and for dose-assignment decisions to de-escalate, stay, escalate or stop.
- If the current dose evinces inconsequential toxicity, then escalate to the next higher dose. Stay at the current dose if the next higher dose evinces prohibitive toxicity.
- If the current dose evinces permissible toxicity, then stay at the current dose.
- If the current dose i evinces prohibitive toxicity and $i > 1$, then de-escalate to the highest dose less than i that is not prohibitively toxic. If $i = 1$ or all doses evince prohibitive toxicity, then stop the trial and declare failure in finding the MTD.
- Upon the trial's conclusion, select a dose with either inconsequential or permissible toxicity that has a posterior mean closest to the target π^* .

4.4 An alternative fully-evaluated model as a comparison

An alternative approach to modeling toxicity severity levels is to directly model the maximum observed toxicity level within the entire patient assessment window for each dose. The individual probabilities of each toxicity level being the maximum are then combined into an average toxicity score (*ATS*) as described above for dose escalation. The method requires the current cohort of patients to be completely observed for model evaluation.

Let \mathbf{y}_i be the maximum observed toxicity level for patient i , where $Y_i = k$ means that the k th level of toxicity is observed as the maximum. Recall that π_k is the probability of observing toxicity level k as the maximum within the patient assessment window. The general form of the contribution to the likelihood attributed to each patient is as follows:

$$L(\mathbf{y}_i) = \prod_{k=1}^K \pi_k^{\mathbf{I}(y_i=k)}$$

Let vector $\pi = (\pi_1, \dots, \pi_K)$. *A priori*, π is distributed as

$$\pi \sim \text{Dirichlet}(\beta_1, \beta_2, \dots, \beta_K).$$

Define $n_k = \sum_{i=1}^n I(Y_i = k)$. Since the multinomial and Dirichlet distributions are conjugate, the posterior distribution of π is

$$\pi \mid \text{Data} \sim \text{Dirichlet}(\beta_1 + n_1, \dots, \beta_K + n_K).$$

Once we deliver the posterior distribution for π , we apply the same procedure for averaging π into *ATS* and the isotonic regression on the posterior *ATS* for dose assignment and trial termination.

4.5 Simulation comparing the proposed design with the alternative

4.5.1 Scenarios

In investigating the operating characteristics of this design, the metrics we use to evaluate the design performance includes (i) the percentage of the trials in which the true MTD is successfully declared, (ii) the average trial length, and (iii) the average number of patients assigned to each dose. Considerations we use to evaluate operating characteristics are contained within the paragraphs that follow.

In the simulation study, we let the accrual rate r follow an exponential distribution with the mean accrual rate of 0.5 or 1 patient per week. Patients are assigned to one of three doses and followed for 6 weeks ($J = 6$). The first cohort of patients always receives the lowest dose. There are 5 toxicity severity levels ($K = 5$). The elicited score for each toxicity level is $w_1 = 0$, $w_2 = 0.25$, $w_3 = 0.5$, $w_4 = 0.75$, $w_5 = 1$. The maximum sample size is 30. The trial will stop early for toxicity when data indicate that the lowest dose is too toxic. For trials that are not stopped early, the last evaluation point occurs when all 30 patients have finished the trial.

For the purpose of generating scenarios, we define the true transition probability matrix P_j through τ ($0 \leq \tau \leq 1$). Table 4.3 shows the structure of P_j determined by τ . The rationale is that the current toxicity level at time j is more likely to stay closer to the previous level at time $j - 1$. This is a huge reduction in dimensionality from 16 to 1, which allows us to show the relationship between τ and the *ATS* through a graph. Figure 4.1 illustrates the negative correlation between τ and the *ATS*. We generate scenarios by specifying individual τ for

4.5 Simulation comparing the proposed design with the alternative

each P_j to achieve the desired toxicity hazard shape over the time intervals while controlling the true ATS within the entire assessment window.

Table 4.3: Structure of the transition probability matrix

$p_{j,k,k'}$	Tox1	Tox2	Tox3	Tox4	Tox5
Tox1	τ	$\frac{1}{2}(1 - \tau)$	$\frac{1}{4}(1 - \tau)$	$\frac{1}{6}(1 - \tau)$	$\frac{1}{12}(1 - \tau)$
Tox2	$\frac{1}{3}(1 - 0.8\tau)$	0.8τ	$\frac{1}{3}(1 - 0.8\tau)$	$\frac{1}{4}(1 - 0.8\tau)$	$\frac{1}{12}(1 - 0.8\tau)$
Tox3	$\frac{1}{6}(1 - 0.6\tau)$	$\frac{1}{3}(1 - 0.6\tau)$	0.6τ	$\frac{1}{3}(1 - 0.6\tau)$	$\frac{1}{6}(1 - 0.6\tau)$
Tox4	$\frac{1}{12}(1 - 0.4\tau)$	$\frac{1}{4}(1 - 0.4\tau)$	$\frac{1}{3}(1 - 0.4\tau)$	0.4τ	$\frac{1}{3}(1 - 0.4\tau)$
Tox5	0	0	0	0	1




Figure 4.1: True value of ATS determined by τ

For each dose, we let the Dirichlet parameters $(\alpha_{j,k,1}, \alpha_{j,k,2}, \dots, \alpha_{j,k,K})$ have the same structure as shown in Table 4.3, so that the priors for the ATS_d are

4.5 Simulation comparing the proposed design with the alternative

(0.14, 0.15, 0.18) and the effective sample size is equivalent to one.

We first divide the scenarios into two groups. For the first group, the ATS_{d1} is fixed at 0.13 and the ATS_{d3} is fixed at 0.53. The ATS_{d2} changes from 0.13 to 0.53 with a total of 11 scenarios. For the second group, the ATS_{d1} is fixed at 0.05 and the ATS_{d2} is fixed at 0.13. The ATS_{d3} changes from 0.13 to 0.53 with a total of 11 scenarios.

Within each group of scenarios, the toxicity hazard is either *constant*^(a), *decreasing*^(b), *increasing*^(c) or *scallop*^(d) over the time intervals. This is achieved by specifying the values of τ in the transition probability matrices P_j such that the corresponding toxicity hazard has the aforementioned shape with the pre-specified ATS.

4.5.2 Sample trial conduct based on the fully-evaluated design and the *DTMS* design

Figure 4.2 shows the plot of a simulated trial based on (a) the fully-evaluated design and (b) the discrete-time multi-state (*DTMS*) design. Dose 2 is the true MTD in this trial, and the average toxicity score for the three doses are (0.13, 0.25, 0.53) with decreasing toxicity hazard. The accrual rate r follows an exponential distribution with mean 1. In figure 4.1(a)(a), each cohort of 3 patients are fully observed. Dose-escalation decisions are based on the patient's maximum toxicity level within 6 weeks. The trial lasts for about 62 weeks. In figure 4.1(b)(b), the patient's weekly toxicity level is recorded for dose-escalation decisions. Accrual is not suspended between cohorts. The trial lasts for about 33 weeks. Escalation to dose 2 occurred at cohort 3 for the fully-evaluated design, and at cohort 4 for the *DTMS* design. Both designs have 2 cohorts of patients on dose 3. The fully-evaluated design has 1 more cohort of patients on the true MTD. Both designs

4.5 Simulation comparing the proposed design with the alternative

successfully declare dose 2 as the MTD.

4.5.3 Simulation results


Accrual rate $R \sim \exp(1)$

The simulations in this section are based on accrual following an exponential distribution with mean accrual rate of 1 patient per week. To investigate the performance of the DTMS design, we use the fully-evaluated design as a comparison and keep all other design parameters the same. Figures 4.3 to 4.6 illustrate the operating characteristics with $\underline{\xi} = 0.2$ and $\bar{\xi} = 0.7$. Each figure represents the simulation results from one of the four toxicity hazard shapes (i.e., constant, decreasing, increasing or scalloping). Each figure comprises two subfigures: (a) the left figure shows the results from scenario group 1, for which the ATS_{d1} is fixed at 0.13, the ATS_{d3} is fixed at 0.53 and the ATS_{d2} is between 0.13 to 0.53; and (b) the results from scenario group 2, for which the ATS_{d1} is fixed at 0.05, the ATS_{d2} is fixed at 0.13 and the ATS_{d3} is between 0.13 to 0.53. The joint bars are the respective results from the fully-evaluated design and the DTMS design. Under scenarios (a) and (b), the top panel illustrates the average proportion of patients assigned to each dose. The bottom panel illustrates the MTD selection percentage for each dose. The unshaded space in the top panel represents the average proportion of patients not assigned to any dose (out of 30) due to early stopping. The unshaded space of the bottom panel marks the percentage of trials in which no dose is declared as the MTD due to early stopping.

For (a) scenario group 1, an ideal design should yield a higher proportion of patients and selection percentage on dose 2 for the scenarios to the left of the ATS^* , and a lower proportion of patients and selection percentage on dose 2 and dose 3 to the right of the ATS^* . For (b) scenario group 2, an ideal design should

4.5 Simulation comparing the proposed design with the alternative

(a) A simulated trial based on the fully-evaluated design



(b) A simulated trial based on the DTMS design




Figure 4.2: Sample trial conduct based on the alternative designs. The x-axis indicates the trial time from the first patient's first visit (FPFV). The y-axis indicates the patient IDs. For each patient, the number at the end of week 6 indicates the maximum observed toxicity level.

4.5 Simulation comparing the proposed design with the alternative

yield a higher proportion of patients and selection percentage on dose 3 for the scenarios to the left of the ATS^* , and a lower proportion of patients and selection percentage on dose 3 to the right of the ATS^* .

The average trial duration based on the fully-evaluated design is approximately 62 weeks. The average trial duration based on the DTMS design is approximately 35 weeks. Overall, both the fully-evaluated design and the DTMS design perform reasonably. All the hazard scenarios yield the same trend in the operating characteristics: More patients are assigned to the lower doses and the MTD selection percentage gradually accumulates toward the lower doses as the higher doses become more toxic. Within scenario group 1, compared to the fully-evaluated design, the DTMS design assigns a slightly greater proportion of patients to the higher dose(s), perhaps because it uses incomplete patient toxicity information at the dose-assessment points. The two alternative designs yield almost the same MTD selection percentage. Within scenario group 2, compared to the fully-evaluated design, the DTMS design assigns a slightly greater proportion of patients to dose 3 as the ATS_{d3} increases, and yields a slightly higher MTD selection percentage on dose 3.

Based on the above observations, for trials with lower accrual and relatively longer assessment windows, the DTMS design significantly decreases trial duration and assigns a slightly greater proportion of patients to the higher doses. For scenarios in which the medium or lower doses are the true MTD (e.g., scenario group 1), the two alternative designs yield similar MTD selection percentage. For scenarios in which the higher doses are the true MTD (e.g., scenario group 2), the DTMS design yields a slightly higher MTD selection percentage.

Figure 4.7 plots the distribution of the posterior mean ATS at the end of 1000 trials for scenario group 1, with the $ATS_{d1} = 0.13$, $ATS_{d2} = 0.21$ and $ATS_{d3} = 0.53$. Figure 4.8 plots the distribution of the posterior mean ATS at

4.5 Simulation comparing the proposed design with the alternative

the end of 1000 trials for scenario group 2, with the $ATS_{d1} = 0.05$, $ATS_{d2} = 0.13$ and $ATS_{d3} = 0.21$. Both designs yield the most accurate mean estimation of the ATS on dose 2. The mean ATS for dose 1 is under estimated and the mean ATS for dose 3 is over-estimated as a result of the isotonic regression on the posterior ATS .

Accrual rate $R \sim \text{Exp}(0.5)$

The simulations in this section are based on accrual following an exponential distribution with mean accrual rate of 0.5 patient per week. All other simulation parameters are unchanged. Figures 4.9 to 4.12 illustrate the operating characteristics of the trial designs, with each figure representing the simulation results from one of the four toxicity hazard shapes. The average trial duration based on the fully-evaluated design is approximately 70 weeks. The average trial duration based on the DTMS design is approximately 63 weeks.

As patient accrual becomes slower, the DTMS design outperforms the fully-evaluated design in scenario group 1 by assigning slightly more patients to dose 1 and achieving a higher MTD selection percentage on dose 1 as dose 2 becomes more toxic. The two alternative designs yield very similar operating characteristics under scenario group 2. Figure 4.13 plots the distribution of the posterior mean ATS at the end of 1000 trials for scenario group 1 with the $ATS_{d1} = 0.13$, $ATS_{d2} = 0.21$, and $ATS_{d3} = 0.53$. Figure 4.14 plots the distribution of the posterior mean ATS at the end of 1000 trials for scenario group 2 with the $ATS_{d1} = 0.05$, $ATS_{d2} = 0.13$, and $ATS_{d3} = 0.21$.

The results show that the DTMS design offers more benefit compared to the alternative design as patient accrual becomes slower. This is because more toxicity information is able to accumulate before the time of the next dose-assignment decision.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.3: Operating characteristics of the fully-evaluated design and the DTMS design for scenarios with constant toxicity hazard. $R \sim \text{exp}(1)$. (a) Scenario group 1; ATS_{d1} is fixed at 0.13, ATS_{d3} is fixed at 0.53 and ATS_{d2} is between 0.13 and 0.53. (b) Scenario group 2; ATS_{d1} is fixed at 0.05, ATS_{d2} is fixed at 0.13 and ATS_{d3} is between 0.13 and 0.53. Under each scenario, the top panel of bars illustrates the average proportion of patients assigned to each dose for the fully-evaluated design and the DTMS design. The bottom panel of bars illustrates the MTD selection percentage on each dose.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.4: Operating characteristics of the fully-evaluated design and the DTMS design for scenarios with decreasing toxicity hazard. $R \sim \text{exp}(1)$. (a) Scenario group 1; ATS_{d1} is fixed at 0.13, ATS_{d3} is fixed at 0.53 and ATS_{d2} is between 0.13 and 0.53. (b) Scenario group 2; ATS_{d1} is fixed at 0.05, ATS_{d2} is fixed at 0.13 and ATS_{d3} is between 0.13 and 0.53. Under each scenario, the top panel of bars illustrates the average proportion of patients assigned to each dose for the fully-evaluated design and the DTMS design. The bottom panel of bars illustrates the MTD selection percentage on each dose.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.5: Operating characteristics of the fully-evaluated design and the DTMS design for scenarios with increasing toxicity hazard. $R \sim \exp(1)$. (a) Scenario group 1; ATS_{d1} is fixed at 0.13, ATS_{d3} is fixed at 0.53 and ATS_{d2} is between 0.13 and 0.53. (b) Scenario group 2; ATS_{d1} is fixed at 0.05, ATS_{d2} is fixed at 0.13 and ATS_{d3} is between 0.13 and 0.53. Under each scenario, the top panel of bars illustrates the average proportion of patients assigned to each dose for the fully-evaluated design and the DTMS design. The bottom panel of bars illustrates the MTD selection percentage on each dose.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.6: Operating characteristics of the fully-evaluated design and the DTMS design for scenarios with scalloping toxicity hazard. $R \sim \exp(1)$. (a) Scenario group 1; ATS_{d1} is fixed at 0.13, ATS_{d3} is fixed at 0.53 and ATS_{d2} is between 0.13 and 0.53. (b) Scenario group 2; ATS_{d1} is fixed at 0.05, ATS_{d2} is fixed at 0.13 and ATS_{d3} is between 0.13 and 0.53. Under each scenario, the top panel of bars illustrates the average proportion of patients assigned to each dose for the fully-evaluated design and the DTMS design. The bottom panel of bars illustrates the MTD selection percentage on each dose.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.7: Distribution of the posterior mean ATS at the end of 1000 trials for scenario group 1 with $R \sim \exp(1)$ based on the fully-evaluated design (top row 1) and the DTMS design (bottom row 2). $ATS.d_1 = 0.13, ATS.d_2 = 0.21, ATS.d_3 = 0.53$.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.8: Distribution of the posterior mean ATS at the end of 1000 trials for scenario group 2 with $R \sim \exp(1)$ based on the fully-evaluated design (top row 1) and the DTMS design (bottom row 2). $ATS.d_1 = 0.05, ATS.d_2 = 0.13, ATS.d_3 = 0.21$.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.9: Operating characteristics of the fully-evaluated design and the DTMS design for scenarios with constant toxicity hazard. $R \sim \exp(0.5)$. (a) Scenario group 1; ATS_{d1} is fixed at 0.13, ATS_{d3} is fixed at 0.53 and ATS_{d2} is between 0.13 and 0.53. (b) Scenario group 2; ATS_{d1} is fixed at 0.05, ATS_{d2} is fixed at 0.13 and ATS_{d3} is between 0.13 and 0.53. Under each scenario, the top panel of bars illustrates the average proportion of patients assigned to each dose for the fully-evaluated design and the DTMS design. The bottom panel of bars illustrates the MTD selection percentage on each dose.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.10: Operating characteristics of the fully-evaluated design and the DTMS design for scenarios with decreasing toxicity hazard. $R \sim \exp(0.5)$. (a) Scenario group 1; ATS_{d1} is fixed at 0.13, ATS_{d3} is fixed at 0.53 and ATS_{d2} is between 0.13 and 0.53. (b) Scenario group 2; ATS_{d1} is fixed at 0.05, ATS_{d2} is fixed at 0.13 and ATS_{d3} is between 0.13 and 0.53. Under each scenario, the top panel of bars illustrates the average proportion of patients assigned to each dose for the fully-evaluated design and the DTMS design. The bottom panel of bars illustrates the MTD selection percentage on each dose.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.11: Operating characteristics of the fully-evaluated design and the DTMS design for scenarios with increasing toxicity hazard. $R \sim \exp(0.5)$. (a) Scenario group 1; ATS_{d1} is fixed at 0.13, ATS_{d3} is between 0.53 and ATS_{d2} is between 0.13 and 0.53. (b) Scenario group 2; ATS_{d1} is fixed at 0.05, ATS_{d2} is fixed at 0.13 and ATS_{d3} is between 0.13 and 0.53. Under each scenario, the top panel of bars illustrates the average proportion of patients assigned to each dose for the fully-evaluated design and the DTMS design. The bottom panel of bars illustrates the MTD selection percentage on each dose.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.12: Operating characteristics of the fully-evaluated design and the DTMS design for scenarios with scalloping toxicity hazard. $R \sim \exp(0.5)$. (a) Scenario group 1; ATS_{d1} is fixed at 0.13, ATS_{d2} is between 0.13 and 0.53. (b) Scenario group 2; ATS_{d1} is fixed at 0.05, ATS_{d2} is fixed at 0.13 and ATS_{d3} is between 0.13 and 0.53. Under each scenario, the top panel of bars illustrates the average proportion of patients assigned to each dose for the fully-evaluated design and the DTMS design. The bottom panel of bars illustrates the MTD selection percentage on each dose.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.13: Distribution of the posterior mean ATS at the end of 1000 trials for scenario group 1 with $R \sim \exp(0.5)$ based on the fully-evaluated design (top row 1) and the DTMS design (bottom row 2). $ATS.d_1 = 0.13, ATS.d_2 = 0.21, ATS.d_3 = 0.53$.

4.5 Simulation comparing the proposed design with the alternative




Figure 4.14: Distribution of the posterior mean ATS at the end of 1000 trials for scenario group 2 with $R \sim \exp(0.5)$ based on the fully-evaluated design (top row 1) and the DTMS design (bottom row 2). $ATS.d_1 = 0.05, ATS.d_2 = 0.13, ATS.d_3 = 0.21$.

4.6 Sensitivity test

This section investigates the sensitivity of the DTMS design to data that are generated from models that are very different from the model described in Section 4.2. In generating these scenarios, we treat the five toxicity severity levels as independent events and specify the mean time to each of the events. The time-to-event data are generated from the exponential model with the prespecified mean. Based on the time-to-event data, if at least one toxicity level occurs within the current time interval, we take the maximum toxicity level as the observed toxicity level for the current time interval. Otherwise, the observed toxicity level will stay unchanged. The mean time to each toxicity level is selected such that the overall ATS_d for each scenario is the value listed in Table 4.4. All other trial conduct parameters stay the same as described in the previous sections. Table 4.4 lists the true ATS scenarios in column 1; the MTD is marked with an asterisk (*). The average number of patients assigned to each dose is listed in column n_d . The operating characteristics show that the DTMS model performs reasonably. For scenario 1, in which the true MTD is dose 1, dose 2 is selected with a higher percentage because the true MTD of dose 2 (0.39) is close to the toxicity target (0.33).

Table 4.4: Operating characteristics of the sensitivity test

	True ATS_d	MTD selection (%)	n_d
Scenario 1	0.17*, 0.39, 0.43	47.2*, 35.2, 9.3	0.40*, 0.39, 0.22
Scenario 2	0.13, 0.17*, 0.43	2.3, 72.6*, 21.5	0.15, 0.47*, 0.37
Scenario 3	0.13, 0.17, 0.20*	1.0, 6.8, 88.6*	0.14, 0.20, 0.66*

4.7 Discussion

This chapter introduced an adaptive phase I dose-finding design based on a new discrete-time multi-state model. The new design categorized toxicity into severity levels and allows each patient's early toxicity levels to inform dose escalation. Compared to other conventional designs that require accrual suspension between cohorts, the new design significantly shortens trial duration.

The toxicity score w_k reflects the severity of the toxicity level k relative to other toxicity levels. The average toxicity score (ATS) combines all toxicity levels via w_k to reflect the overall toxicity. The isotonic regression process guarantees that the *ATS* is monotone across doses and allows the *ATS* to borrow information across doses.

If investigators are more interested in characterizing the maximum observed toxicity level up to the current time interval, we can model the transition of the maximum observed toxicity level up to the current time interval by assuming that $p(j, k, k') = 0$ for $k > k'$. The transition probability matrices are therefore upper triangle matrices with reduced dimensions.

The DTMS model is ideal for phase I cancer trials that require longer treatment cycles or close monitoring of the toxicity each patient experiences. Computer simulations should be used to direct the accrual rate in order to achieve desirable operating characteristics. A possible expansion of this design would be to allow for a temporary suspension of accrual when predictive probabilities show prohibitively high risks of toxicity for the expected doses of future patients (Bekele *et al.* 2008).

5

Summary

This dissertation explored phase I oncology trials from three perspectives: In Chapter 2 we investigated the current alternative Bayesian dose-assignment rules based on two probability models. In Chapter 3 we described a phase I dose-finding design based on a time-to-DLT model with flexible dose-assessment points. In Chapter 4 we proposed an adaptive phase I dose-finding design based on a discrete-time multi-state model that incorporates the toxicity level and allows patients' partial toxicity information to inform decisions about dose-escalation.

Chapter 2 provided a review and comparison of the four Bayesian decision rules that summarize the posterior distribution of toxicity in different ways such that the selected dose's probability of toxicity is close enough to the target (0.33) with appropriate overdose control criteria. A modified "3+3" design was also introduced for comparison as it remains widely used in phase I dose-finding trials. Of the four Bayesian decision rules, ODC selects the next dose based on the current dose's posterior distribution of toxicity. The other three rules (CRMr, MaxTI and Loss) select the next doses based on all doses' posterior distribution of toxicity. Unlike CRRr, for which decisions are based on the point estimation

of the posterior probability of toxicity, the other three Bayesian rules incorporate uncertainty in summarizing the posterior probability of toxicity. In the simulation study, we performed intra-rule and inter-rule comparisons for the Bayesian decision rules based on the beta/binomial model and the logistic regression model, and compared the results with those from the modified “3+3” algorithm. The modified “3+3” algorithm is conservative, targeting the toxicity rate between 0.16 and 0.33; therefore, it usually underestimates the MTD. The simulation results showed that there is no single rule that outperforms the other rules under all the scenarios. Instead, the four Bayesian decision rules perform similarly in regard to their operating characteristics, with some rules slightly outperforming the others under certain scenarios, which we outlined in Chapter 2.

The dose-finding design based on the time-to-DLT model, introduced in Chapter 3, associates time with toxicity information in estimating the probability of DLT at the end of cycle 1. Dose-escalation decisions are made whenever a cycle-1 DLT occurs, or two months after the previous dose-assessment point. The simulation study showed that the estimation of the end-of-cycle-1 DLT rate is highly affected by the shape of the toxicity hazard and the follow-up time for the design based on the time-to-DLT model. For trials in which more late-onset toxicities are expected, a longer follow-up time (up to 3 cycles) yields a more conservative MTD selection. These late-onset toxicities are ignored by the design based on a logistic regression model that uses patients’ binary toxicity information up to cycle 1, which yields more aggressive MTD selection.

The dose-finding design based on the discrete-time multi-state (DTMS) model, introduced in Chapter 4, categorizes toxicity into severity levels and considers the transitions between multiple toxicity severity levels. The model allows patients’ early toxicity states to inform dose escalation. Thus there is no need to suspend the enrollment of the next patient or small cohort of patients into the

trial before the previous patient or cohort has been completely followed. Dose-escalation decisions are based on the average toxicity score, which combines all toxicity levels via the toxicity score w_k that reflects the severity of the toxicity level k relative to other toxicity levels. Compared to the fully-evaluated design that models the maximum observed toxicity level within the patients' entire assessment window, the DTMS design significantly shortens trial duration. As a trade-off, the DTMS design assigns a slightly greater proportion of patients to the higher dose(s), perhaps because it uses incomplete patient toxicity information at the dose-assessment points. The DTMS design offers more benefit compared to the alternative design as patient accrual becomes slower because more toxicity information is able to accumulate before the time of the next dose-assignment decision.

A model-based phase I dose-finding design includes two frameworks: a statistical model and a dose-escalation rule. The statistical model uses the available toxicity information to provide the current estimation of the probability of toxicity. The dose-escalation rule makes decisions based on the distribution of the toxicity. This dissertation explored phase I dose-finding designs in cancer trials from these two frameworks and proposed new designs that target the problems existing with the traditional designs. Different combinations of the alternative decision rules and the statistical models may be explored to identify the best operating characteristics for a specific trial.

6

Bibliography

- Babb J, Rogatko A, Zacks S (1998). Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine* **17**, 1103-1120.
- Babb J, Rogatko A (2001). Patient-specific dosing in cancer phase I clinical trials. *Statistics in Medicine* **20**, 2079-2090.
- Bekele BN, Thall PF (2004). Dose-finding based on multiple toxicities in a soft tissue sarcoma trial. *Journal of the American Statistical Association* **99**, 26-35.
- Bekele BN, Ji Y, Shen Y, Thall PF. (2008). Monitoring late onset toxicities in phase I trials using predicted risks. *Biostatistics* **9**, 442-457.
- Bekele BN, Li Y, Ji Y (2010). Risk-group-specific dose finding based on an average toxicity score. *Biometrics* **66**, 541-548.
- Berry DA (2004). Bayesian statistics and the efficiency and ethics of clinical trials. *Statistical Science* **19**, 175-187. .

- Berry DA (2005). Introduction to Bayesian methods III: use and interpretation of Bayesian tools in design and analysis. *Clinical Trials* **2**, 295-300.
- Berry DA (2006). Bayesian clinical trials. *Nature Reviews Drug Discovery* **5**, 27-36.
- Berry DA, Mller P, Grieve AP, Smith M, Parke T, Blazek R, Mitchard N, Krams M (2002). Adaptive Bayesian designs for dose-ranging drug trials. In Gatsonis C, Kass RE, Carlin B, Carriquiry A, Gelman A, Verdinelli I, West M (eds), *Case Studies in Bayesian Statistics, Volume V*. Springer-Verlag, New York, pp. 99-181.
- Cheung YK, Chappell R (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56**, 1177-1182.
- Cheung YK, Chappell R (2002). A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics* **58**, 671-674.
- Chevret S (2006). *Statistical Methods for Dose-finding Experiments*. John Wiley & Sons, Ltd, West Sussex, England.
- Chung SC, Schulz M (2007). Bayesian designs for clinical trials in early drug development. *Journal of Clinical Research Best Practices* **3**(6).
- Conaway MR, Dunbar S, Peddada SD (2004). Designs for single- or multiple-agent phase I trials. *Biometrics* **60**, 661-669.
- de Leeuw J, Hornik K, Mair P (2009). Isotone optimization in R: Pool-adjacent-violators algorithm (PAVA) and active set methods. *Journal of Statistical Software* **32**(5).

- De Moor CA, Higdon DM, Hilsenbeck SG, Clark GM, Von Hoff DD (1996). Incorporating toxicity severity level information in the continual reassessment method. Duke University Working Paper 96-03. Available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.56.5183>
- Gelman A, Carlin JB, Stern HS, Rubin DB (2004). *Bayesian Data Analysis, Second Edition*. Chapman and Hall/CRC, Boca Rotan, FL.
- Ji Y, Li Y, Bekele BN (2007a). Dose-finding in phase I clinical trials based on toxicity probability intervals. *Clinical Trials* **4**, 235-244.
- Ji Y, Li Y, Yin G (2007b). Bayesian dose-finding in phase I clinical trials based on a new statistical framework. *Statistica Sinica* **17**, 531-547.
- Li Y, Bekele BN, Ji Y, Cook JD (2008). Dose-schedule finding in phase I/II clinical trials using a Bayesian isotonic transformation. *Statistics in Medicine* **27**, 4895-4913.
- 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, Fisher L (1990). Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* **46**, 33-48.
- Putter H, Fiocco M, Geskus RB (2007). Tutorial in biostatistics: competing risks and multi-state models. *Statistics in Medicine* **26**, 2389-2430.
- Robertson T, Wright FT, Dykstra RL (1988). *Order Restricted Statistical Inference*. John Wiley & Sons, New York.

- Rogatko A, Tighiouart M (2007). Novel and efficient translational clinical trial designs in advanced prostate cancer. In Chung LWK, Isaacs WB, Simons JW (eds.), *Prostate Cancer: Biology, Genetics, and the New Therapeutics, Second Edition*, pp.487-495.
- Rogatko (2009) - cited in chapter 1.
- Spiegelhalter DJ, Freedman LS, Parmar MK (1994). Bayesian approaches to randomized trials (with discussion). *Journal of the Royal Statistical Society, Series A*, **157**, 357-416.
- Spiegelhalter DJ, Myles JP, Jones DR, Abrams KR (1999). An introduction to bayesian methods in health technology assessment. *BMJ*, **319**, 508-511.
- Spiegelhalter DJ, Abrams KR, Myles JP (2004). *Bayesian Approaches to Clinical Trials and Health-care Evaluation*. John Wiley & Sons, Ltd, West Sussex, England.
- Storer B (1989). Design and analysis of phase I clinical trials. *Biometrics* **45**, 925-937.
- Thall PF, Lee SJ (2003). Practical model-based dose-finding in phase I clinical trials: methods based on toxicity. *International Journal of Gynecological Cancer* **13**, 251-261.
- Thall PF, Millikan RE, Mueller P, Lee SJ (2003). Dose-finding with two agents in phase I oncology trials. *Biometrics* **59**, 487-496.
- Von Hoff DD, Kuhn J, Clark GM (1984). Design and conduct of phase I trials. In Buyse ME, Staquet MJ, Sylvester RJ (eds), *Cancer Clinical Trials: Methods and Practice*. Oxford University Press, New York, pp. 210-220.

- Whitehead J (1997). Bayesian decision procedures with applications to dose finding studies. *International Journal of Pharmaceutical Medicine* **11**, 201-208.
- Yuan Z, Chappell R, Bailey H (2007). The continual reassessment method for multiple toxicity grades: a Bayesian quasi-likelihood approach. *Biometrics* **63**, 173-179.

Vita

Lin Yang was born in Qinhuangdao, China on October 17, 1977, the Daughter of Xingfan Shi and Zeqing Yang. She entered Shanghai JiaoTong University (SJTU) in Shanghai, China. She received the degree of Bachelor of Sciences with a major in Electrical Engineering from SJTU in May, 1999. She entered Purdue University, Indianapolis in 2001 and received the degree of Master of Sciences with a major in Electrical and Computer Engineering from Purdue University in May, 2003. In August of 2005 she entered The University of Texas Health Science Center at Houston Graduate School of Biomedical Sciences. She works as a senior statistician in Novartis Institutes for Biomedical Research in Cambridge, MA since December, 2010.

Permanent address:

35 Lowell Street

Somerville, Massachusetts 02143