New Statistical Methods for Phase I Clinical Trials of a Single Agent

by

Daniel G. Muenz

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biostatistics)
in The University of Michigan
2017

Doctoral Committee:

Professor Thomas M. Braun, Co-Chair
Professor Jeremy M. G. Taylor, Co-Chair
Research Professor Mousumi Banerjee
Assistant Professor Alison M. Mondul
Research Associate Professor Matthew J. Schipper

Daniel G. Muenz

dmuenz@umich.edu

ORCID iD: 0000-0002-6701-2298

For my father z"l

Acknowledgments

With deep gratitude, I thank Drs. Thomas Braun and Jeremy Taylor, my advisors and committee co-chairs, for their extensive, thoughtful guidance. Whenever I veered off course, their insight and patience brought me back.

To my other committee members, Drs. Mousumi Banerjee, Alison Mondul, and Matthew Schipper, thank you for your valuable questions and feedback, which have honed my thinking and improved the quality of this dissertation.

I must also thank my mother, Janet Muenz, who stubbornly maintained her belief that I would eventually finish this dissertation. As usual, she was right.

This research was partially supported by National Institutes of Health grant T32 CA-83654.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

CAR conditional autoregressive

CRM continual reassessment method

DLT dose-limiting toxicity

iid independent and identically distributed

LLT low-level toxicity

LTP latent toxicity process

MH Metropolis-Hastings

MTD maximum tolerated dose

OU Ornstein-Uhlenbeck

ABSTRACT

The primary goal of phase I clinical trials in oncology is to determine a safe and possibly effective dose of a treatment, and to recommend this dose for further testing in larger trials. With chemotherapeutic treatments, the risk of severe dose-limiting toxicity (DLT) is the primary concern, assuming that the probability of efficacy will necessarily increase with dose. A phase I trial then seeks the highest dose with acceptable risk of DLT, called the maximum tolerated dose (MTD). Increasing the dose beyond the MTD would lead to unacceptable risk, while decreasing the dose would decrease the probability of benefit. In contrast, many newer therapies are molecularly targeted, where the probabilities of DLT and efficacy may plateau or even rise and then fall after some threshold. In this setting, a phase I trial must account for both toxicity and efficacy in identifying an optimal dose.

In this dissertation, we present three new approaches to modeling data in phase I trials of a single agent. Our methods improve on current practice by making more use of commonly available data. First, for chemotherapies, we investigate the utility of counting multiple DLTs per patient, in addition to counting lower-level toxicities (LLT). Typically, methods for phase I trials model only binary DLT responses, and ignore LLTs. We find that using event counts and including LLTs increases the probability of correctly identifying the MTD, particularly when the MTD is not among the highest dose levels being considered.

Second, we consider chemotherapies that are administered over multiple cycles, where dosage may vary across cycles. Multi-cycle treatments are often analyzed using only DLTs observed in the first cycle, ignoring DLTs from later cycles and

thus potentially underestimating the DLT rates. We develop a latent process model, representing a continuous level of toxicity over time, which rises and falls after each administration, but which is only observed discretely in each cycle as either no toxicity, LLT, or DLT. The process is inspired by the pharmacokinetics of drug absorption and clearance. We use our model to re-estimate the MTD at the end of adaptive trials that originally used only first-cycle data, and we find that our model typically increases the probability of correctly identifying the MTD. Our method can also recommend how to adjust a patient's dose mid-treatment, to attain a target DLT rate.

Third, we develop a method for molecularly targeted therapies, incorporating both DLTs and efficacy responses and allowing the rates of both responses to vary flexibly with dose. In particular, we adopt the conditional autoregressive model, which allows us to share information between dose levels without imposing any functional form on the dose-response curves. We find that our method can adapt to a variety of dose-toxicity and dose-efficacy patterns, and often performs at least as well as competing methods.

CHAPTER 1

Introduction

In this dissertation, we present new statistical methods for phase I clinical trials of a single agent. These small trials, which provide the first opportunity to test a new treatment in humans, traditionally seek to identify a maximum tolerated dose (MTD) from a pre-specified set of dose levels. The MTD is the highest dose with acceptable risk of dose-limiting toxicity (DLT), which is an adverse event severe enough to prevent a patient from continuing treatment. A key feature of phase I trials is that, historically, they measure only toxicity and ignore efficacy. This omission can be justified for certain classes of drugs, in particular cytotoxic agents [1], by assuming that as dose increases, so do the probabilities of both DLT and therapeutic effect. Therefore, the MTD, based solely on toxicity data, is optimal in the sense that increasing the dose beyond it would lead to unacceptable risk of toxicity, while decreasing the dose would decrease the efficacy. When a phase I trial finishes, the estimated MTD can then be further tested in a larger-scale phase II trial, in which both safety and efficacy data are collected.

Phase I trials are conducted in an iterative manner, with the following steps: (1) a small cohort (typically one to three patients) is enrolled together and assigned the same dose; (2) the cohort is followed for some time to collect toxicity data; and (3) a new cohort is enrolled at a dose level that is believed to be closer to the MTD, based on data from previous patients. The steps are repeated until either we believe

the procedure has converged to the MTD, or until a desired number of patients have been observed. When the trial is complete, we obtain a final estimate of the MTD. In essence, a phase I method consists of a decision rule, telling us which dose to assign to the next cohort and when to stop the trial, and optionally, a statistical model, that estimates DLT probabilities for each dose and provides these estimates as inputs for the decision rule. The set of dose levels to be tested is typically fixed before the trial starts.

A classic, widely employed method which uses only a rule and no statistical model, is the 3+3 method. It was first described in the context of phase I trials by Storer [2] in 1989, although its use in other contexts goes back to at least the 1940s [3]. In the 3+3 method, we enroll three patients at the lowest dose level, which is presumed safe based on extrapolation from animal models [1], and we follow each patient for the occurrence of a DLT. If no patients had a DLT, we enroll another three patients at the next higher dose level. If exactly one patient had a DLT, we enroll the next cohort of three at the current dose level. Once two or more DLTs are observed at a dose level given to three or six patients, the trial stops, and the dose just below this overly toxic one is declared the MTD.

An advantage of this method, and one reason for its longevity, is that it is simple to implement. A clinician conducting a trial can carry out the method without assistance from a statistician and without need of statistical software. Additionally, by slowly escalating the dose, the 3+3 method should result in fewer overdoses than a method that preallocates a small, fixed number of patients at each dose level. A limitation of the 3+3 method is that it does not target a particular, desired DLT probability. Instead, it chooses the dose level one below the dose which has an observed DLT proportion $\geq 1/3$, based on data from three or six patients. This point demonstrates that, without a probability model, the 3+3 method does not acknowledge that a low dose level with, say, a DLT probability of 0.1 could produce two or three DLTs

by chance, thereby invoking the stopping criterion and leading to the selection of a dose that may be subtherapeutic. Similarly, a high, unsafe dose level may by chance produce no DLTs in a small cohort, leading us to escalate to an even more unsafe dose.

To address these limitations, O'Quigley, Pepe, and Fisher [4] developed an alternative framework, using statistical modeling of DLT probabilities. This method, known as the continual reassessment method (CRM), has been extensively refined and expanded since initial publication [5–8], although the core idea remains unchanged. To begin, the CRM requires an initial guess of the DLT probability for each dose being considered. These guesses are often referred to as a "probability skeleton," which we denote by π_{0j}, for $j = 1, \ldots, J$, where J is the number of dose levels. A model is then specified where the DLT probabilities π_j at all doses are linked via a single parameter α. That is, $\pi_j = f(\alpha, \pi_{0j})$ for some function f, with a common choice for f being the power model, $f(\alpha, \pi_{0j}) = \pi_{0j}^{\exp(\alpha)}$. These probabilities feed into a Bernoulli likelihood, and estimation of α is typically carried out with Bayesian methods, providing posterior estimates of the π_j, denoted $\hat{\pi}_j$. The $\hat{\pi}_j$ are recalculated whenever a new patient outcome is collected. A Bayesian analysis is useful for the CRM in part because it allows estimation of DLT probabilities early in the trial, before any DLTs are observed. A maximum likelihood approach would not be able to estimate α in this context, although it is possible to switch to maximum likelihood once DLTs are observed [8].

Once we have the $\hat{\pi}_j$, we can use them in a decision rule to guide dose selection. For example, if we want to target a 0.3 probability of DLT, then the next patient could be assigned to the dose j where $\hat{\pi}_j$ is closest to 0.3. Other decision rules are possible with the CRM, but the key feature is to use model-based estimates of π_j to make decisions regarding the location of the MTD. This stands in contrast with the 3+3 method, where the target is, in a sense, the range from 0.17 to 0.26 [9], as

opposed to a single specifiable number. We also note that, although the CRM uses only a single parameter to characterize the dose-toxicity curve, such a parsimonious approach is sufficient for the goal of estimating the MTD [4,10]. That is, we only need to estimate the value of the dose-toxicity curve at the MTD, and a single parameter will allow this.

In this dissertation, each of the new methods we introduce builds on the CRM in an attempt to address a particular shortcoming in current phase I methodology, and to improve on current practice by making more use of commonly available data. In Chapter 2, we present methods that model the number of toxicities a patient experiences across different body systems, as opposed to the simple binary response used in the 3+3 design and CRM. Other researchers have developed methods to incorporate more complex toxicity data [11–14], but our new methods do so in a way that is both unique and appealingly simple, with one method in particular providing a closed form estimator of DLT probabilities. In Chapter 3, building on the work of Legedza and Ibrahim [15], Doussau et al. [16], Zhang and Braun [17], and Fernandes et al. [18], we construct a model for treatments that are administered over multiple cycles, where dosage may vary across cycles. Currently, multi-cycle treatments are often analyzed by ignoring the longitudinal nature of the toxicity data and using only DLTs observed in the first cycle, even though DLTs occurring later are clinically important. Finally, in Chapter 4 we develop a method to jointly optimize DLT and efficacy rates, when the assumption of strictly increasing rates of either DLT or efficacy may not hold. Other methods have been introduced to address this setting [19–25], but we believe our method is unique in that it shares information across adjacent dose levels, but does not impose any restrictions on the shape or direction of the dose-toxicity and dose-efficacy curves.

To elaborate further, in Chapter 2 we introduce three closely related models, each one building off the previous one. The setting is traditional, evaluating a patient's

response to a single administration, and assuming an increasing dose-toxicity curve. For the first model, we count the number of DLTs a patient has and model the count with a Poisson distribution whose mean increases with dose. Our goal is to estimate the probability of having any DLT (whether a patient has one DLT or five, they will by definition still have to stop their treatment), but by using counts rather than binary data we hope to provide more efficient estimation of the MTD. For models two and three, we incorporate the number of low-level toxicities (LLTs) a patient has, where we define an LLT as an event that is medically significant but not dose-limiting. Our goal is not to predict LLT counts, although technically this is possible, but rather to improve the efficiency in estimating the MTD and its corresponding DLT rate. (Note that throughout this dissertation, we use the terms "DLT rate" and "DLT probability" interchangeably). We compare the performance of our models with that of the CRM in a simulation study.

In Chapter 3, we consider a treatment that requires multiple administrations, where dosage may change in response to the previous administration. In contrast to the single endpoint used in most phase I statistical methods, here we have longitudinal toxicity data, measured in each cycle of administration. Additionally, even though we are still interested ultimately in whether a DLT will occur or not, using toxicity data with more levels (i.e., LLT or DLT) can provide helpful information. For example, if a patient has an LLT in cycle 1, they may be more likely to have a DLT in cycle 2 than if they had no toxicity at all in cycle 1. In current practice, typically only DLTs in the first cycle are modeled, allowing faster accrual of patients, but potentially causing the DLT rates to be underestimated. A retrospective analysis of multiple phase I trials found that many patients do not have a DLT until after the first cycle [26].

We propose using a latent stochastic process model for these longitudinal toxicity data where, after each administration, the process tends to increase over time (meaning higher toxicity), reach a peak, and then fall back to a baseline (perhaps as the

drug is cleared from the body). As the process rises and falls, it may cross thresholds corresponding to the observed levels of toxicity. If we use a trinary outcome (0 = no toxicity, 1 = LLT, 2 = DLT) for a treatment given in m cycles, then the observed data for each patient is (1) their sequence of outcomes (e.g., 0102 or 00010) which either has length m or terminates early with a 2, and (2) their sequence of dose assignments, which might vary due to toxicity or efficacy.

The observed outcome in each cycle is assumed to correspond to the maximum of the continuous latent toxicity process (LTP) within that cycle. That is, if the LTP spends any time above the highest threshold in a given cycle, then we will observe that cycle as a 2 (DLT). If the process always stays below the lowest threshold, we will observe a 0 (no toxicity). Otherwise, we observe a 1 (LLT).

Beyond estimating the MTD, our LTP model can answer the question of what dose to assign to a patient in their next cycle, given their data observed in earlier cycles. However, given the number of parameters in our model, the DLT rates cannot be estimated with enough precision early in the trial, when few patients have enrolled, to satisfactorily guide dose selection for the next cohort. Thus, we consider using our method after the trial has completed and sufficient data are available. The trial may have been originally run with the 3+3 method or the CRM, ignoring all but the first cycle. Our method can then re-estimate the MTD using data from all cycles.

In Chapter 4, we present a method inspired by recent developments of molecularly-targeted anti-cancer agents. The mechanism of these drugs may be such that increasing the dose beyond a threshold no longer increases the probabilities of DLT or efficacy, and perhaps after a threshold these probabilities even begin to fall [27, 28]. Thus, collecting toxicity data alone will be insufficient for selecting the optimal dose. For example, if efficacy has plateaued, we may be able to reduce the dose – and potentially the DLT rate – without sacrificing efficacy. We develop a method for this setting that incorporates DLT and efficacy response data, and that does not force a

particular shape or direction on the dose-response curves. However, crucially for our small-data context, our method still borrows data across dose levels.

In particular, we use the conditional autoregressive (CAR) model, typically used in the geospatial analysis of lattice data [29–31], and adapt it for binomial data. We apply this model twice, once for DLT rates and once for efficacy rates, where the estimation is completely separate for the two responses. As the models for DLT and efficacy rates are identical, we describe the model for DLT rates, noting that the same description holds for efficacy. In our model, the patients' binary DLT data, aggregated over patients, are binomial. The set of binomial probabilities, representing DLT rates, are transformed to be unbounded, and are given a multivariate-normal prior with a covariance structure taken from the CAR model. This model treats our set of candidate dose levels as forming a path graph, with connections between nearest dose levels in each direction, so a dose has one or two first-order "neighbors," depending on if it is at the edge or in the interior of the set of dose levels. Thus, we smooth the DLT rates across doses by borrowing information foremost from first-order neighbors, and to a lesser extent from higher-order neighbors. To select the optimal dose according to both estimated DLT and efficacy rates, we adopt the "desirability" metric of Thall and Cook [22], but modify it to promote exploration among the doses. We evaluate our method with a simulation study, in settings where the true dose-response curves either plateau, peak, or simply increase as usual.

Finally, we conclude in Chapter 5 by summarizing the contributions of our methods, and noting how they compare with existing methods.

CHAPTER 2

Modeling Adverse Event Counts in Phase I Clinical Trials of a Cytotoxic Agent

2.1 Introduction

Phase I clinical trials are small studies whose aim is to identify the maximum tolerated dose (MTD) of a treatment, i.e., the highest dose with an acceptable probability of severe dose-limiting toxicity (DLT). The selected dose is then recommended for additional, larger scale testing in a phase II trial. For cytotoxic agents such as chemotherapy, a key feature of phase I trials is that they measure only toxicity and ignore efficacy [1]. The omission of efficacy data can be justified by assuming that as dose increases, so do the probabilities of both DLT and therapeutic effect. Therefore, the MTD, based solely on toxicity data, is optimal in the sense that increasing the dose beyond it would lead to unacceptable risk of toxicity, while decreasing the dose would decrease the efficacy.

Phase I trials typically enroll patients sequentially in small cohorts, adaptively changing what dose to assign to the next cohort. At the end of the trial, one of the prespecified doses will be selected as the MTD. Two competing methods for implementing this framework are the 3+3 design [2], which uses a rule-based algorithm for escalating the dose towards the MTD, and the continual reassessment method (CRM) [4, 6–8], which adaptively updates a statistical model of DLT rates for each dose and uses

model-based estimates of those rates to select the next dose. Both the 3+3 design and the CRM use a binary indicator of whether a patient had a DLT.

While most phase I trials have used variants of the 3+3 design or CRM with binary data [32, 33], richer data are routinely collected. We refer to the "full data" as a vector of event grades for each patient, capturing all events a patient has across multiple event types, e.g., infections, gastrointestinal disorders, or hematologic events. Each grade is an integer from 0 to 5, with higher grades indicating greater severity, according to the Common Terminology Criteria for Adverse Events [34]. A DLT is commonly taken to be an event of grade 3 or higher, although this may vary by protocol.

A number of authors have explored using this richer toxicity data in different ways. Bekele and Thall [11] developed a method to incorporate the full data, adapting a multivariate ordinal probit model. The model estimates, for each dose level, the probabilities of observing each grade for every event type. Rather than base their adaptive dose-finding algorithm on DLT rates, Bekele and Thall defined a new metric, the "total toxicity burden (TTB)," as a sum of weighted event grade probabilities, with weights elicited from physicians and probabilities estimated by the model. The algorithm then seeks a pre-specified target TTB to define the MTD. The appeal of this approach is that it never dichotomizes the toxicity information, either before modeling or after in selecting the MTD, and it provides a different target than the DLT rate. A drawback, however, is that it is not straightforward to think in the TTB scale, and thus choosing and explaining the target TTB is challenging.

Yuan et al. [12] developed a quasi-continuous toxicity score to reflect the highest-grade event a patient experiences. A quasi-Bernoulli likelihood is used to carry out estimation and inform dose escalation decisions. This method worked well in simulation studies, but similar to the CRM it only uses one event per patient. Thus, Chen et al. [13] extended Yuan et al. to the case where patients may have multiple

toxicities with varying grades. They used isotonic regression to estimate the average toxicity score at each dose level. Alternative quasi-continuous toxicity scores have been proposed by Ezzalfani et al. [35] and Wang and Ivanova [36]. Lee et al. [37] proposed using multiple constraints on the percentiles of a toxicity score, as opposed to the typical single constraint of the score being less than a threshold. Their multi-constraint framework can work with any of the aforementioned scores. As with Bekele et al., using these methods with novel scores requires redefining the MTD.

Van Meter et al. [14] proposed using the original ordinal adverse event grades (0 to 5) in a proportional odds model, finding similar performance to the CRM in how often the correct MTD was selected. As with Yuan et al., the response for each patient was a single grade, which could be taken as the maximum grade across all the patient's events.

We propose three new models that make use of phase I data in a way that we have not previously seen in the literature. Specifically, one model requires all of a patient's events to be categorized as DLTs or not, and models the total count of DLTs per patient. Our two other models require that events be categorized as DLTs or low-level toxicities (LLTs), which we define as events of medical significance but which should not prevent a patient from continuing treatment, and we model the total number of DLTs and LLTs per patient. As a motivation for our model, we note that in a review of 54 single-agent phase I trials, Postel-Viney et al. [38] found a total of 2,084 patients experienced 24,918 adverse events, including 189 patients who experienced 300 DLTs. These numbers imply that, on average, patients experienced 12 events, and that 9% of patients had an average of 1.6 DLTs. Thus, data with multiple events and in particular multiple DLTs per patient are available. Our method uses the same adaptive dose-finding framework as the CRM, but incorporates these richer data in an attempt to better estimate the DLT rates.

2.2 Data and Notation

Let N_i be the DLT count, and M_i the total event count (DLTs and LLTs) for patient i, so the patient experiences a DLT if $N_i > 0$. These counts are obtained from the full data, with a clinician categorizing each adverse event as a DLT or LLT. A patient may experience multiple adverse events, both across and within event types, so in principle N_i and M_i are unbounded above. With data from n patients, let $N_+ = \sum_{i=1}^{n} N_i$ and $M_+ = \sum_{i=1}^{n} M_i$. Let $X_i \in \{1, \ldots, J\}$ be the dose level assigned to patient i among J pre-specified levels. We define $r_+ = \sum_{i=1}^{n} r_{X_i}$ as the total scaled dose given to patients. The scaled dose values $r_1, \ldots, r_J$, will be defined in the next section.

A common feature in phase I designs is the need for investigators to specify, *a priori*, a so-called probability skeleton. This represents the best prior guess for the DLT rate for each candidate dose, which we denote π_{0j}, $j = 1, \ldots, J$. To ensure that the probability of DLT increases monotonically with dose, we require $\pi_{01} < \cdots < \pi_{0J}$.

2.3 Models and Methods

In this section we specify our three models, in increasing order of complexity, and we show how each model is used to estimate DLT rates.

2.3.1 Model 1: Using DLT Counts

Our first model uses only DLT counts, not LLT counts. We model the DLT count as an increasing function of dose, $N_i|\beta, X_i{=}j \sim \text{Poisson}(\beta r_j)$, where r_j is an increasing function of j that rescales the doses to agree with the skeleton. We will show later how to calculate r_j. As with the CRM, our model has only one parameter, β, which we give a Bayesian prior, $\beta \sim \text{Gamma}(\sigma_\beta^{-2}, \sigma_\beta^{-2})$, such that, *a priori*, $\text{E}(\beta) = 1$ and $\text{sd}(\beta) = \sigma_\beta$.

According to Model 1, the probability of a patient having any DLTs at dose j is

$$\pi_j(\beta) = \Pr(N_i > 0|\beta, j) = 1 - \Pr(N_i = 0|\beta, j) = 1 - e^{-\beta r_j}.$$

Thus, we equate the prior mean of $\pi_j(\beta)$ with the skeleton, $\pi_{0j} = \mathrm{E}[\pi_j(\beta)] = 1 - \left(1 + r_j\sigma_\beta^2\right)^{-\sigma_\beta^{-2}}$. Solving this equation for r_j gives $r_j = \sigma_\beta^{-2}\Big[(1-\pi_{0j})^{-\sigma_\beta^2} - 1\Big]$. The hyperparameter σ_β controls the amount of prior uncertainty around the skeleton.

2.3.2 Model 1: Likelihood and Estimation

The likelihood for data from n patients is $L_\beta = \prod_{i=1}^n e^{-\beta r_{X_i}}(\beta r_{X_i})^{N_i}/N_i! \propto e^{-\beta r_+}\beta^{N_+}$. Using a gamma prior $p(\beta) \propto e^{-\sigma_\beta^{-2}\beta}\beta^{\sigma_\beta^{-2}-1}$, the posterior distribution of β is $\beta|Data \sim \mathrm{Gamma}(\tilde{a}_\beta, \tilde{b}_\beta)$, where $\tilde{a}_\beta = \sigma_\beta^{-2} + N_+$ and $\tilde{b}_\beta = \sigma_\beta^{-2} + r_+$. We estimate the DLT rate, specifically the probability of having any DLTs at dose j, denoted $\hat{\pi}_j$, as the posterior mean of $\pi_j(\beta)$: $\hat{\pi}_j = \mathrm{E}[\pi_j(\beta)|Data] = 1 - (1 + r_j/\tilde{b}_\beta)^{-\tilde{a}_\beta}$. Thus, we have a closed-form Bayesian estimator for DLT rates, and no numerical integration or Markov chain Monte Carlo is necessary.

2.3.3 Model 2: Using DLT and LLT Counts

Our second model incorporates more information than Model 1. First, we use a Poisson distribution to model a patient's total event count M_i, such that $M_i|\theta, X_i = j \sim \mathrm{Poisson}(\theta r_j)$, $\theta > 0$, where again r_j rescales the dose to agree with the skeleton. Note that we reuse the symbol r_j to emphasize the similarity with Model 1, but the calculation of r_j is separate for the two models.

Next, conditional on M_i, the number of DLTs N_i is given a binomial distribution, $N_i|q, M_i, X_i \sim \mathrm{Binomial}(M_i, q)$. The parameter q is the probability that a given event is a DLT rather than a LLT. Note that, although we condition on dose X_i, we do not include them on the right-hand side. Thus we are assuming that the proportion of DLTs is constant across doses; we will relax this assumption in Model 3. We assume

that θ and q are *a priori* independent with marginal priors $\theta \sim \text{Gamma}(\sigma_\theta^{-2}, \sigma_\theta^{-2})$ and $q \sim \text{Beta}(a_q, b_q)$.

The modeled probability of a patient having any DLTs at dose j is

$$\begin{aligned}
\pi_j(\theta, q) &= \Pr(N_i > 0 | \theta, q, j) \\
&= 1 - \sum_{m=0}^{\infty} \Pr(N_i = 0 | q, m, j) \Pr(M_i = m | \theta, j) \\
&= 1 - \sum_{m=0}^{\infty} (1-q)^m e^{-\theta r_j} (\theta r_j)^m / m! \\
&= 1 - e^{-\theta q r_j}.
\end{aligned}$$

Thus, we calculate r_j by solving the equation

$$\begin{aligned}
\pi_{0j} &= \text{E}[\pi_j(\theta, q)] = 1 - \text{E}[e^{-\theta q r_j}] \\
&= 1 - \int_0^1 \left[\int_0^\infty e^{-\theta q r_j} p(\theta)\, d\theta \right] p(q)\, dq \\
&= 1 - \int_0^1 \left(1 + q r_j \sigma_\theta^2\right)^{-\sigma_\theta^{-2}} p(q)\, dq \\
&= 1 - \frac{{}_2F_1\left(b_q, \sigma_\theta^{-2}; a_q + b_q; r_j / [r_j + \sigma_\theta^{-2}]\right)}{(1 + r_j \sigma_\theta^2)^{\sigma_\theta^{-2}}}
\end{aligned}$$

where ${}_2F_1$ is the Gaussian hypergeometric function [39], and $p(\cdot)$ denotes a density. This equation can be solved numerically for r_j, for any choice of skeleton π_{0j} and hyperparameters σ_θ, a_q, and b_q.

2.3.4 Model 2: Likelihood and Estimation

With data for n patients, the full likelihood is

$$L_{\theta,q} = \prod_{i=1}^{n} \Pr(M_i, N_i | \theta, q, X_i) = \prod_{i=1}^{n} \Pr(M_i | \theta, X_i) \prod_{i=1}^{n} \Pr(N_i | q, M_i, X_i) = L_\theta L_q$$

where $L_\theta = \prod_{i=1}^{n} e^{-\theta r_{X_i}} (\theta r_{X_i})^{M_i} / M_i! \propto e^{-\theta r_+} \theta^{M_+}$ and $L_q \propto \prod_{i=1}^{n} q^{N_i} (1-q)^{M_i - N_i} = q^{N_+} (1-q)^{M_+ - N_+}$. Due to the separability of the likelihood and the prior independence of the parameters, we can estimate θ separately from q. Using a gamma prior for

θ, we have $\theta|Data \sim \text{Gamma}(\tilde{a}_\theta, \tilde{b}_\theta)$, where $\tilde{a}_\theta = \sigma_\theta^{-2} + M_+$ and $\tilde{b}_\theta = \sigma_\theta^{-2} + r_+$. For q, using a beta prior gives us $q|Data \sim \text{Beta}(\tilde{a}_q, \tilde{b}_q)$, where $\tilde{a}_q = a_q + N_+$ and $\tilde{b}_q = b_q + M_+ - N_+$.

Finally, we estimate the posterior DLT rate at dose j as $\hat{\pi}_j = \text{E}[\pi_j(\theta, q)|Data] = 1 - {}_2F_1(\tilde{b}_q, \tilde{a}_\theta; \tilde{a}_q + \tilde{b}_q; r_j/[r_j + \tilde{b}_\theta])/(1 + r_j/\tilde{b}_\theta)^{\tilde{a}_\theta}$. In the R programming language, [40] the function ${}_2F_1$ is available in the `gsl` package [41] under the name `hyperg_2F1`.

2.3.5 Model 3: Using DLT and LLT Counts, Version Two

Our third model is very similar to Model 2. The key difference is that q, the probability of an event being a DLT rather than an LLT, is assumed to vary with dose. Specifically, we model $q = q_j(\alpha) = \text{expit}(-3 + e^\alpha r_j)$. The requirement that the coefficient of r_j is positive ensures that the log-odds of an event being a DLT increases with increasing dose level j. We now place a prior on α instead of q: $\alpha \sim \text{Normal}(\alpha_0, \sigma_\alpha^2)$.

According to Model 3, the probability of a patient having any DLTs at dose j is $\pi_j(\theta, \alpha) = 1 - e^{-\theta q_j(\alpha) r_j}$. The rescaled doses r_j are found by numerically solving $\pi_{0j} = \text{E}[\pi_j(\theta, \alpha)] = 1 - \text{E}[e^{-\theta q_j(\alpha) r_j}] = 1 - \text{E}[(1 + q_j(\alpha) r_j \sigma_\theta^2)^{-\sigma_\theta^{-2}}]$, where the last expectation is with respect to the distribution of α only.

2.3.6 Model 3: Likelihood and Estimation

The likelihood factors into separate functions of the two parameters, $L_{\theta,\alpha} = L_\theta L_\alpha$. The function L_θ is the same as in Model 2; the function L_α is

$$\begin{aligned} L_\alpha &= \prod_{i=1}^{n} \Pr(N_i|q, M_i, X_i) \\ &\propto \prod_{i=1}^{n} \left(\frac{e^{-3+e^\alpha r_{X_i}}}{1 + e^{-3+e^\alpha r_{X_i}}} \right)^{N_i} \left(\frac{1}{1 + e^{-3+e^\alpha r_{X_i}}} \right)^{M_i - N_i} \\ &\propto \exp\left(e^\alpha \sum_{i=1}^{n} r_{X_i} N_i \right) \prod_{i=1}^{n} \left(1 + e^{-3+e^\alpha r_{X_i}}\right)^{-M_i}. \end{aligned}$$

Due to the separability of the likelihood and the prior independence of the parameters, we can estimate θ separately from α. The posterior for θ is the same as in Model 2: $\theta|Data \sim \text{Gamma}(\tilde{a}_\theta, \tilde{b}_\theta)$, where $\tilde{a}_\theta = \sigma_\theta^{-2} + M_+$ and $\tilde{b}_\theta = \sigma_\theta^{-2} + r_+$. For α, however, a closed-form posterior is not available. To obtain the posterior, we use a simple discrete grid-based approximation, which is computationally efficient for a single parameter [42].

Finally, we estimate the posterior DLT rate at dose j as $\hat{\pi}_j = \text{E}[\pi_j(\theta, \alpha)|Data] = 1 - \text{E}[e^{-\theta q_j(\alpha) r_j}|Data] = 1 - \text{E}[(1 + q_j(\alpha) r_j/\tilde{b}_\theta)^{-\tilde{a}_\theta}|Data]$, where the final expectation is with respect to the posterior distribution of α.

2.3.7 Dose-Finding Algorithm

As a trial is run, patients are enrolled sequentially, and with data from each new patient we update our estimates of the DLT rates. Although different dose-finding algorithms are possible using our models, we describe a simple one, which we will evaluate in the simulation study. First, we choose a set of J dose levels, and we set a target DLT rate τ, denoting the maximum acceptable probability of a patient having any DLTs. We enroll the first patient at the lowest dose, level 1, and follow the patient for a desired time. Using the patient's data and one of our three models, we estimate the DLT rates $\hat{\pi}_1, \ldots, \hat{\pi}_J$. Our updated estimate of the MTD, denoted $\widehat{\text{MTD}}$, is the dose j that minimizes $|\hat{\pi}_j - \tau|$. The next patient is then assigned either to $\widehat{\text{MTD}}$ or to the lowest dose not yet assigned, whichever is smaller. This rule simply prevents the sequence of dose assignments from skipping over untested dose levels. We continue to enroll new patients and update $\widehat{\text{MTD}}$ using all the patients' data. Once we have observed a desired number of patients, the trial concludes and we obtain a final estimate of the MTD.

2.4 Simulation Study and Comparison with CRM

We simulate adaptive trials using each of our proposed models, and we record how often the true MTD is selected. To set a reference point for the performance of our models, we include a comparison with a CRM model, described in the next section.

2.4.1 CRM model for comparison

For comparison, we will use a CRM model that is related to all of our proposed models. Specifically, let $Y_i = I(N_i > 0)$ be a binary indicator of whether patient i had a DLT. Recall that, according to our three models, the probability of DLT is, respectively, $\pi_j(\beta) = 1 - e^{-\beta r_j}$, $\pi_j(\theta, q) = 1 - e^{-\theta q r_j}$, and $\pi_j(\theta, \alpha) = 1 - e^{-\theta q_j(\alpha) r_j}$. For the CRM, with a single parameter ϕ, we have to specify a model for $\Pr(Y_i = 1|\phi, j) = \Pr(N_i > 0|\phi, j)$ directly. We use $\Pr(Y_i = 1|\phi, j) = 1 - e^{-\phi r_j}$, $\phi > 0$, due to its functional similarity to our models, which will facilitate jointly calibrating the hyperparameters of each model. Our prior on ϕ is $\phi \sim \text{Gamma}(\sigma_\phi^{-2}, \sigma_\phi^{-2})$, and the rescaled dose is $r_j = \sigma_\phi^{-2}\left[(1 - \pi_{0j})^{-\sigma_\phi^2} - 1\right]$. The posterior of ϕ is obtained with a discrete grid-based approximation, which is then used to calculate the posterior DLT rates $\hat{\pi}_j = 1 - \text{E}[e^{-\phi r_j}|Data]$.

2.4.2 Simulating a Patient's Toxicity Data

In our simulation study, for the data-generating mechanism we assume that there are $K = 15$ distinct event types, each of which can experience either no event, an LLT, or a DLT, and there is a patient-specific frailty influencing the probabilities of those outcomes. The observed data for a patient are the total number of DLTs and LLTs across the K event types. We chose to use $K = 15$ as a middle ground among the number of event types we have seen in published trials [11, 38, 43], however we will also conduct a small sensitivity analysis with $K = 7$.

Letting $Y_{ik} \in \{0 = \text{no toxicity}, 1 = \text{LLT}, 2 = \text{DLT}\}$ be the response for patient i in event type $k = 1, \ldots, K$, we use a cumulative logit model, $\operatorname{logit} \Pr(Y_{ik} \geq y | X_i) = \eta_{iy} = \alpha_y + \gamma_i + r_{X_i}$ for $y = 1, 2$, where $\gamma_i \sim \text{Normal}(0, \gamma_{\text{sd}}^2)$ is a random patient-specific frailty; X_i is the dose; and r_{X_i} rescales the dose to obtain a desired probability of DLT. The multinomial probabilities $p_{iy} = \Pr(Y_{ik} = y | X_i)$ implied by this model are $p_{i0} = 1 - \operatorname{expit}(\eta_{i1})$, $p_{i1} = \operatorname{expit}(\eta_{i1}) - \operatorname{expit}(\eta_{i2})$, and $p_{i2} = \operatorname{expit}(\eta_{i2})$. Note that the linear predictor η_{iy} does not depend on the event type k, that is, for simplicity we make all event types have the same multinomial probabilities.

The number of DLTs for patient i is $N_i = \sum_{k=1}^{K} I(Y_{ik} = 2) \sim \text{Binomial}(K, p_{i2})$ and the total event count is $M_i = \sum_{k=1}^{K} I(Y_{ik} \in \{1, 2\}) \sim \text{Binomial}(K, p_{i1} + p_{i2})$. If we want dose j to have a population average DLT rate of π_j, then we select r_j such that $\pi_j = \text{E}[\Pr(N_i > 0)] = \text{E}[1 - (1 - \operatorname{expit}(\eta_{i2}))^K]$, where the expectation is with respect to γ_i. Since we rescale the doses, shifting the intercepts α_1 and α_2 by some constant δ has no effect on the probabilities. Thus, without loss of generality, we fix $\alpha_2 = 0$, and we require $\alpha_1 > 0$ to ensure the probabilities are positive. Once calibrated, we draw $Y_{ik} \sim \text{Multinomial}(1; p_{i0}, p_{i1}, p_{i2})$ for $k = 1, \ldots, K$. Then $N_i = \sum_{k=1}^{K} I(Y_{ik} = 2)$ and $M_i = \sum_{k=1}^{K} I(Y_{ik} \in \{1, 2\})$.

2.4.3 Simulation Scenarios

Using our data-generating simulation model and dose-finding algorithm, we can simulate full trials. In each simulated trial, there are $J = 5$ candidate dose levels, we enroll n patients one at a time, and the target DLT rate is $\tau = 0.3$. To setup our models, we use a skeleton of $(0.1, 0.2, 0.3, 0.4, 0.5)$. We must also specify some model-specific tuning parameters: For Model 1, we use $\sigma_\beta = 1$. For Model 2, $\sigma_\theta = 0.8$, $a_q = 2$, and $b_q = 8$. For Model 3, $\sigma_\theta = 0.8$, $\alpha_0 = 0.8$, and $\sigma_\alpha = 0.7$. For the CRM model, $\sigma_\phi = 1$. These hyperparameter values were chosen so that the prior DLT rate distributions at each dose are close across the four models, as shown in Figure 2.1.

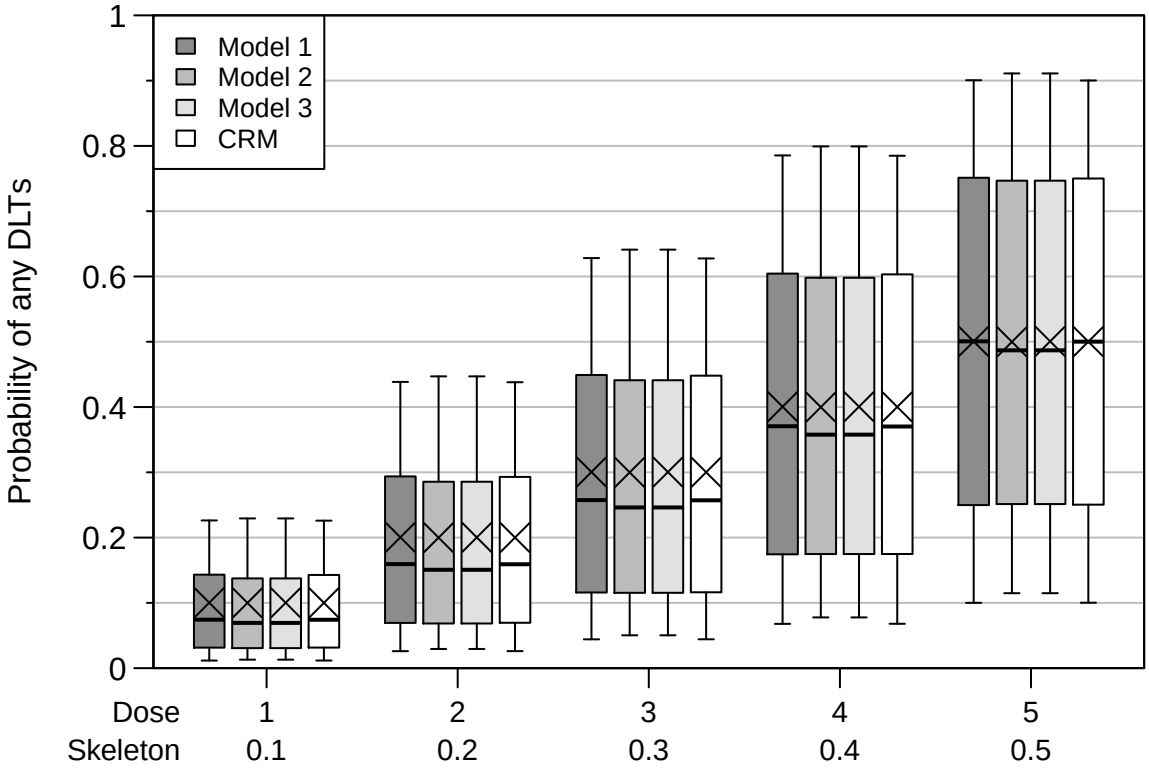

Figure 2.1. Prior distributions of DLT rates at each dose, for all four models. The boxplots show the 10th, 25th, 50th, 75th, and 90th percentiles, and the $\times$ marks the mean. By design, the prior means equal the skeleton.

We consider a variety of scenarios by varying three factors: (1) the sample size n, with either 30 or 60 patients; (2) the true DLT rates for all J doses, thereby determining which dose is the true MTD; and (3) the parameters for our simulation model, α_1 and γ_{sd}. All scenarios use $K = 15$ event types. For the scenarios we call s-A, $s = 1, \ldots, J$, we set $\alpha_1 = 2$ and $\gamma_{\text{sd}} = 0.5$. For scenarios s-B, we set $\alpha_1 = 2.5$ and $\gamma_{\text{sd}} = 0.75$. The relevance of these parameters is that, in the A scenarios, the simulated event counts M_i are typically small, while in the B scenarios M_i is more likely to reach its upper bound of 15 (i.e., there is an event in all 15 event types). Thus the B scenarios present more of a challenge for our Models 1, 2, and 3, which have no concept of an upper bound for the event count. For example, in scenario 3-A, with true DLT rates $(0.1, 0.2, 0.3, 0.4, 0.5)$, the probability of a patient having at least 10 events at the different doses are $(0, 0.00004, 0.0006, 0.0033, 0.012)$. In contrast, for

scenario 3-B, with the same DLT rates, the probabilities of having at least 10 events are $(0.0002, 0.0038, 0.018, 0.049, 0.11)$. Table 2.1 shows the true DLT rates for each scenario. Note that in scenarios s-A and s-B, the true MTD is dose s.

2.4.4 Simulation Results

For each simulation scenario and statistical model, we simulated 10,000 adaptive trials. The results are reported in Tables 2.2 and 2.3, for 30 and 60 patients, respectively, with the following summary statistics: the percentage of trials in which each dose is selected as the MTD; the average number of patients (out of 30 or 60) assigned to each dose during a trial; and a score that measures how far the true DLT rate of the selected MTD is from the target τ. We actually consider two scores, defined as $\text{score}_k = \sum_{j=1}^{J} \hat{s}_j \, \text{loss}_k(\pi_j, \tau)$, where $\hat{s}_j$ is the percentage of trials in which dose j is selected as the MTD, π_j is the true DLT rate at dose j, $\text{loss}_1(\pi, \tau) = |\pi - \tau|$, and $\text{loss}_2(\pi, \tau) = |\pi - \tau| I(\pi \leq \tau) + 2|\pi - \tau| I(\pi > \tau)$. Note that loss_1 is symmetric in π around τ, while loss_2 penalizes $\pi > \tau$ more than $\pi \leq \tau$, i.e., there is a higher penalty for picking dose levels above the target rather than below. For both scores, smaller values are better, and a score of 0 is ideal.

In Table 2.2, we see that, in the A scenarios, our models consistently do at least as well and often better than the CRM, picking the correct MTD more often and treating more patients at the correct MTD. In the B scenarios, the results are more

Table 2.1. True DLT rates for each simulation scenario.

	Dose level				
Scenarios	1	2	3	4	5
1-A and 1-B	**0.30**	0.40	0.50	0.60	0.70
2-A and 2-B	0.20	**0.30**	0.40	0.50	0.60
3-A and 3-B	0.10	0.20	**0.30**	0.40	0.50
4-A and 4-B	0.07	0.14	0.21	**0.30**	0.40
5-A and 5-B	0.06	0.12	0.18	0.24	**0.30**

Table 2.2. Simulation results with 30 patients. Each line summarizes 10,000 simulated adaptive trials of 30 patients with 5 dose levels. The target DLT rate is 0.3. The *italicized* rows show the true DLT rates. The **bold** entries highlight the results at the true MTD.

		Dose Selection Percentage (Number of Patients Treated)						
Scenario	Model	1	2	3	4	5	$Score_1$	$Score_2$
		0.3	*0.4*	*0.5*	*0.6*	*0.7*		
	1	**57.3 (13.4)**	40.0 (13.0)	2.7 (2.9)	0.0 (0.6)	0.0 (0.2)	4.5	9.1
1-A	2	**73.1 (18.9)**	25.8 (9.3)	1.1 (1.6)	0.0 (0.2)	0.0 (0.0)	2.8	5.6
	3	**63.2 (15.6)**	35.6 (12.6)	1.2 (1.7)	0.0 (0.1)	0.0 (0.0)	3.8	7.6
	CRM	**51.5 (11.9)**	43.4 (13.8)	4.9 (3.5)	0.2 (0.7)	0.0 (0.2)	5.4	10.8
	1	**67.2 (15.6)**	31.2 (11.3)	1.6 (2.4)	0.1 (0.5)	0.0 (0.2)	3.5	6.9
1-B	2	**83.8 (22.2)**	15.5 (6.7)	0.6 (1.0)	0.0 (0.1)	0.0 (0.0)	1.7	3.3
	3	**83.7 (21.2)**	16.2 (8.3)	0.1 (0.6)	0.0 (0.0)	0.0 (0.0)	1.6	3.3
	CRM	**50.9 (11.7)**	44.2 (13.9)	4.8 (3.5)	0.2 (0.7)	0.0 (0.2)	5.4	10.9
		0.2	*0.3*	*0.4*	*0.5*	*0.6*		
	1	14.4 (5.1)	**61.0 (15.2)**	22.7 (7.3)	1.9 (1.8)	0.1 (0.6)	4.1	6.8
2-A	2	22.7 (7.7)	**58.6 (15.1)**	17.6 (5.9)	1.1 (1.1)	0.0 (0.2)	4.3	6.2
	3	14.5 (5.5)	**63.9 (16.3)**	20.6 (7.3)	1.0 (0.8)	0.0 (0.0)	3.7	6.0
	CRM	14.3 (4.9)	**57.1 (14.6)**	25.0 (7.7)	3.6 (2.1)	0.1 (0.7)	4.7	7.9
	1	20.7 (6.5)	**61.6 (15.3)**	16.6 (6.2)	1.0 (1.5)	0.0 (0.5)	4.0	5.8
2-B	2	37.6 (11.9)	**51.5 (13.4)**	10.3 (3.9)	0.6 (0.7)	0.0 (0.1)	4.9	6.1
	3	29.1 (9.1)	**63.2 (17.2)**	7.5 (3.6)	0.2 (0.2)	0.0 (0.0)	3.7	4.5
	CRM	13.8 (4.8)	**57.6 (14.8)**	25.3 (7.7)	3.2 (2.1)	0.2 (0.7)	4.6	7.8
		0.1	*0.2*	*0.3*	*0.4*	*0.5*		
	1	0.5 (1.5)	26.4 (9.1)	**53.5 (12.2)**	17.5 (5.3)	2.1 (1.9)	4.9	7.1
3-A	2	0.8 (1.7)	29.3 (9.8)	**53.2 (12.4)**	15.3 (4.9)	1.5 (1.2)	4.9	6.8
	3	0.4 (1.5)	22.4 (8.2)	**59.6 (14.5)**	17.0 (5.4)	0.6 (0.5)	4.1	6.0
	CRM	0.8 (1.6)	26.3 (9.1)	**49.7 (11.6)**	20.1 (5.5)	3.2 (2.2)	5.4	8.1
	1	1.1 (1.8)	33.6 (10.3)	**50.5 (11.5)**	13.6 (4.6)	1.3 (1.7)	5.2	6.8
3-B	2	2.3 (2.9)	40.2 (12.3)	**46.2 (10.7)**	10.4 (3.4)	0.9 (0.8)	5.7	6.9
	3	0.9 (1.9)	39.0 (12.4)	**53.6 (13.3)**	6.4 (2.3)	0.0 (0.1)	4.7	5.4
	CRM	0.8 (1.6)	27.2 (9.2)	**48.6 (11.3)**	20.0 (5.6)	3.3 (2.2)	5.6	8.2
		0.07	*0.14*	*0.21*	*0.3*	*0.4*		
	1	0.0 (1.1)	3.7 (3.8)	33.2 (9.4)	**44.8 (9.5)**	18.2 (6.2)	5.4	7.2
4-A	2	0.0 (1.1)	3.3 (3.4)	33.4 (9.7)	**47.0 (10.4)**	16.2 (5.4)	5.2	6.8
	3	0.0 (1.1)	2.7 (3.1)	29.3 (9.5)	**53.1 (12.1)**	14.9 (4.2)	4.6	6.0
	CRM	0.0 (1.2)	4.9 (4.2)	30.1 (9.0)	**44.1 (9.1)**	20.8 (6.5)	5.6	7.7
	1	0.1 (1.2)	5.9 (4.5)	37.7 (9.9)	**41.8 (9.0)**	14.5 (5.4)	5.8	7.2
4-B	2	0.1 (1.4)	8.3 (5.5)	41.2 (10.9)	**39.4 (8.5)**	11.0 (3.7)	6.2	7.3
	3	0.0 (1.2)	5.3 (4.5)	44.7 (12.9)	**45.3 (10.0)**	4.8 (1.4)	5.3	5.8
	CRM	0.0 (1.2)	5.1 (4.4)	31.6 (9.0)	**42.1 (8.9)**	21.3 (6.5)	5.8	7.9
		0.06	*0.12*	*0.18*	*0.24*	*0.3*		
	1	0.0 (1.1)	1.5 (2.8)	14.1 (6.1)	33.1 (8.1)	**51.2 (11.9)**	4.0	4.0
5-A	2	0.0 (1.1)	0.9 (2.3)	12.0 (5.7)	35.4 (8.9)	**51.7 (12.1)**	3.7	3.7
	3	0.0 (1.1)	0.9 (2.2)	11.6 (6.0)	38.8 (10.4)	**48.7 (10.3)**	3.9	3.9
	CRM	0.0 (1.1)	2.1 (3.2)	15.6 (6.5)	31.8 (7.6)	**50.4 (11.5)**	4.2	4.2
	1	0.0 (1.1)	2.4 (3.0)	16.9 (6.6)	35.4 (8.2)	**45.3 (11.1)**	4.6	4.6
5-B	2	0.0 (1.2)	2.7 (3.5)	20.5 (7.7)	37.8 (8.8)	**38.8 (8.8)**	5.2	5.2
	3	0.0 (1.1)	1.6 (3.0)	21.9 (9.0)	48.5 (11.4)	**27.9 (5.5)**	5.8	5.8
	CRM	0.0 (1.1)	2.0 (3.2)	15.4 (6.4)	32.1 (7.7)	**50.5 (11.6)**	4.1	4.1

Table 2.3. Simulation results with 60 patients. Each line summarizes 10,000 simulated adaptive trials of 60 patients with 5 dose levels. The target DLT rate is 0.3. The *italicized* rows show the true DLT rates. The **bold** entries highlight the results at the true MTD.

		Dose Selection Percentage (Number of Patients Treated)						
Scenario	Model	1	2	3	4	5	$Score_1$	$Score_2$
		0.3	*0.4*	*0.5*	*0.6*	*0.7*		
	1	**68.3 (32.3)**	31.5 (23.8)	0.2 (3.2)	0.0 (0.6)	0.0 (0.2)	3.2	6.4
1-A	2	**80.3 (41.9)**	19.5 (16.0)	0.2 (1.9)	0.0 (0.2)	0.0 (0.0)	2.0	4.0
	3	**72.7 (36.0)**	27.1 (22.0)	0.1 (1.9)	0.0 (0.1)	0.0 (0.0)	2.7	5.5
	CRM	**61.4 (28.9)**	37.7 (26.0)	0.9 (4.2)	0.0 (0.7)	0.0 (0.2)	4.0	7.9
	1	**79.4 (37.6)**	20.5 (19.2)	0.1 (2.5)	0.0 (0.5)	0.0 (0.2)	2.1	4.2
1-B	2	**90.1 (48.2)**	9.8 (10.6)	0.0 (1.1)	0.0 (0.1)	0.0 (0.0)	1.0	2.0
	3	**91.1 (47.5)**	8.9 (11.9)	0.0 (0.6)	0.0 (0.0)	0.0 (0.0)	0.9	1.8
	CRM	**61.0 (28.7)**	37.9 (26.2)	1.1 (4.3)	0.0 (0.7)	0.0 (0.2)	4.0	8.0
		0.2	*0.3*	*0.4*	*0.5*	*0.6*		
	1	12.6 (9.3)	**72.4 (35.6)**	14.9 (12.7)	0.1 (1.9)	0.0 (0.5)	2.8	4.3
2-A	2	19.1 (13.8)	**68.8 (34.6)**	12.0 (10.2)	0.2 (1.3)	0.0 (0.2)	3.1	4.4
	3	12.4 (9.2)	**73.4 (37.2)**	14.1 (12.6)	0.1 (0.9)	0.0 (0.0)	2.7	4.1
	CRM	11.8 (8.7)	**67.1 (33.6)**	20.5 (14.5)	0.6 (2.5)	0.0 (0.7)	3.4	5.5
	1	19.5 (12.7)	**72.0 (35.4)**	8.5 (9.9)	0.0 (1.6)	0.0 (0.5)	2.8	3.7
2-B	2	32.7 (22.3)	**61.6 (30.4)**	5.6 (6.4)	0.1 (0.8)	0.0 (0.1)	3.9	4.4
	3	27.1 (17.9)	**68.7 (36.5)**	4.2 (5.5)	0.0 (0.2)	0.0 (0.0)	3.1	3.6
	CRM	11.8 (8.8)	**67.4 (33.6)**	20.2 (14.4)	0.6 (2.5)	0.0 (0.7)	3.3	5.5
		0.1	*0.2*	*0.3*	*0.4*	*0.5*		
	1	0.0 (1.6)	20.8 (16.1)	**66.6 (30.1)**	12.3 (10.0)	0.2 (2.2)	3.4	4.6
3-A	2	0.0 (1.7)	22.2 (17.3)	**66.8 (30.6)**	10.8 (9.0)	0.1 (1.4)	3.3	4.4
	3	0.1 (1.5)	18.1 (14.3)	**70.3 (33.9)**	11.4 (9.7)	0.1 (0.6)	3.0	4.2
	CRM	0.2 (1.7)	20.1 (15.9)	**62.4 (28.4)**	16.8 (11.3)	0.6 (2.7)	3.8	5.6
	1	0.2 (2.0)	29.6 (19.6)	**62.9 (28.8)**	7.3 (7.7)	0.0 (1.9)	3.7	4.5
3-B	2	0.4 (3.2)	36.2 (23.9)	**57.0 (26.1)**	6.3 (5.9)	0.1 (0.9)	4.3	5.0
	3	0.1 (2.0)	31.7 (22.8)	**64.1 (31.4)**	4.1 (3.8)	0.0 (0.1)	3.6	4.0
	CRM	0.1 (1.7)	19.8 (15.8)	**63.2 (28.9)**	16.3 (11.0)	0.6 (2.7)	3.8	5.5
		0.07	*0.14*	*0.21*	*0.3*	*0.4*		
	1	0.0 (1.1)	0.8 (4.5)	28.1 (18.6)	**59.9 (25.4)**	11.2 (10.4)	3.8	4.9
4-A	2	0.0 (1.1)	0.8 (4.0)	26.9 (18.3)	**61.6 (27.0)**	10.7 (9.6)	3.6	4.7
	3	0.0 (1.1)	0.7 (3.5)	24.7 (17.9)	**64.6 (29.8)**	10.0 (7.8)	3.3	4.3
	CRM	0.0 (1.2)	1.3 (5.1)	25.6 (17.5)	**57.4 (24.2)**	15.8 (12.0)	4.1	5.7
	1	0.0 (1.2)	2.0 (5.7)	37.5 (21.3)	**53.8 (23.3)**	6.8 (8.4)	4.4	5.1
4-B	2	0.0 (1.4)	2.2 (6.8)	39.2 (23.0)	**52.3 (22.6)**	6.3 (6.2)	4.5	5.1
	3	0.0 (1.2)	1.4 (5.4)	40.9 (26.0)	**54.3 (24.7)**	3.4 (2.7)	4.2	4.6
	CRM	0.0 (1.2)	1.2 (5.1)	26.4 (17.8)	**56.3 (23.9)**	16.1 (12.0)	4.2	5.8
		0.06	*0.12*	*0.18*	*0.24*	*0.3*		
	1	0.0 (1.1)	0.2 (3.0)	7.1 (9.6)	36.6 (18.6)	**56.2 (27.8)**	3.1	3.1
5-A	2	0.0 (1.1)	0.1 (2.3)	5.9 (8.5)	37.2 (19.7)	**56.8 (28.5)**	3.0	3.0
	3	0.0 (1.1)	0.1 (2.3)	5.7 (8.5)	39.1 (21.9)	**55.1 (26.2)**	3.0	3.0
	CRM	0.0 (1.1)	0.3 (3.5)	7.7 (9.7)	34.5 (17.9)	**57.5 (27.8)**	3.0	3.0
	1	0.0 (1.1)	0.2 (3.4)	10.6 (11.1)	42.1 (19.9)	**47.0 (24.5)**	3.8	3.8
5-B	2	0.0 (1.2)	0.4 (4.1)	12.5 (12.6)	44.4 (21.0)	**42.8 (21.1)**	4.2	4.2
	3	0.0 (1.1)	0.3 (3.3)	11.7 (13.9)	52.6 (26.9)	**35.4 (14.8)**	4.6	4.6
	CRM	0.0 (1.1)	0.2 (3.4)	7.3 (9.7)	34.1 (17.7)	**58.4 (28.1)**	3.0	3.0

mixed. For scenarios 2-B through 4-B, our models typically do around as well as the CRM in selecting the correct MTD. However, when not picking the correct MTD our models have a clear tendency to select doses lower than the MTD, while the CRM appears more likely to pick higher doses. This is reflected in the score_2 column, where our models consistently have lower scores. This trend is more pronounced in scenarios 1-B and 5-B, where the true MTD is the lowest or highest dose, respectively. In 1-B, our models do much better than the CRM, but this is likely because our models prefer lower doses. Similarly, in 5-B the CRM does considerably better, but this is likely because the CRM prefers higher doses.

It is difficult to identify a clear winner among the models across all scenarios. If we ignore scenario 5-B, then Model 3, which incorporates LLTs and allows the probability of DLT versus LLT to vary with dose, is the preferred model. This validates the idea that counting events and using LLTs can improve estimation of the MTD. However, Model 3 is clearly not preferred in scenario 5-B, in which the highest dose is the MTD and patients commonly experience many events, sometimes reaching the maximum of 15 events. We consider this scenario somewhat unlikely, however, as a well-designed trial will typically have the true MTD among the middle dose levels.

2.4.5 Example Trial

We illustrate the behavior of Model 3 by presenting a full simulated trial. We simulated the trial data using scenario 4-A described previously, where the true DLT rates are $(0.07, 0.14, 0.21, 0.3, 0.4)$. The target DLT rate is 0.3, so that dose 4 is the true MTD. Our skeleton is (0.1, 0.2, 0.3, 0.4, 0.5), implying that *a priori* the MTD is believed to be dose 3. The results from simulating 30 patients are shown in Table 2.4. First we note that, looking at the last row of the table, the trial correctly concludes that dose 4 is the MTD. Indeed, although there is still some movement, the trial has mostly fixed on dose 4 by patient 12. Based on the results in Table 2.2, Model 3 will

pick the correct dose in scenario 4-A 53.1% of the time, so this example trial is in line with expectations. We also note that only one patient has multiple DLTs in this trial, although 18 patients have more than one total event, so our model does have data to work with. At the end of the trial, beyond selecting the correct MTD, we have also estimated the DLT rates well: $(0.07, 0.12, 0.18, 0.27, 0.38)$.

2.4.6 Sensitivity Analysis

We conduct a small sensitivity analysis, shown in Table 2.5, investigating how model performance varies with changing model parameters and with the number of event types K. In particular, starting from the 2-A scenario, we consider the following three variations: "Larger variance," in which we increase the prior variances of model parameters ($\sigma_\beta = 1.4$ for Model 1, $\sigma_\theta = 1.15$, $a_q = 1$, and $b_q = 4$ for Model 2, $\sigma_\theta = 1.3$, $\alpha_0 = 2$, and $\sigma_\alpha = 1.1$ for Model 3, and $\sigma_\phi = 1.4$ for the CRM); "Smaller variance," in which we decrease the prior variances of model parameters ($\sigma_\beta = 0.7$ for Model 1, $\sigma_\theta = 0.56$, $a_q = 4$, and $b_q = 16$ for Model 2, $\sigma_\theta = 0.65$, $\alpha_0 = 3$, and $\sigma_\alpha = 0.8$ for Model 3, and $\sigma_\phi = 0.7$ for the CRM); and "$K = 7$ event types," in which we simply reduce the number of events types from 15 to 7. We simulated 10,000 adaptive trials of 30 patients with each model, and for each described variation. We observe that all models are sensitive to the prior variances of model parameters, with the "smaller variance" models outperforming the "larger variance" models by about 10% in terms of selecting the correct MTD. In contrast, reducing the number of event types made minimal difference compared to the 2-A results in Table 2.2. These results suggest that the models, including the CRM, should be carefully calibrated given a variety of plausible DLT rates, while the precise number of event types is less important.

Table 2.4. Example of a single simulated trial of 30 patients and 5 dose levels, using Model 3 and scenario 4-A. The true DLT rates are $(0.07, 0.14, 0.21, 0.3, 0.4)$ and the target is 0.3, so dose 4 is the true MTD. Our initial guesses for the DLT rates are $(0.1, 0.2, 0.3, 0.4, 0.5)$, implying that initially we suspect dose 3 is the MTD. However, our trials always start the first patient on dose 1.

	Dose	Total		Estimated DLT Rates					
Patient	Given	Events	DLTs	Dose 1	2	3	4	5	$\widehat{\text{MTD}}$
1	1	1	0	0.10	0.21	0.32	0.44	0.55	2
2	2	0	0	0.07	0.16	0.25	0.36	0.47	3
3	3	4	0	0.08	0.16	0.25	0.37	0.51	3
4	3	3	0	0.08	0.16	0.25	0.37	0.50	3
5	3	5	1	0.11	0.20	0.31	0.45	0.59	3
6	3	1	0	0.09	0.18	0.28	0.41	0.55	3
7	3	0	0	0.08	0.16	0.25	0.37	0.51	3
8	3	0	0	0.07	0.14	0.23	0.34	0.47	4
9	4	6	1	0.09	0.16	0.26	0.37	0.52	3
10	3	2	0	0.08	0.16	0.25	0.36	0.50	3
11	3	1	0	0.08	0.15	0.23	0.34	0.48	4
12	4	3	0	0.08	0.14	0.22	0.32	0.45	4
13	4	1	0	0.07	0.13	0.21	0.30	0.42	4
14	4	0	0	0.07	0.12	0.19	0.28	0.40	4
15	4	2	1	0.07	0.13	0.21	0.30	0.43	4
16	4	2	0	0.07	0.13	0.20	0.29	0.41	4
17	4	1	0	0.07	0.12	0.19	0.28	0.39	4
18	4	5	1	0.07	0.13	0.20	0.30	0.43	4
19	4	1	0	0.07	0.12	0.19	0.29	0.41	4
20	4	0	0	0.06	0.12	0.19	0.27	0.39	4
21	4	4	0	0.07	0.12	0.18	0.27	0.38	4
22	4	2	0	0.06	0.11	0.18	0.26	0.37	4
23	4	6	0	0.07	0.12	0.18	0.26	0.36	4
24	4	2	0	0.07	0.11	0.17	0.25	0.35	4
25	4	2	1	0.07	0.12	0.18	0.26	0.37	4
26	4	0	0	0.06	0.11	0.17	0.25	0.36	4
27	4	2	0	0.06	0.11	0.17	0.25	0.35	5
28	5	6	2	0.07	0.13	0.19	0.28	0.40	4
29	4	2	0	0.07	0.12	0.19	0.28	0.39	4
30	4	2	0	0.07	0.12	0.18	0.27	0.38	4

Table 2.5. Sensitivity of models to tuning parameters and number of event types. Each line summarizes 10,000 simulated adaptive trials of 30 patients with 5 dose levels. The target DLT rate is 0.3 and the true DLT rates are $(0.2, 0.3, 0.4, 0.5, 0.6)$, so dose 2 is the true MTD. Results should be compared with the scenario 2-A results in Table 2.2.

		Dose Selection Percentage (Number of Patients Treated)						
Scenario	Model	1	2	3	4	5	Score_1	Score_2
		0.2	*0.3*	*0.4*	*0.5*	*0.6*		
	1	18.8 (7.0)	**56.2 (13.8)**	22.7 (6.8)	2.2 (1.8)	0.0 (0.6)	4.6	7.3
Larger	2	24.6 (8.9)	**53.7 (13.1)**	19.7 (6.0)	1.9 (1.5)	0.0 (0.4)	4.8	7.2
Variance	3	22.9 (8.1)	**59.9 (16.0)**	16.7 (5.5)	0.6 (0.4)	0.0 (0.0)	4.1	5.9
	CRM	20.3 (7.5)	**51.9 (12.9)**	24.2 (6.8)	3.5 (2.1)	0.2 (0.7)	5.2	8.4
	1	7.3 (2.6)	**67.3 (16.8)**	23.8 (8.5)	1.6 (1.8)	0.0 (0.3)	3.4	6.1
Smaller	2	17.9 (5.5)	**66.1 (17.9)**	15.3 (5.9)	0.6 (0.6)	0.0 (0.0)	3.5	5.1
Variance	3	2.6 (1.7)	**67.7 (16.7)**	29.2 (11.1)	0.4 (0.5)	0.0 (0.0)	3.3	6.3
	CRM	5.9 (2.2)	**62.7 (15.9)**	28.6 (9.5)	2.7 (2.1)	0.1 (0.3)	4.0	7.4
	1	12.8 (4.7)	**61.0 (15.3)**	23.7 (7.6)	2.4 (1.8)	0.1 (0.5)	4.2	7.0
$K = 7$	2	17.8 (6.4)	**59.2 (15.0)**	21.2 (6.9)	1.7 (1.5)	0.1 (0.3)	4.3	6.7
Event Types	3	12.4 (4.9)	**60.1 (15.2)**	25.7 (8.5)	1.8 (1.4)	0.0 (0.1)	4.2	7.1
	CRM	14.4 (4.9)	**57.5 (14.8)**	24.9 (7.6)	3.0 (2.0)	0.1 (0.6)	4.6	7.7

2.4.7 Model Properties

To better understand the performance of our models, we investigate the bias in estimating DLT rates with each model. For simplicity, we simulate non-adaptive trials of 10 or 20 patients all receiving the same dose level, and we consider all combinations of true DLT rate and prior guess from 0 to 0.7 in increments of 0.025. Patient responses are simulated using the A and B scenarios described previously.

Figure 2.2 shows the bias for both the CRM model and our Models 1, 2, and 3. In each plot, along the diagonal from bottom left to top right, the true DLT rate and prior guess agree, so we would expect minimal bias near this region. Above the diagonal, the prior guess exceeds the truth, so we would expect positive bias, and vice versa below the diagonal. In Table 2.2, we note that our models tend to choose doses below the true MTD more often than the CRM. As a plausible explanation, we

see in Figure 2.2 that our models tend to be more positively biased than the CRM (particularly Models 2 and 3). That is, our models overestimate the DLT rates more than the CRM, and thus the dose selection algorithm picks lower doses. However, large differences between the truth and the skeleton tend to lead to greater bias for the CRM than for our models.

We also note that our Models 2 and 3 tend to select the true MTD less often in the B scenarios than in the A scenarios. This is likely due to a difference between the data-generating model and Models 2 and 3. Specifically, the data-generating model creates LLT and DLT event counts for a patient, and the sum of these two counts cannot exceed 15, thereby making the two counts negatively correlated. However, Models 2 and 3 implicitly force this correlation to be zero. In the A scenarios, both counts are usually low, so they do not influence each other much, and there is minimal conflict with our models. The B scenarios, in contrast, generate higher event counts for a patient, leading to stronger correlations and greater conflict with our models. If we increase the number of event types to $K = 60$ in our generative model, so that the upper bound is almost never realized in scenarios A and B, then the differences in dose selection percentages between scenarios A and B decreases, although using 60 event types would not be realistic.

Additionally, we note that the DLT count N_i either explicitly, as in Model 1, or implicitly, as in Models 2 and 3 after marginalizing over M_i, follows a Poisson distribution in our proposed models, while in the data-generating model N_i is binomial. This discrepancy impacts the performance of our models, yet it has no effect on the CRM, which does not attempt to model event counts. For this reason, the CRM should perform exactly the same in scenarios A and B, within Monte Carlo simulation error.

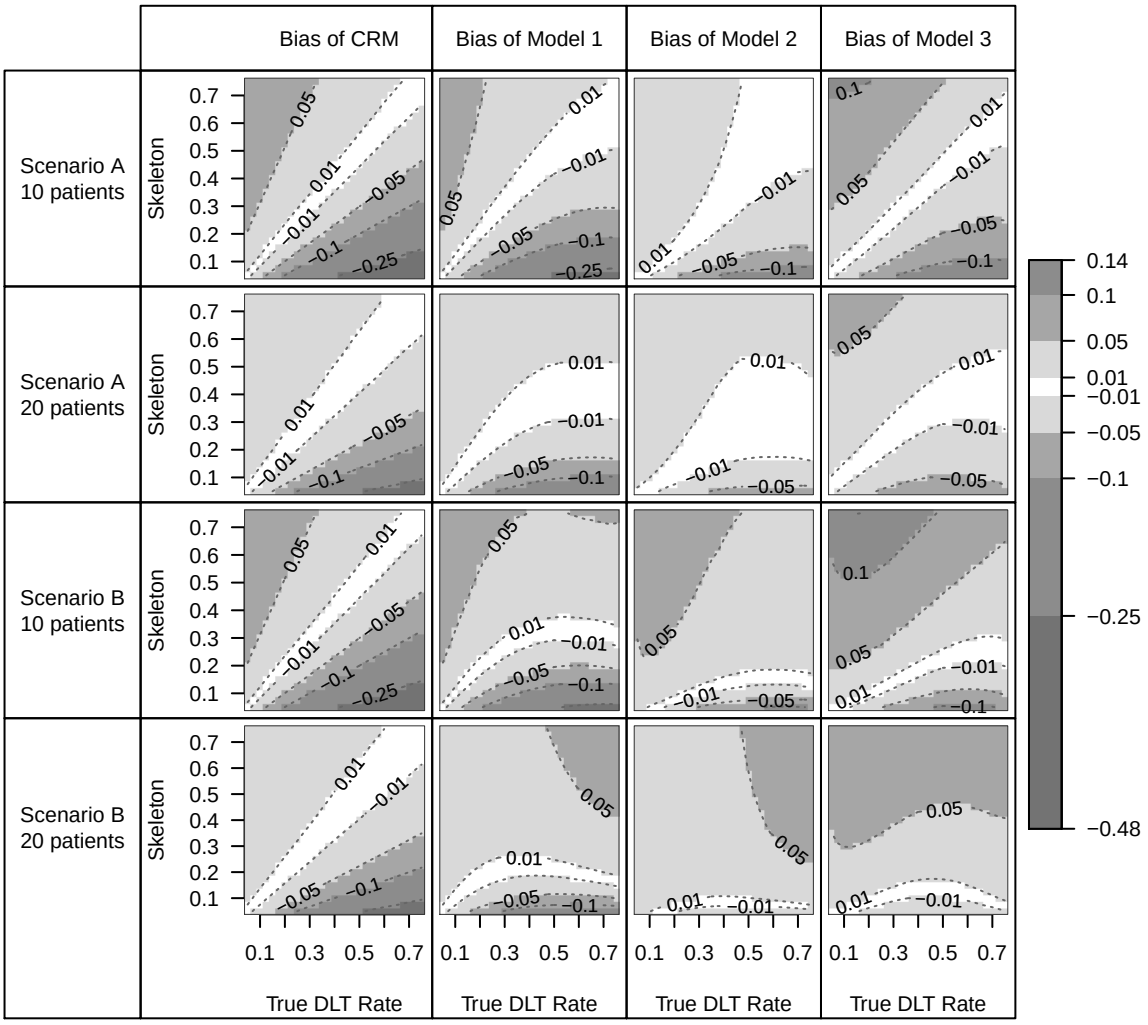

Figure 2.2. Bias in estimating DLT rates for CRM model and Models 1, 2, and 3. Each pixel in each plot is based on 10,000 simulated trials under scenarios A and B of either 10 or 20 patients. In a trial, all patients receive the same dose level, i.e., the trial is not adaptive, and the true DLT rate is the same. The purpose of this figure is to show, for every combination of true DLT rate and prior guess (i.e. skeleton), how well the models do in estimating the DLT rate.

2.5 Discussion

Models 1, 2, and 3 all make use of phase I trial data that are routinely collected but rarely analyzed or incorporated into the dose selection algorithm. Model 1 makes fewer assumptions than Models 2 and 3, placing a Poisson distribution on the number of DLTs, but disregarding the number of LLTs. Models 2 and 3 jointly model a patient's DLT and LLT counts.

All three models perform comparably to, and sometimes better than, the CRM across a variety of scenarios in our simulation study. Nonetheless, the simulations did suggest that including total number of toxicities and DLTs might lead to conservative dose decisions when the true MTD is at the highest dose. However, we do observe better operating characteristics for some of our models when the MTD is one of the lower doses.

The data-generating model in our simulations was designed both to be biologically plausible and to disagree with our three proposed models, thereby creating fair test scenarios. Although we believe our models generally compare favorably with the CRM, in scenario 5-B, Models 2 and 3 perform poorly relative to both the CRM and to Model 1. The import of this differential depends on how plausible one believes scenario 5-B to be. If the true MTD is rarely the highest dose under consideration, then Model 3 clearly bests both the CRM and our Models 1 and 2. This suggests that we are indeed gaining useful information and efficiency by modeling events counts and incorporating LLTs.

In a real trial, patients may drop out early, or it may be desirable to enroll a new patient before the previous patients have finished follow-up. Our models can be extended to handle partial follow-up by scaling the expected number of toxicities in accordance with follow-up, in a manner similar to the time-to-event CRM [5]. Further research is needed to recommend appropriate scaling weights.

CHAPTER 3

Phase I Clinical Trial Method for a Multi-Cycle Cytotoxic Agent

3.1 Introduction

In most phase I clinical trial methods, each patient provides a toxicity outcome at a single timepoint. The outcome could be binary, as in the continual reassessment method (CRM), or more complex, reflecting various degrees of severity. There is only a small literature, however, on methods for longitudinal outcomes, obtained from treatments administered multiple times. In each cycle of administration, a patient may experience toxicity, and furthermore the dose may change between cycles. If a patient has a dose-limiting toxicity (DLT), the treatment schedule will likely stop and no further cycles will be administered. However, if the patient has a low-level toxicity (LLT), the treatment schedule might continue, although perhaps at a lower dose.

In current practice, toxicities recorded after the first cycle are often ignored, with estimates of DLT rates based only on the first cycle. The primary reason for this practice is expediency: patients can be evaluated more quickly, and thus additional patients can be enrolled sooner, leading to a shorter trial. However, a patient may not experience a DLT until after the first cycle, and if ignored, DLT rates will likely be underestimated. In a retrospective study of phase I trials of multi-cycles treatments, Postel-Vinay et al. [26] found that among 445 patients, 57% of severe toxicities oc-

curred after the first cycle, and 50% of patients had their worst toxicity after the first cycle.

One simple way to use toxicity outcomes from all cycles is to record whether a patient had a DLT in any cycle, and then use the CRM. However, if the patient's dose changed over the cycles, it is unclear if the DLT should be associated with the original dose assignment or with the dose level given just before the DLT occurred. To use the CRM, we must choose between these two possibilities. However, if we believe the DLT is attributable to the cumulative effect of all preceding treatment cycles, as Simon et al. argued [44], then the CRM can no longer be used, and newer methods are needed.

To address this setting, Legedza and Ibrahim [15] proposed using logistic regression to model the probability of DLT, accounting for cumulative dose in later cycles. The mechanism of accounting for cumulative dose was inspired by the pharmacokinetics of drug absorption and clearance. Their model has the form $\text{logit}\{\Pr(\text{DLT}_{ic})\} = \epsilon + \beta \log(d_{ic} + e^{-\lambda} D_{ic})$, where DLT_{ic} is the observed binary indicator of DLT for patient i in cycle c, d_{ic} is the dose for that patient and cycle, D_{ic} is the patient's cumulative dose prior to cycle c, ϵ is an intercept, β is a dose-response parameter, and λ is the drug clearance rate. Only ϵ and β need to be estimated; λ is assumed known. The authors also considered removing ϵ from the model, in which case only β must be estimated. We note, however, that although the authors refer to ϵ as a random effect, it is not subject-specific, and thus even when included the model does not account for within-subject correlation between repeated outcomes. However, the model does allow a patient's dose to vary across cycles.

Doussau et al. [16] in 2013 proposed using a proportional odds mixed-effects model, with three outcome levels in each cycle: no toxicity, LLT, or DLT. Their model has the form $\text{logit}\{\Pr(Y_{ic} \leq k)\} = \alpha_k - \beta_1 d_i - \beta_2 \tau_c - u_i$ for $k = 0, 1$, where Y_{ic} is an ordinal, three-level response variable, with 0 meaning no toxicity, 1 meaning LLT,

and 2 meaning DLT for patient i in cycle c; d_i is the dose for patient i, assumed constant throughout the trial; τ_c is the start time of cycle c; α_0 and α_1 are response level-specific intercepts; β_1 is a dose-response parameter; β_2 is a time-effect parameter; and $u_i \sim \mathcal{N}(0, \sigma_0^2)$ is a random subject-specific intercept, to account for the correlation of a patient's repeated measurements. For parsimony, the authors proposed forcing $\beta_2 = 0$ during the adaptive dose-finding phase of the trial, and then estimating it at trial's end, once all patients have been evaluated. While this model captures within-subject correlation, it does not allow dose to change between cycles, and it only indirectly accounts for a cumulative dose effect via the time effect.

Zhang and Braun [17] in 2013 developed a cure rate model, which estimates the time to DLT from multiple administrations, assuming a fraction of patients are "cured," i.e., will never experience DLT. The hazard of DLT at time t is given by $g(t) = \sum_{c=1}^{m} \theta_{ic} f(t - \tau_c)$, where m is the number of cycles, θ_{ic} is the cure parameter for a patient given dose d_{ic}, and f is the Weibull density. The hazard $g(t)$ from multiple administrations is thus formulated as the sum of hazards from individual dosings. Note that the toxic effect of a dose given at time τ_c eventually wears off according to the shape of $f(t - \tau_c)$. The probability of DLT by the time of the final patient visit τ_{m+1} is thus $1 - \exp\{-\int_0^{\tau_{m+1}} g(t)\, dt\}$. The authors proposed using their model to simultaneously optimize the dose and the schedule of doses. This model allows dose to vary across cycles, and it accounts for cumulative dose by summing individual hazards. However, it does not account for within-subject correlation across cycles, and furthermore it does not leverage the potentially informative LLTs: the hazard of DLT may be higher for a patient with previous LLTs than for one without.

In 2016, Fernandes et al. [18] developed a two-state Markov model, with states DLT or no DLT in each cycle, the latter denoted $\overline{\text{DLT}}$. The model specifies the conditional probability of DLT in cycle c, given that a patient had no previous DLTs, with the form $\Pr(\text{DLT}_{ic} | \overline{\text{DLT}}_{i,c-1}, \ldots, \overline{\text{DLT}}_{i,1}) = 1 - \exp\{-\alpha(d_{ic} - \rho d_{ic}^{\ddagger})^{+} - \beta D_{ic} d_{ic}\}$,

where d_{ic} and D_{ic} are the current and prior cumulative dose for patient i in cycle c; $d_{ic}^{\ddagger} = \max\{d_{i,1}, \ldots, d_{i,c-1}\}$ is the highest dose patient i received prior to cycle c; the function $(x)^+$ equals x if $x > 0$, and is 0 otherwise; and the dose-response parameters to be estimated are α, ρ, and β. Since DLT is considered an absorbing state, there is no need to model the probability of transitioning from DLT to $\overline{\text{DLT}}$. This model accounts for within-subject correlation by explicitly modeling state transition probabilities across cycles; it allows dose to vary across cycles; and it accounts for cumulative dose. However, the model does not make use of LLTs, although a sequel explores including LLTs in a three-state model [45].

Most recently, Yin et al. [46] proposed using the quasi-continuous total toxicity profile (TTP) of Ezzalfani et al. [35] as a longitudinal outcome. TTP is a weighted sum of the grades for all of a patient's toxicities, with weights elicited from experts, and higher values indicating increased toxicity. TTP is normalized to lie between 0 and 1, producing nTTP, which is used in the linear mixed effect model $\text{nTTP}_{ic} = \beta_0 + \beta_1 d_i + \beta_2 \tau_c + u_i + \varepsilon_{ic}$. As in Doussau et al. [16], this model accounts for within-subject correlation via the random effect u_i, and it incorporates LLTs into the calculation of TTP, yet it does not allow dose to change between cycles. Additionally, the model cannot estimate DLT rates, and instead requires the MTD to be redefined in terms of a target nTTP value.

To address the limitations of the preceding models, we develop a model that accounts for within-patient correlation, allows dose to change, incorporates cumulative dose, and includes LLTs to improve the estimation of DLT rates. We have two goals in mind for our model: (1) the model should accurately estimate the maximum tolerated dose (MTD) at the end of a trial, and (2) it should help us decide, for a patient in the middle of their treatment plan, what dose to recommend for their next cycle, given their previously observed outcomes.

3.2 Latent Toxicity Process

We propose using a latent stochastic process to model toxicity over time. Our thinking is that the observed toxicity outcome in each cycle, recorded as either "no toxicity," LLT, or DLT, reflects an underlying unobserved continuous process. Conceptually, the continuous process may represent the amount of drug in a patient's body, or the amount of damage done to the body, with the amount rising after a new administration, and falling as the dose is cleared from the body. As the process rises and falls, it may pass thresholds corresponding to the observed, discrete measures of toxicity.

We define the model as follows. Let $Z(t) = \mu(t) + X(t)$ be the stochastic latent toxicity process (LTP) at time t, which is the sum of a deterministic function, $\mu(t)$, and another stochastic process, $X(t)$. The deterministic part changes over time in response to dose.

With m planned administrations, let τ_c for $c = 1, 2, \ldots, m$ be the patient treatment times after enrollment, with $\tau_1 = 0$. The dose given at time τ_c is d_c; we explain how to choose the numeric value of d_c in Section 3.5. The final patient visit occurs at time τ_{m+1}, at which no dose is given, so there is no corresponding d_{m+1}. Then we define

$$\mu(t) = \sum_{c:\tau_c \le t} d_c h(t - \tau_c) \quad \text{where} \quad h(t) = \left(\frac{t}{\alpha}\right)^{\lambda} \exp\left\{-\lambda\left(\frac{t}{\alpha} - 1\right)\right\} I(t > 0). \quad (3.1)$$

The sum above is over all dose cycles given up to time t. The function $h(t)$ controls the shape of the dose-toxicity response from a single administration. It starts at 0 when $t = 0$, reaches its peak when $t = \alpha$, and then falls back to 0 as $t \to \infty$, assuming $\lambda > 0$. The parameter λ controls how concentrated the dose effect is around the peak time α; higher values of λ imply a shorter timeframe in which toxicities are likely to occur. The motivation for our definition of $\mu(t)$ comes from pharmacokinetics, in which a drug is absorbed into the body, reaches a maximum concentration in the blood, and is eventually cleared out by the kidneys. Summing over all previously

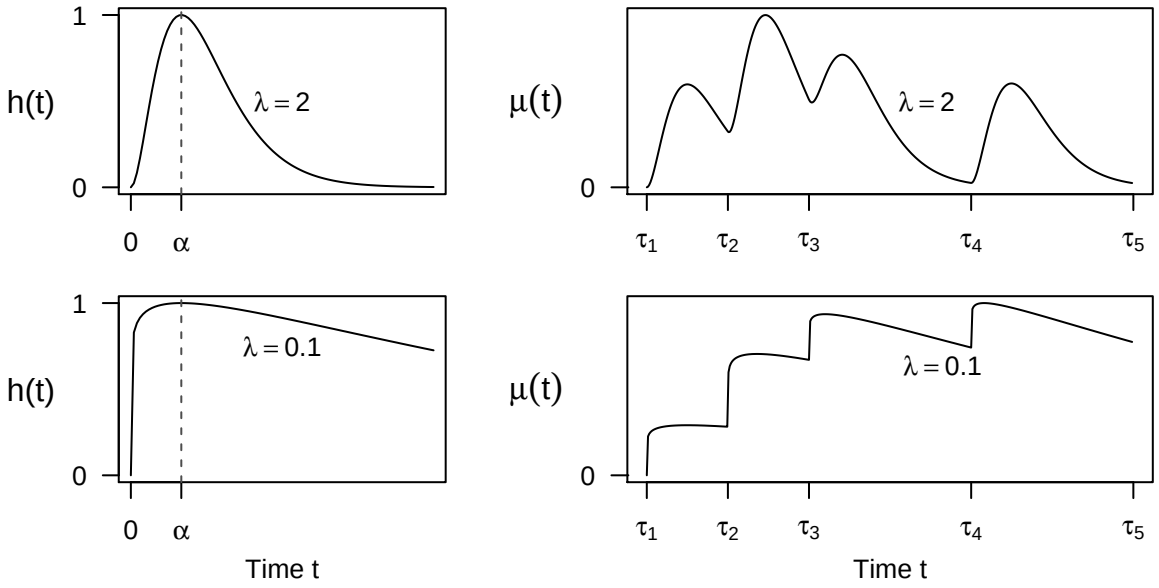

Figure 3.1. Deterministic dose-toxicity curves: $h(t)$ shows the response to a single administration, and $\mu(t)$ the response to multiple administrations.

administered doses allows us to account for their cumulative effect, similar to the summing of hazards in Zhang and Braun [17]. A depiction of $h(t)$ and $\mu(t)$, with $\lambda = 2$ and $\lambda = 0.1$, is in Figure 3.1. Note that with smaller λ, the process has a slower return to zero, leading to a greater cumulative dose effect.

For the stochastic process $X(t)$, which does not depend on dose, we use a mean-zero Ornstein-Uhlenbeck (OU) process, defined by the stochastic differential equation $dX(t) = -\beta X(t)dt + \sigma dW(t)$, with boundary condition $X(0) = 0$ and where $W(t)$ is a Weiner process [47]. This process behaves similarly to Brownian motion, except that it has a tendency to return to zero. Thus, adding $\mu(t)$ and $X(t)$ produces a process $Z(t)$ that broadly follows the trajectory implied by $\mu(t)$, but has the flexibility to deviate from and return to $\mu(t)$. For any two times $s < t$, the OU differential equation implies

$$X(t)|X(s) \sim \mathcal{N}\left(X(s)e^{-\beta(t-s)},\ \frac{\sigma^2}{2\beta}\left(1 - e^{-2\beta(t-s)}\right)\right),$$

and combining this with Equation (3.1) gives

$$Z(t)|Z(s) \sim \mathcal{N}\left(\mu(t) + [Z(s) - \mu(s)]e^{-\beta(t-s)},\ \frac{\sigma^2}{2\beta}\left(1 - e^{-2\beta(t-s)}\right)\right). \tag{3.2}$$

Using Equation (3.2), and noting that $Z(0) = \mu(0) + X(0) = 0$, we can simulate the full process $Z(t)$ for a given set of dose levels and parameter values $\lambda, \alpha, \beta, \sigma$.

In a real trial, we do not observe this latent toxicity process. Instead, at the end of each dose cycle (i.e., $\tau_2, \ldots, \tau_{m+1}$), we observe a discretized measure which reflects how high the latent process reached in the preceding cycle. That is, within each cycle $c \in \{1, 2, \ldots, m\}$ the latent process has a maximum, $Z_c^\vee = \max\{Z(t) : t \in (\tau_c, \tau_{c+1}]\}$, and we relate this to the observed discrete variable Y_c by

$$Y_c = \begin{cases} 0 & \text{if } Z_c^\vee < \gamma_1 \\ 1 & \text{if } Z_c^\vee \in [\gamma_1, \gamma_2) \\ 2 & \text{if } Z_c^\vee \geq \gamma_2 \end{cases} \tag{3.3}$$

where 0, 1, and 2 refer to no toxicity, LLT, and DLT, respectively, and γ_1 and γ_2 are threshold parameters.

The original toxicity data are typically graded by clinicians on a discrete scale from 0 to 5 in order of increasing toxicity [34]. For our model, we require these six values to be mapped into the three used in Equation (3.3); we leave the particular mapping up to clinicians. Typically, $Y_c = 2$ if toxicities are graded 3, 4, or 5. If a patient has multiple toxicities within a cycle, Y_c is coded based on the highest grade. Ultimately, we are interested in whether a DLT ever occurs across the cycles, but using intermediate toxicity levels (e.g., LLTs) provides more information to pinpoint a patient's latent process over time. In principle our model could keep the six original levels separate, providing more precise information on the location of the process, but each additional level requires an additional threshold parameter in the definition of Y_c. As a compromise, we use three levels, with two threshold parameters.

The parameters in this model are $\lambda, \alpha, \beta, \sigma, \gamma_1, \gamma_2$, all of which must be positive, and with the requirement that $\gamma_2 > \gamma_1$. In Figure 3.2, we depict two simulated processes $Z(t)$ with the same deterministic component $\mu(t)$. Note that one patient has DLTs in cycles 2 and 3, while the other patient has LLTs but not DLTs. In

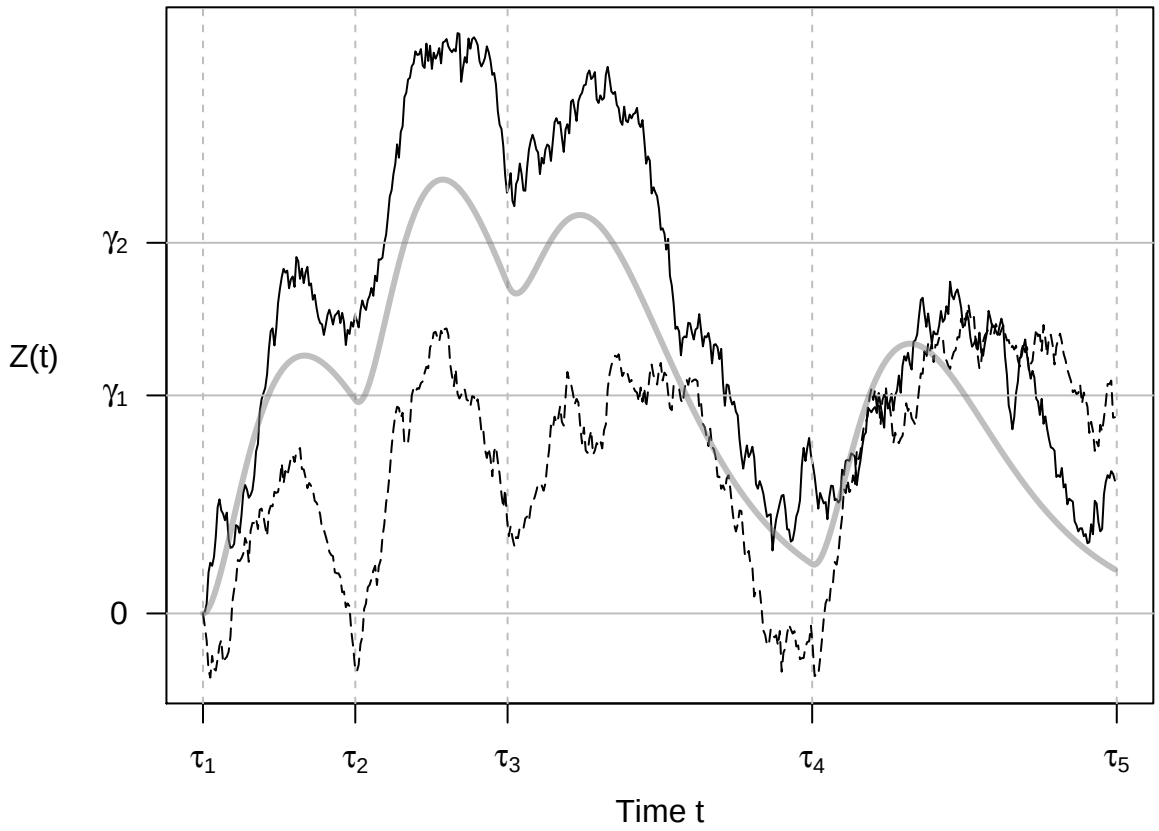

Figure 3.2. Two simulated stochastic latent toxicity processes for a patient receiving four dose administrations at times τ_1, τ_2, τ_3, and τ_4. The processes share a common deterministic component $\mu(t)$, shown as the smooth gray curve.

reality, a patient experiencing DLT would likely receive no further treatment, and we would not observe Y_c for cycles after the DLT. To clarify the model, in Figure 3.3 we present a causal diagram of all model parameters, latent processes, and observed data, showing how they relate to each other.

3.3 Specifying the Likelihood

In the preceding formulation, for the sake of simplicity we omitted a subscript i needed to indicate a specific patient. However, to specify the likelihood of all model parameters, we formulate the model in greater generality. Each patient may not receive the full number of cycles m, typically due to experiencing a DLT, and the

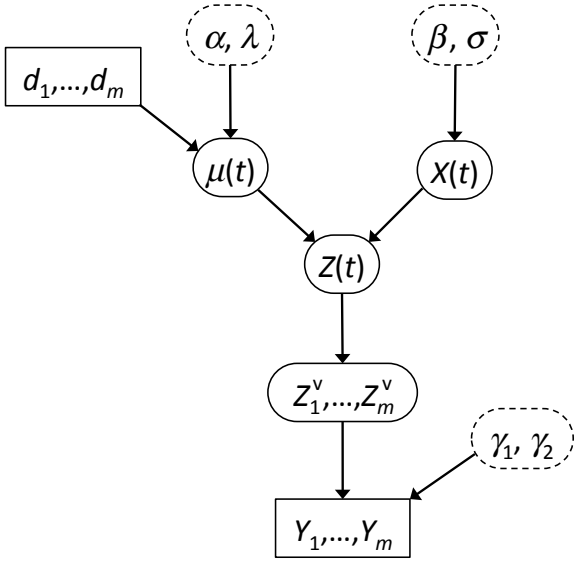

Figure 3.3. Causal diagram of latent toxicity process model. Arrows indicate the direction of causality. Parameters are surrounded by curved, dashed lines. The latent processes are surrounded by curved, solid lines. And the observed data are surrounded by rectangular, solid lines. The parameters are shared among patients, while the data and latent processes are unique to each patient.

sequence of dose assignments will likely vary between patients. Thus, let $m_i \leq m$ be the number of cycles and $d_{i1}, d_{i2}, \ldots, d_{i,m_i}$ the doses for patient i, with $i = 1, \ldots, n$. For simplicity, we assume that the timing of dose administrations is the same across patients, with dose cycle c starting at time τ_c after enrollment.

Ideally, we would derive the distribution of $Z_{ic}^{\vee}$, the maximum of our latent process for patient i in cycle c, to determine the likelihood, which bypasses the underlying continuous process $Z_i(t)$. However, because the distribution of $Z_{ic}^{\vee}$ is analytically intractable, we take the following approach. Noting that τ_{m+1} is the maximum time any patient is observed, we divide the interval $[0, \tau_{m+1}]$ into a fine grid of K times, $0 = t_1, t_2, \ldots, t_K = \tau_{m+1}$, all equally spaced with constant difference $\Delta_K = \tau_{m+1}/(K-1)$. Let $K_i \leq K$ be the number of time points t_k, $k = 1, \ldots, K$, that are no greater than τ_{m_i+1}, i.e., the time points that are in the observation period of patient i. We define $Z_{ik} = Z_i(t_k)$, and letting $\mathcal{C}_c = (\tau_c, \tau_{c+1}]$ denote the window for treatment cycle c, we have $Z_{ic}^{\vee} \approx \max\{Z_{ik} : t_k \in \mathcal{C}_c\}$. As $K \to \infty$, this discrete approximation becomes

exact. In the cases we consider in Section 3.8, we find that $K = 70$ is sufficient for $m = 3$ cycles, and $K = 140$ is sufficient for $m = 6$ cycles.

The observed outcome Y_{ic} is defined as in Equation (3.3), but with the i subscript where appropriate. The mean function $\mu(t)$ in Equation (3.1) is similarly extended to $\mu_{ik} = \mu_i(t_k)$. For convenience, let $\boldsymbol{Y}$ and $\boldsymbol{Z}$ be the vectors of all Y_{ic} and Z_{ik}, respectively, and let $\boldsymbol{\theta} = (\alpha, \lambda, \beta, \sigma)$ and $\boldsymbol{\gamma} = (\gamma_1, \gamma_d)$, where $\gamma_d = \gamma_2 - \gamma_1$. Then, given $\boldsymbol{\theta}$ and $\boldsymbol{\gamma}$, the joint density of $(\boldsymbol{Y}, \boldsymbol{Z})$ is

$$\begin{aligned}
[\boldsymbol{Y}, \boldsymbol{Z}|\boldsymbol{\theta}, \boldsymbol{\gamma}] &= [\boldsymbol{Y}|\boldsymbol{Z}, \boldsymbol{\theta}, \boldsymbol{\gamma}][\boldsymbol{Z}|\boldsymbol{\theta}, \boldsymbol{\gamma}] \\
&= [\boldsymbol{Y}|\boldsymbol{Z}, \boldsymbol{\gamma}][\boldsymbol{Z}|\boldsymbol{\theta}] \\
&= \prod_{i=1}^{n} \left(\prod_{c=1}^{m_i} [Y_{ic}|Z_{ic}^{\vee}, \boldsymbol{\gamma}] \times \prod_{k=2}^{K_i} [Z_{ik}|Z_{i,k-1}, \boldsymbol{\theta}] \right).
\end{aligned} \tag{3.4}$$

Note that the last product starts at $k = 2$ rather than $k = 1$ because Z_{i1} is fixed at $Z_i(0) = 0$. The notation $[\cdot|\cdot]$ used above and in the rest of this chapter denotes a probability density or mass function $p(\cdot|\cdot)$.

We can think of Equation (3.4) as the full data likelihood. The densities $[Y_{ic}|Z_{ic}^{\vee}, \boldsymbol{\gamma}]$ and $[Z_{ik}|Z_{i,k-1}, \boldsymbol{\theta}]$ can be readily specified given our model formulation:

$$[Y_{ic}|Z_{ic}^{\vee}, \boldsymbol{\gamma}] = \begin{cases} \mathbf{1}(Z_{ic}^{\vee} < \gamma_1) & \text{if } Y_{ic} = 0 \\ \mathbf{1}\big(\gamma_1 \le Z_{ij}^{\vee} < \gamma_1 + \gamma_d\big) & \text{if } Y_{ic} = 1 \\ \mathbf{1}(Z_{ic}^{\vee} \ge \gamma_1 + \gamma_d) & \text{if } Y_{ic} = 2 \end{cases}$$

and

$$[Z_{ik}|Z_{i,k-1}, \boldsymbol{\theta}] = \frac{1}{s}\phi\left(\frac{Z_{ik} - \mu_{ik} - (Z_{i,k-1} - \mu_{i,k-1})e^{-\beta\Delta_K}}{s} \right) \tag{3.5}$$

where $s^2 = \frac{\sigma^2}{2\beta}\big(1 - e^{-2\beta\Delta_K}\big)$ and $\phi(\cdot)$ is the standard normal density. Note that $[Y_{ic}|Z_{ic}^{\vee}, \boldsymbol{\gamma}]$ is a degenerate density, always equal to either 1 or 0 because Y_{ic} is completely determined by $Z_{ic}^{\vee}$ and $\boldsymbol{\gamma}$, according to Equation (3.3).

3.4 Parameter Estimation

We take a Bayesian approach to parameter estimation, and use the Metropolis-Hastings (MH) algorithm [42, 48] to obtain samples from the posterior density

$$[\boldsymbol{Z}, \boldsymbol{\theta}, \boldsymbol{\gamma} | \boldsymbol{Y}] \propto [\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{\theta}, \boldsymbol{\gamma}] = [\boldsymbol{Y}, \boldsymbol{Z} | \boldsymbol{\theta}, \boldsymbol{\gamma}][\boldsymbol{\theta}, \boldsymbol{\gamma}] = [\boldsymbol{Y} | \boldsymbol{Z}, \boldsymbol{\gamma}][\boldsymbol{Z} | \boldsymbol{\theta}][\boldsymbol{\theta}, \boldsymbol{\gamma}].$$

Here $[\boldsymbol{\theta}, \boldsymbol{\gamma}]$ is the joint prior distribution for our parameters, and we assume that all parameters are *a priori* independent. Our desire for prior independence provides our rationale for using (γ_1, γ_d) instead of (γ_1, γ_2), as the latter could not reasonably be assumed independent, since $\gamma_2 > \gamma_1$. We use the following priors:

$$\begin{array}{lll} \alpha \sim \text{Log-}\mathcal{N}\big(\mu_\alpha, \sigma_\alpha^2\big), & \beta \sim \text{Log-}\mathcal{N}\big(\mu_\beta, \sigma_\beta^2\big), & \gamma_1 \sim \mathcal{N}\big(\mu_{\gamma_1}, \sigma_{\gamma_1}^2\big) I(0, \infty), \\ \lambda \sim \text{Log-}\mathcal{N}\big(\mu_\lambda, \sigma_\lambda^2\big), & \sigma^2 \sim \text{InverseGamma}(a_\sigma, b_\sigma), & \gamma_d \sim \mathcal{N}\big(\mu_{\gamma_d}, \sigma_{\gamma_d}^2\big) I(0, \infty), \end{array} \tag{3.6}$$

where Log-$\mathcal{N}(\cdot, \cdot)$ is the lognormal distribution, and $\mathcal{N}(\cdot, \cdot) I(a, b)$ is the truncated normal distribution, restricted to the interval (a, b). The prior means are thus $\mathrm{E}[\alpha] = \exp(\mu_\alpha + \sigma_\alpha^2/2)$, $\mathrm{E}[\lambda] = \exp(\mu_\lambda + \sigma_\lambda^2/2)$, $\mathrm{E}[\beta] = \exp(\mu_\beta + \sigma_\beta^2/2)$, $\mathrm{E}[\sigma^2] = b_\sigma/(a_\sigma - 1)$ for $a_\sigma > 1$, $\mathrm{E}[\gamma_1] = \mu_{\gamma_1} + \sigma_{\gamma_1} \phi(\mu_{\gamma_1}/\sigma_{\gamma_1})/\Phi(\mu_{\gamma_1}/\sigma_{\gamma_1})$ where $\Phi(\cdot)$ is the standard normal cumulative density, and $\mathrm{E}[\gamma_d] = \mu_{\gamma_d} + \sigma_{\gamma_d} \phi(\mu_{\gamma_d}/\sigma_{\gamma_d})/\Phi(\mu_{\gamma_d}/\sigma_{\gamma_d})$.

For σ, γ_1, and γ_2, the chosen priors are semi-conjugate, i.e., the full conditional distributions belong to the same family as the priors, so we can use Gibbs sampling steps for these parameters within the MH algorithm. For α, λ, and β, closed-form full conditionals are not available, so we use regular MH sampling with truncated normal proposals. First we derive the full conditionals for γ_1, γ_d, and σ:

$$\begin{aligned} [\gamma_1 | \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{\theta}, \gamma_d] \propto [\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{\theta}, \boldsymbol{\gamma}] = [\boldsymbol{Y}, \boldsymbol{Z} | \boldsymbol{\theta}, \boldsymbol{\gamma}][\boldsymbol{\theta}, \boldsymbol{\gamma}] &\propto \left(\prod_{i=1}^{n} \prod_{c=1}^{m_i} [Y_{ic} | Z_{ic}^{\vee}, \boldsymbol{\gamma}] \right) [\gamma_1] \\ &= [\gamma_1] \prod_{i,c: Y_{ic}=0} \mathbf{1}(Z_{ic}^{\vee} < \gamma_1) \prod_{i,c: Y_{ic}=1} \mathbf{1}(\gamma_1 \le Z_{ic}^{\vee} < \gamma_1 + \gamma_d) \prod_{i,c: Y_{ic}=2} \mathbf{1}(Z_{ic}^{\vee} > \gamma_1 + \gamma_d) \\ &= [\gamma_1] \, \mathbf{1}\Big(\gamma_1 > Z^{\vee(0)}\Big) \mathbf{1}\Big(Z^{\vee(1)} - \gamma_d \le \gamma_1 < Z^{\wedge(1)}\Big) \mathbf{1}\Big(\gamma_1 < Z^{\wedge(2)} - \gamma_d\Big) \\ &= [\gamma_1] \, \mathbf{1}\Big(\max\{Z^{\vee(0)}, Z^{\vee(1)} - \gamma_d\} \le \gamma_1 < \min\{Z^{\wedge(1)}, Z^{\wedge(2)} - \gamma_d\} \Big) \end{aligned}$$

where $Z^{\vee(\ell)} = \max_{i,c:Y_{ic}=\ell}\{Z^{\vee}_{ic}\}$ and $Z^{\wedge(\ell)} = \min_{i,c:Y_{ic}=\ell}\{Z^{\vee}_{ic}\}$. Similarly,

$$\begin{aligned}
[\gamma_d|\boldsymbol{Y},\boldsymbol{Z},\boldsymbol{\theta},\gamma_1] &\propto \left(\prod_{i=1}^{n}\prod_{c=1}^{m_i}[Y_{ic}|Z^{\vee}_{ic},\boldsymbol{\gamma}]\right)[\gamma_d] \\
&= [\gamma_d]\prod_{i,c:Y_{ic}=1}\mathbf{1}(\gamma_1 \le Z^{\vee}_{ic} < \gamma_1+\gamma_d)\prod_{i,c:Y_{ic}=2}\mathbf{1}(Z^{\vee}_{ic} > \gamma_1+\gamma_d) \\
&= [\gamma_d]\,\mathbf{1}\Big(Z^{\vee(1)}-\gamma_1 \le \gamma_d < Z^{\wedge(2)}-\gamma_1\Big).
\end{aligned}$$

Thus

$$\begin{aligned}
\gamma_1|\boldsymbol{Y},\boldsymbol{Z},\boldsymbol{\theta},\gamma_d &\sim \mathcal{N}\big(\mu_{\gamma_1},\sigma^2_{\gamma_1}\big)I(\max\{Z^{\vee(0)},Z^{\vee(1)}-\gamma_d\},\min\{Z^{\wedge(1)},Z^{\wedge(2)}-\gamma_d\}), \\
\gamma_d|\boldsymbol{Y},\boldsymbol{Z},\boldsymbol{\theta},\gamma_1 &\sim \mathcal{N}\big(\mu_{\gamma_d},\sigma^2_{\gamma_d}\big)I(Z^{\vee(1)}-\gamma_1, Z^{\wedge(2)}-\gamma_1).
\end{aligned}$$

Once we have sampled γ_1 and γ_d, then we have $\gamma_2 = \gamma_1+\gamma_d$. For σ^2 we have

$$\begin{aligned}
[\sigma^2|\boldsymbol{Y},\boldsymbol{Z},\alpha,\lambda,\beta,\boldsymbol{\gamma}] &\propto [\boldsymbol{Z}|\sigma^2,\alpha,\lambda,\beta][\sigma^2] \\
&= \left(\prod_{i=1}^{n}\prod_{k=2}^{K_i}[Z_{ik}|Z_{i,k-1},\sigma^2,\alpha,\lambda,\beta]\right)[\sigma^2].
\end{aligned}$$

Combining this expression with Equation (3.5) and the prior for σ^2, some algebra shows that $\sigma^2|\boldsymbol{Y},\boldsymbol{Z},\alpha,\lambda,\beta,\boldsymbol{\gamma} \sim \text{InverseGamma}(\tilde{a}_\sigma,\tilde{b}_\sigma)$, where

$$\begin{aligned}
\tilde{a}_\sigma &= a_\sigma + \frac{1}{2}\sum_{i=1}^{n}K_i - \frac{n}{2} \quad \text{and} \\
\tilde{b}_\sigma &= b_\sigma + \frac{\beta}{1-e^{-2\beta\Delta_K}}\sum_{i=1}^{n}\sum_{k=2}^{K_i}\big[Z_{ik}-\mu_{ik}-(Z_{i,k-1}-\mu_{i,k-1})e^{-\beta\Delta_K}\big]^2.
\end{aligned}$$

Although we might not think of the Z_{ik}, $i = 1,\ldots,n$ and $k = 2,\ldots,K_i$, as parameters, they are unobserved and thus we need to sample them just like parameters. For all Z_{ik} such that $t_k \in \mathcal{C}_c$,

$$\begin{aligned}
[Z_{ik}|\boldsymbol{Y},\boldsymbol{Z}\setminus Z_{ik},\boldsymbol{\theta},\boldsymbol{\gamma}] &\propto [\boldsymbol{Y}|\boldsymbol{Z},\boldsymbol{\gamma}][\boldsymbol{Z}|\boldsymbol{\theta}] \\
&\propto [Y_{ic}|Z^{\vee}_{ic},\boldsymbol{\gamma}] \times \begin{cases} [Z_{ik}|Z_{i,k-1},\boldsymbol{\theta}][Z_{i,k+1}|Z_{ik},\boldsymbol{\theta}] & \text{if } k < K_i \\ [Z_{ik}|Z_{i,k-1},\boldsymbol{\theta}] & \text{if } k = K_i. \end{cases}
\end{aligned}$$

Thus $Z_{ik}|\boldsymbol{Y}, \boldsymbol{Z} \setminus Z_{ik}, \boldsymbol{\theta}, \boldsymbol{\gamma} \sim \mathcal{N}(M_{ik}, V_{ik})I(L_{ik}, U_{ik})$, where

$$M_{ik} = \begin{cases} \mu_{ik} + [(Z_{i,k-1} - \mu_{i,k-1}) + (Z_{i,k+1} - \mu_{i,k+1})]\frac{e^{-\beta\Delta_K}}{1+e^{-2\beta\Delta_K}} & \text{if } k < K_i \\ \mu_{ik} + [Z_{i,k-1} - \mu_{i,k-1}]e^{-\beta\Delta_K} & \text{if } k = K_i, \end{cases}$$

$$V_{ik} = \begin{cases} s^2/(1 + e^{-2\beta\Delta_K}) & \text{if } k < K_i \\ s^2 & \text{if } k = K_i, \end{cases}$$

$$L_{ik} = \begin{cases} \gamma_1 & \text{if } Y_{ic} = 1 \text{ and } \max\{Z_{ik'} : t_{k'} \in \mathcal{C}_c \setminus t_k\} < \gamma_1 \\ \gamma_2 & \text{if } Y_{ic} = 2 \text{ and } \max\{Z_{ik'} : t_{k'} \in \mathcal{C}_c \setminus t_k\} < \gamma_2 \\ -\infty & \text{otherwise,} \end{cases}$$

$$U_{ik} = \begin{cases} \gamma_1 & \text{if } Y_{ic} = 0 \\ \gamma_2 & \text{if } Y_{ic} = 1 \text{ and } \max\{Z_{ik'} : t_{k'} \in \mathcal{C}_c \setminus t_k\} < \gamma_2 \\ \infty & \text{otherwise.} \end{cases}$$

That is, the full conditional distribution for Z_{ik} is a truncated normal, so Gibbs sampling will work for Z_{ik} as well.

Next we describe how to use MH to obtain posterior samples of α, λ, and β. For α, with a proposed value α^* from proposal distribution J_α, the acceptance ratio is

$$\begin{aligned} r_\alpha &= \frac{[\boldsymbol{Z}, \alpha^*, \lambda, \beta, \sigma, \boldsymbol{\gamma}|\boldsymbol{Y}]}{[\boldsymbol{Z}, \alpha, \lambda, \beta, \sigma, \boldsymbol{\gamma}|\boldsymbol{Y}]} \frac{J_\alpha(\alpha|\alpha^*)}{J_\alpha(\alpha^*|\alpha)} \\ &= \frac{[\boldsymbol{Z}|\alpha^*, \lambda, \beta, \sigma]}{[\boldsymbol{Z}|\alpha, \lambda, \beta, \sigma]} \frac{[\alpha^*]}{[\alpha]} \frac{J_\alpha(\alpha|\alpha^*)}{J_\alpha(\alpha^*|\alpha)} \\ &= \left(\prod_{i=1}^{n} \prod_{k=2}^{K_i} \frac{[Z_{ik}|Z_{i,k-1}, \alpha^*, \lambda, \beta, \sigma]}{[Z_{ik}|Z_{i,k-1}, \alpha, \lambda, \beta, \sigma]} \right) \frac{[\alpha^*]}{[\alpha]} \frac{J_\alpha(\alpha|\alpha^*)}{J_\alpha(\alpha^*|\alpha)}. \end{aligned}$$

In the last expression, the first term can be calculated with Equation (3.5), the second term comes from the prior distribution of α, and the third term is the proposal distribution, which we have chosen. With r_α computed, we use it as follows. Draw $u \sim \text{Uniform}(0, 1)$. If $u < r_\alpha$, we accept the proposed value α^*, using it as the next value in our Markov chain. Otherwise, we reject α^* and reuse the previous value of α in the chain.

The acceptance ratios for λ and β, which are derived and used analogously to that for α, are

$$r_\lambda = \left(\prod_{i=1}^{n} \prod_{k=2}^{K_i} \frac{[Z_{ik}|Z_{i,k-1}, \alpha, \lambda^*, \beta, \sigma]}{[Z_{ik}|Z_{i,k-1}, \alpha, \lambda, \beta, \sigma]} \right) \frac{[\lambda^*]}{[\lambda]} \frac{J_\lambda(\lambda|\lambda^*)}{J_\lambda(\lambda^*|\lambda)},$$

$$r_\beta = \left(\prod_{i=1}^{n} \prod_{k=2}^{K_i} \frac{[Z_{ik}|Z_{i,k-1}, \alpha, \lambda, \beta^*, \sigma]}{[Z_{ik}|Z_{i,k-1}, \alpha, \lambda, \beta, \sigma]} \right) \frac{[\beta^*]}{[\beta]} \frac{J_\beta(\beta|\beta^*)}{J_\beta(\beta^*|\beta)}.$$

We use truncated normal proposal distributions for α, λ, and β.

After the Markov chain has converged and we have obtained a large sample from the joint posterior of all parameters, we estimate the parameters by their posterior means, denoted $\hat{\alpha}$, $\hat{\lambda}$, $\hat{\beta}$, $\hat{\sigma}$, $\hat{\gamma_1}$, and $\hat{\gamma_2}$.

3.5 Model Calibration

Typically in phase I trials, the dose values used in statistical models are not literally the dosage, e.g., milligrams, of a drug, but rather come from a probability skeleton. The skeleton is the *a priori* best guess of the probability of DLT at each dose, and it is used to calibrate the model. We take the same approach here, with "probability of DLT" meaning the probability of a DLT ever occurring across the multiple cycles of administration, assuming the same dose level is given in every cycle.

Let J be the number of dose levels under investigation, and π_{0j}, $j = 1, \ldots, J$ be the skeleton, where $\pi_{01} < \cdots < \pi_{0J}$. We take the following approach to calibrate our model to agree with the skeleton for dose level j. Recall that the latent $Z(t)$ is the sum of $\mu(t)$ and $X(t)$, where only $\mu(t)$ depends on dose. With the same dose $d_j^* > 0$ across cycles, Equation (3.1) simplifies to $\mu(t) = d_j^* \sum_{c:\tau_c \le t} h(t - \tau_c)$. Our task is then to find the value d_j^* such that $\pi_{0j} = \Pr(Z(t) \ge \gamma_2 \text{ for any } t \in [0, \tau_{m+1}]|\boldsymbol{\theta}, \boldsymbol{\gamma})$, where m is the number of planned cycles. We can solve this equation numerically for each j as follows:

1. For a given value of d_j^*, simulate $Z(t)$ over a fine time grid in $[0, \tau_{m+1}]$ using Equation (3.2), setting α, λ, β, and σ to their prior means.
2. Check if $Z(t)$ exceeds γ_2 for any t, i.e., if there was a DLT, where γ_2 is set to its prior mean $\mathrm{E}[\gamma_2] = \mathrm{E}[\gamma_1] + \mathrm{E}[\gamma_d]$.
3. Repeat steps 1 and 2 many (e.g. 10,000) times, and count the proportion in which a DLT is observed.
4. If the observed proportion of DLTs is close enough to π_{0j}, we are done; otherwise, adjust d_j^* and return to step 1.

With these values calculated, if patient i received dose level j in cycle c, then we set $d_{ic} = d_j^*$.

3.6 Estimating DLT Probabilities

Once we have estimated the parameters of our model, we can estimate the DLT probability at each dose. As in the preceding section, by "DLT probability" at dose j, which we denote π_j, we mean the probability that a patient receiving dose j in every cycle will ever have a DLT. That is, $\pi_j = \Pr(Z(t) \geq \gamma_2 \text{ for any } t \in [0, \tau_{m+1}] | \boldsymbol{\theta}, \boldsymbol{\gamma})$. We can estimate this probability numerically, for dose level j as follows:

1. Simulate $Z(t)$ over a fine time grid in $[0, \tau_{m+1}]$, assuming constant dose level j, and using $\hat{\alpha}$, $\hat{\lambda}$, $\hat{\beta}$, $\hat{\sigma}$ as the model parameters.
2. Check if $Z(t)$ exceeds $\hat{\gamma_2}$ for any t, i.e., if there was a DLT.
3. Repeat steps 1 and 2 many (e.g. 10,000) times, and count the proportion in which a DLT is observed. This is $\hat{\pi}_j$.

3.7 Adjusting a Patient's Dose

With a multi-cycle treatment, it is natural to consider adjusting a patient's dose level across cycles. For example, if a patient has gone through the first few cycles with no toxicity, but has not obviously benefitted from the treatment (e.g., as measured by tumor shrinkage), we might consider increasing the dose for later cycles. Our model can show how the probability of DLT will be altered.

Assume a patient received dose j for the first m_1 cycles, with outcomes $Y_1, \ldots, Y_{m_1}$. Then the following algorithm will estimate the probability of DLT in later cycles $m_1 + 1, \ldots, m$, under any dose scheme for those cycles.

1. Simulate $Z(t)$ over a fine time grid in $[0, \tau_{m+1}]$, using $\hat{\alpha}$, $\hat{\lambda}$, $\hat{\beta}$, $\hat{\sigma}$ as the model parameters. For cycles $1, \ldots, m_1$, the dose level is j. For cycles $m_1 + 1, \ldots, m$, the dose is whatever we want to consider changing to.
2. If the simulated $Z(t)$ does not agree with the observed $Y_1, \ldots, Y_{m_1}$, then discard it and repeat step 1.
3. Check if $Z(t) > \hat{\gamma_2}$ for any $t > \tau_{m_1+1}$, i.e., if there was a DLT in a later cycle.
4. Repeat steps 1-3 many (e.g. 10,000) times, and count the proportion in which a DLT is observed, ignoring the processes discarded in step 2. This proportion is our estimate of the DLT probability.

The above algorithm can be carried out for multiple proposed dose schemes, to compare the probability of DLT in each scheme.

3.8 Simulation Study

We conduct a simulation study, to investigate the LTP model's ability to select the correct MTD using data from a trial run with the CRM. The CRM will only use patient data from cycle 1, using the power model $\Pr(Y_{i1} = 2|\text{dose } j) = \nu_{0j}^{\exp(\psi)}$ with

$\psi \sim \mathcal{N}(0, \sigma_\psi^2 = 1.34)$. The prior variance $\sigma_\psi^2 = 1.34$ is the default in the R package `dfcrm` [40, 49]. The skeleton ν_{0j}, $j = 1, \ldots, J$, is our prior guess for the probability of DLT in cycle 1 with dose j. To set ν_{0j}, we first choose π_{0j}, the skeleton over all m planned cycles, and we assume that DLTs are independent and identically distributed (*iid*) across cycles. Thus $\pi_{0j} = 1 - (1 - \nu_{0j})^m$, so we use $\nu_{0j} = 1 - (1 - \pi_{0j})^{1/m}$. After calculating $\hat{\psi}$ as the posterior mean of ψ, the estimated probability of DLT in cycle 1 is $\nu_{0j}^{\exp(\hat{\psi})}$, and the estimated probability of DLT extrapolated to all m cycles is $\hat{\pi}_j^{\text{CRM}} = 1 - [1 - \nu_{0j}^{\exp(\hat{\psi})}]^m$. Although we do not truly believe the *iid* assumption, it allows us to use the CRM to estimate full-cycle DLT rates using only cycle 1 DLTs. A simple alternative would be to set $\nu_{0j} = \pi_{0j}$ and take $\hat{\pi}_j^{\text{CRM}} = \nu_{0j}^{\exp(\hat{\psi})}$, but this leads to much worse performance for the CRM than our approach.

We simulate trials with $n = 30$ and $n = 60$ patients, for treatments with $m = 3$ and $m = 6$ cycles, with evenly spaced dose administrations at times $0, 1, \ldots, m - 1$ and with final follow-up at time m. Data are simulated with the LTP model, using "true" parameter values $\alpha = 0.3$, $\lambda = 2$, $\beta = 2$, $\sigma = 0.7$, $\gamma_1 = 0.75$, and $\gamma_2 = 1$ for $m = 3$; and $\alpha = 0.2$, $\lambda = 1.5$, $\beta = 1.5$, $\sigma = 0.6$, $\gamma_1 = 0.75$, and $\gamma_2 = 1$ for $m = 6$. These values were chosen to generate outcome data consistent with the findings in Postel-Vinay et al. [26]. In particular, for patients who have a DLT, the model-based probability of DLT occurring after cycle 1 falls between 0.55 and 0.75, as depicted in Figure 3.4 for $m = 3$, while the published number is 0.57. To calibrate the model to achieve a desired DLT rate π_j for dose j, we calculate a value d_j^* as described in Section 3.5, except we use the true parameter values instead of prior means, and we want the probability of DLT to equal π_j rather than the skeleton π_{0j}.

The simulated trials will consider $J = 5$ dose levels, and we examine 5 scenarios for the true DLT rates, one for each dose level to be the MTD. We define the "DLT rate" for a given dose level as the probability of ever having a DLT across the m cycles, if the same dose level is administered in every cycle. The true DLT rates for

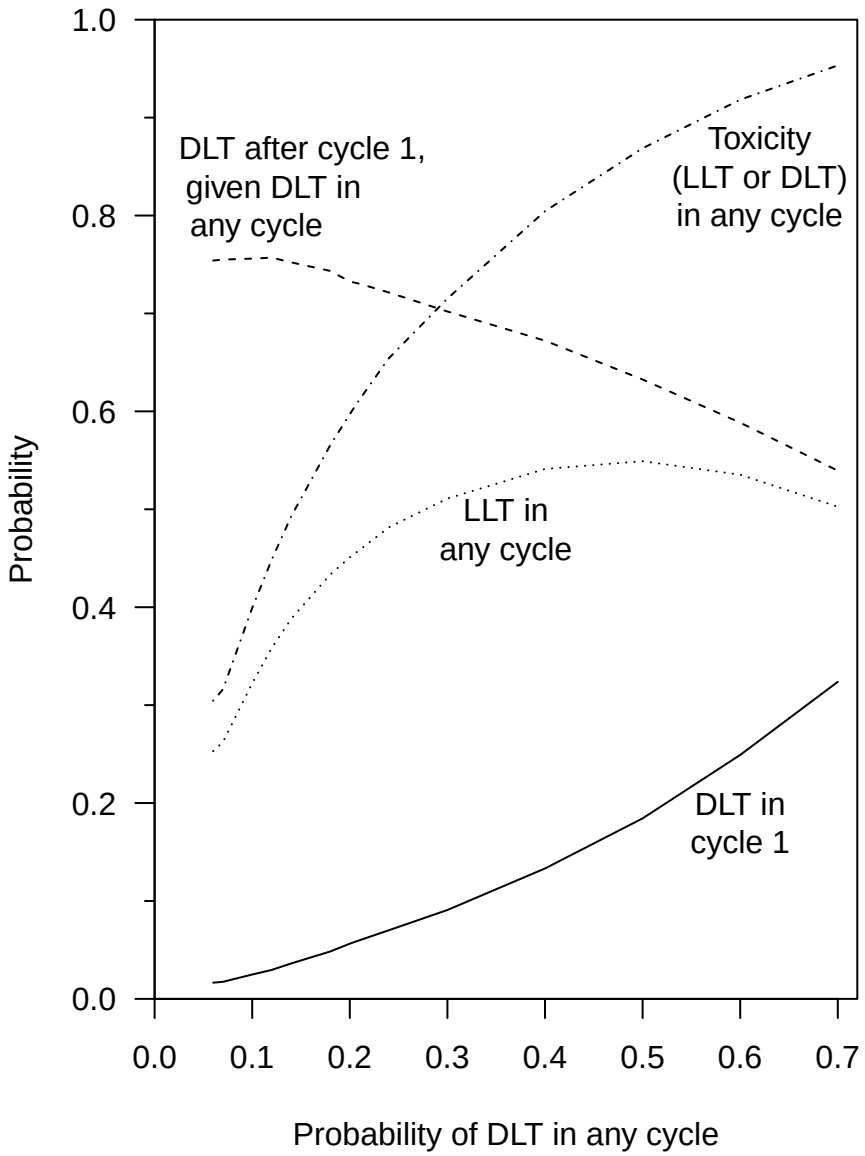

Figure 3.4. Probabilities of toxicity by cycle and overall, for a three-cycle treatment. The probabilities are calculated with the latent toxicity process model, assuming a constant dose across cycles. The patient visit times are 0, 1, 2, and 3, with dose administrations at the first three visits. The model parameters were set to $\alpha = 0.3$, $\lambda = 2$, $\beta = 2$, $\sigma = 0.7$, $\gamma_1 = 0.75$, $\gamma_2 = 1$. Note that if a patient has a DLT, their treatment is stopped and we do not observe any subsequent toxicities. Thus, we can only observe an LLT if it is not preceded by a DLT, and a DLT can only be observed after cycle 1 if there was no DLT in cycle 1.

the dose levels in each scenario are given in Table 3.1. The target DLT rate is always $\tau = 0.3$.

To simulate one trial, we start the first patient at dose level 1 and simulate their multi-cycle outcomes with the LTP model. After calculating $\hat{\pi}_j^{\text{CRM}}$, $j = 1, \ldots, J$, as described above, we estimate the MTD, denoted $\widehat{\text{MTD}}$, as the dose j that minimizes $|\hat{\pi}_j^{\text{CRM}} - \tau|$. The next patient will then be assigned to either the $\widehat{\text{MTD}}$, or to one dose higher than the current patient, whichever is smaller. This process is repeated for n patients, after which we obtain a final $\widehat{\text{MTD}}$ from the CRM. With the complete trial data, including LLTs and data from all cycles, we then fit the LTP model, which produces its own estimates $\hat{\pi}_j$, $j = 1, \ldots, J$, and a corresponding estimate of the MTD. Note that, if a patient had a DLT in cycle c, then their later outcomes are considered unobserved. Additionally, if a patient had LLTs in two consecutive cycles, then we lowered this patient's dose by one level for the next cycle. However, this dose de-escalation does not factor into the adaptive dose-finding algorithm with the CRM. For each scenario, we simulate 5,000 adaptive trials with the CRM. We then subsample 500 completed trials in each scenario and re-estimate the MTD with the LTP model.

To set up the CRM and LTP models, we use the same skeleton of $(\pi_{01}, \ldots, \pi_{0J}) = (0.1, 0.2, 0.3, 0.4, 0.5)$ for all scenarios. Note that the skeleton matches the truth for

Table 3.1. True probability of DLT over all cycles, for each simulation scenario and dose level. Since the target DLT rate is 0.30, the true MTD in scenario j is dose level j, for $j = 1, \ldots, 5$.

	Dose level				
Scenario	1	2	3	4	5
1	**0.30**	0.40	0.50	0.60	0.70
2	0.20	**0.30**	0.40	0.50	0.60
3	0.10	0.20	**0.30**	0.40	0.50
4	0.07	0.14	0.21	**0.30**	0.40
5	0.06	0.12	0.18	0.24	**0.30**

scenario 3. As mentioned above, we use $\sigma_\psi^2 = 1.34$ for the prior variance of the CRM parameter. For the LTP model, we assume that some of the parameters are known. In particular, we propose eliciting α and λ from experts, as these parameters have natural interpretations related to the timing to toxicity described in Section 3.2. Legedza and Ibrahim [15] used a similar simplifying assumption for their longitudinal model. We initially tried to estimate α and λ, but found that the Markov chain would not converge in a reasonable time. Additionally, it is likely that the data are only weakly informative for these parameters, as we do not observe when within a cycle a toxicity occurred. Thus, for estimation we set α and λ to their true values, respectively 0.3 and 2 for $m = 3$, and 0.2 and 1.5 for $m = 6$.

To continue setting up the LTP model, for β, γ_1, and γ_d we specify the hyper-parameters in Equation (3.6), using $(\mu_\beta, \sigma_\beta) = (-0.2, 0.63)$, $(\mu_{\gamma_1}, \sigma_{\gamma_1}) = (0.92, 0.6)$, and $(\mu_{\gamma_d}, \sigma_{\gamma_d}) = (0.057, 0.6)$. The prior means of these parameters are thus $\mathrm{E}[\beta] = 1$, $\mathrm{E}[\gamma_1] = 1$, and $\mathrm{E}[\gamma_d] = 0.5$. For σ, rather than estimate it as proposed, we fix it at the value in $\Omega = \{0.2, 0.4, 0.6, 0.8\}$ that minimizes the "score" across all scenarios with $n = 60$, where score $= \sum_{j=1}^{J} \hat{s}_j |\pi_j - \tau|$ and $\hat{s}_j$ is the percentage of trials in which the model selects dose j as the MTD in a given scenario. The score thus measures the average distance between the target and the true DLT rate of the estimated MTD; smaller scores are better and zero is perfect. We average the score across all scenarios and choose $\sigma \in \Omega$ to minimize it. Note that, for $m = 3$, Ω does not exactly contain the true value of 0.7. This process yields $\sigma = 0.6$ for both $m = 3$ and $m = 6$, which we use for both $n = 30$ and $n = 60$. As with α and λ, we found that estimating σ with Metropolis-Hastings led to unacceptably slow convergence. Thus, the only parameters to be estimated with MH in the LTP model are β, γ_1, and γ_d. For each trial, we generated 100,000 samples of these parameters, and calculated posterior means using the last 50,000. Each sample was based on simulating the latent process at $K = 70$ time points for $m = 3$, and at $K = 140$ time points for $m = 6$. In small

test runs, we found that MH samples of this size and using these K were sufficient to achieve convergence and to precisely estimate the posterior means. Generating 100,000 samples took approximately 30 seconds on our computer with $n = 30$, and 60 seconds with $n = 60$ for $m = 3$; computational time roughly doubled for $m = 6$.

3.8.1 Simulation Results

Simulation results are presented in Table 3.2 for $m = 3$ and in Table 3.3 for $m = 6$. For the CRM and LTP model, we report how often each dose level is selected as the MTD, and the corresponding score. For the CRM, we additionally report how many patients, out of n, are on average treated at each dose. This tells us, on average, how many patients at each dose the LTP model has access to when re-estimating the MTD at the end of a trial.

The CRM does poorly in scenarios 2–4 because it has few events to work with in cycle 1. In scenario 4 with 60 patients and 3 cycles, the CRM selects the correct MTD in only 16.3% of trials (although it selects an MTD within one level of the truth 96.5% of the time). The true DLT rates in this scenario are $(0.07, 0.14, 0.21, 0.30, 0.40)$, however, the true probabilities of DLT in cycle 1, as depicted in Figure 3.4, are (0.017, 0.036, 0.06, 0.09, 0.13). The proximity of these first-cycle probabilities with each other makes it difficult to distinguish between them, particularly with a small sample size. In scenario 5 with 3 cycles, the cycle 1 probabilities are compressed even more, and the CRM tends to substantially underestimate the overall DLT rates. Indeed, for dose level 5, the CRM on average estimates the DLT rate to be 0.16 for $n = 60$, lower than the true rate of 0.3. However, because dose level 5 is the highest level under consideration, the CRM cannot escalate further and ends up selecting dose level 5, the correct MTD, 92.5% of the time with 60 patients. The LTP model tends to more accurately estimate DLT rates than the CRM, but paradoxically this leads the LTP model to select the correct MTD less often in scenario 5. The explanation

Table 3.2. Simulation results for $m = 3$ cycles with $n = 30$ and $n = 60$ patients. For each scenario and patient count n, we simulated 5,000 adaptive trials with the CRM. We then subsampled 500 of these completed trials and re-estimated the MTD with the latent toxicity process (LTP) model. In all scenarios, the target DLT rate is 0.3. The *italicized* rows show the true DLT rates, i.e., the probability of DLT over all cycles. The **bold** entries highlight the results at the true MTD.

Scenario	Number of Patients	Model	Dose Selection Percentage (Number of Patients Treated) 1	2	3	4	5	Score
1			*0.3*	*0.4*	*0.5*	*0.6*	*0.7*	
	$n = 30$	CRM	**63.4 (16.5)**	16.4 (4.5)	11.0 (3.3)	5.7 (2.4)	3.5 (3.3)	6.9
		LTP	**77.6**	21.0	1.4	0.0	0.0	2.4
	$n = 60$	CRM	**71.8 (37.1)**	17.8 (9.5)	7.3 (6.1)	2.4 (3.6)	0.7 (3.7)	4.3
		LTP	**85.4**	14.6	0.0	0.0	0.0	1.5
2			*0.2*	*0.3*	*0.4*	*0.5*	*0.6*	
	$n = 30$	CRM	32.5 (10.3)	**19.5 (4.8)**	17.5 (4.1)	12.7 (3.6)	17.8 (7.1)	12.9
		LTP	34.4	**48.6**	15.0	2.0	0.0	5.3
	$n = 60$	CRM	30.2 (19.8)	**26.0 (11.7)**	22.0 (10.1)	13.9 (7.8)	7.8 (10.6)	10.4
		LTP	33.4	**57.2**	9.4	0.0	0.0	4.3
3			*0.1*	*0.2*	*0.3*	*0.4*	*0.5*	
	$n = 30$	CRM	11.6 (6.0)	15.8 (4.3)	**17.4 (4.2)**	16.5 (4.1)	38.7 (11.3)	13.3
		LTP	4.4	43.4	**42.2**	8.6	1.4	6.4
	$n = 60$	CRM	5.7 (8.6)	14.5 (8.8)	**22.6 (10.4)**	24.7 (10.6)	32.4 (21.6)	11.6
		LTP	1.4	43.6	**52.4**	2.6	0.0	4.9
4			*0.07*	*0.14*	*0.21*	*0.3*	*0.4*	
	$n = 30$	CRM	3.6 (3.6)	6.6 (2.9)	9.3 (3.2)	**13.0 (3.6)**	67.5 (16.7)	9.5
		LTP	0.2	12.0	42.2	**35.2**	10.4	6.8
	$n = 60$	CRM	0.9 (4.2)	2.6 (4.2)	7.1 (5.9)	**16.3 (8.2)**	73.1 (37.5)	8.6
		LTP	0.2	6.0	44.2	**46.2**	3.4	5.3
5			*0.06*	*0.12*	*0.18*	*0.24*	*0.3*	
	$n = 30$	CRM	1.6 (2.6)	2.6 (2.1)	4.1 (2.2)	5.7 (2.6)	**85.9 (20.4)**	1.7
		LTP	0.4	4.2	14.6	36.8	**44.0**	4.8
	$n = 60$	CRM	0.4 (2.9)	1.0 (2.6)	2.1 (3.2)	4.0 (4.2)	**92.5 (47.1)**	0.8
		LTP	0.0	1.2	6.6	24.2	**68.0**	2.5

Table 3.3. Simulation results for $m = 6$ cycles with $n = 30$ and $n = 60$ patients. For each scenario and patient count n, we simulated 5,000 adaptive trials with the CRM. We then subsampled 500 of these completed trials and re-estimated the MTD with the latent toxicity process (LTP) model. In all scenarios, the target DLT rate is 0.3. The *italicized* rows show the true DLT rates, i.e., the probability of DLT over all cycles. The **bold** entries highlight the results at the true MTD.

Scenario	Number of Patients	Model	Dose Selection Percentage (Number of Patients Treated) 1	2	3	4	5	Score
1	$n = 30$		*0.3*	*0.4*	*0.5*	*0.6*	*0.7*	
		CRM	**66.1 (17.8)**	15.4 (4.2)	9.8 (2.8)	5.3 (2.1)	3.4 (3.1)	6.4
		LTP	**83.2**	16.1	0.7	0.0	0.0	1.7
	$n = 60$	CRM	**75.3 (39.7)**	14.8 (8.8)	7.1 (5.2)	2.1 (3)	0.6 (3.4)	3.8
		LTP	**88.6**	11.4	0.0	0.0	0.0	1.1
2	$n = 30$		*0.2*	*0.3*	*0.4*	*0.5*	*0.6*	
		CRM	31.5 (10.6)	**18.9 (4.4)**	17.6 (3.9)	13.4 (3.4)	18.6 (7.6)	13.2
		LTP	31.8	**51.7**	14.5	2.0	0.0	5.0
	$n = 60$	CRM	27.8 (18.7)	**26.9 (12.1)**	22.9 (10.8)	14.3 (8.4)	7.9 (10.0)	10.3
		LTP	28.4	**63.1**	8.5	0.0	0.0	3.7
3	$n = 30$		*0.1*	*0.2*	*0.3*	*0.4*	*0.5*	
		CRM	9.4 (5.8)	17.1 (4.4)	**17.1 (3.9)**	16.9 (4.2)	39.5 (11.6)	13.2
		LTP	4.4	39.1	**47.3**	8.1	1.1	5.8
	$n = 60$	CRM	5.7 (9.1)	15.6 (9.9)	**24.5 (10.9)**	24.5 (10.8)	29.6 (19.3)	11.1
		LTP	1.1	39.6	**56.7**	2.6	0.0	4.4
4	$n = 30$		*0.07*	*0.14*	*0.21*	*0.3*	*0.4*	
		CRM	3.4 (3.1)	6.1 (2.8)	10.1 (2.9)	**16.3 (5.2)**	64.1 (16.0)	9.1
		LTP	0.1	11.9	38.3	**40.1**	9.7	6.3
	$n = 60$	CRM	0.7 (3.2)	2.1 (4.1)	7.5 (6.6)	**20.4 (10.3)**	69.3 (35.8)	8.1
		LTP	0.4	5.9	38.5	**52.3**	2.9	4.8
5	$n = 30$		*0.06*	*0.12*	*0.18*	*0.24*	*0.3*	
		CRM	2.3 (2.9)	2.8 (1.7)	4.3 (2.0)	8.6 (3.2)	**81.9 (20.1)**	2.1
		LTP	0.2	3.9	11.3	32.9	**51.7**	4.1
	$n = 60$	CRM	0.9 (3.9)	1.4 (3.1)	3.1 (3.8)	5.1 (5.1)	**89.4 (44.1)**	1.2
		LTP	0.0	0.8	6.2	15.9	**77.1**	1.8

is that for dose levels 4 and 5, with true DLT rates of 0.24 and 0.3, respectively, the LTP-based DLT rate estimates hover around the target rate of 0.3, and sometimes dose 4 is closer, sometimes dose 5 is closer. In contrast, since the CRM substantially underestimates the DLT rate of dose 5, it is nearly always closest to the target. A similar explanation applies for scenario 1, in which the CRM tends to overestimate the DLT rates and thus often selects the lowest dose, which is the correct MTD. This dynamic holds for both the 3 and 6 cycle treatments.

In scenarios 2–4, the CRM selects the true MTD between 13% and 26.9% of the time. Correspondingly, it treats only around 4 patients for $n = 30$ and 10 patients for $n = 60$ at the true MTD. The LTP model thus has few patients treated at the MTD to use when estimating DLT rates, however it still manages to more than double the selection percentages of the true MTD, compared to the CRM. Additionally, we note that increasing the sample size from 30 to 60 patients increases the true MTD selection percentages by about 10% for the LTP model in scenarios 1–4, and by about 25% in scenario 5. Lastly, with 6 cycles instead of 3, the LTP model selects the true MTD an additional 3–9% of the time. That is, with more cycles, patients provide more data for the LTP model, leading to better DLT rate estimates.

3.9 Discussion

We have developed a model for longitudinal toxicity responses in phase I trials, that accounts for within-subject correlation and cumulative dose effect by positing a continuous latent toxicity process. The process is designed to mimic the pharmacokinetics of drug absorption and clearance, tending to increase after a new dose is administered and the drug is absorbed, before falling off as the drug is cleared. Our model was inspired by earlier work from Legedza and Ibrahim [15] and Zhang and Braun [17], particularly in our definition of the mean dose function $\mu(t)$. However, while these

earlier works assumed that a patient's responses across cycles were mutually independent, we naturally account for the correlation through the Ornstein-Uhlenbeck stochastic process. Additionally, borrowing from Doussau et al. [16], we incorporate LLTs in our model, allowing us to more precisely determine the location of a patient's latent process over time, and thus improving the estimation of DLT rates. However, our use of a trinary ordinal response (i.e., no toxicity, LLT, and DLT), while convenient, is somewhat arbitrary. The original ordinal responses in phase I trials typically have six levels, and our model can be extended to handle more levels by adding more threshold parameters. The utility of adding more levels should be explored.

In our simulations, we found considerable bias in the CRM when using responses only from cycle 1, even though we tried to improve the method by modeling cycle 1 DLT rates and then extrapolating to m cycles. In practice, phase I trials run with the CRM may simply equate the cycle 1 DLT rate with the full m-cycle rate, leading potentially to even greater bias because many DLTs do not occur until later cycles [26]. Even so, our model, using data collected through the CRM, can significantly improve on the CRM in estimating the MTD.

We suspect that our model would achieve further gains if the CRM had been run with a randomization step in the dose-finding algorithm. That is, instead of always moving towards the current estimate of the MTD, the algorithm would allow a possibility of staying at the current dose or moving in the "wrong" direction, thereby promoting more exploration among the doses. Given that the CRM is biased in our context but still often selects an MTD near the truth, randomization may increase the number of patients treated near the true MTD and thus provide our model with better data when re-estimating the MTD.

Initially, we intended to use our model during the adaptive enrollment phase of a trial, guiding dose selection for new patients just as the CRM does. However, we found that the LTP model has too many parameters, even assuming α and λ

are known, to precisely estimate DLT rates during the trial, when limited data are available. Thus, we have proposed using our model to retrospectively analyze data from completed trials, and in this context we found our model to work well.

In our analysis, we sought to identify the MTD, a single dose level administered over multiple cycles. However, our model can also estimate the DLT rate for any dose algorithm, which specifies how a patient's dose level should change over the cycles in response to the patient's outcomes. For example, the algorithm might specify starting at a high dose, and then tapering down if a patient has an LLT. Our LTP model could help identify which dose level to start at and by how many levels to jump down after an LLT. Because of this flexibility, our model could theoretically be used adaptively during a trial to determine if dose changes for enrolled patients could be made to focus their risk of toxicity as close to the target DLT rate as possible. However, the reality of such an approach would likely be limited by the actual amount of data available as patients are accrued.

Finally, we note that carrying out the estimation for our model required overcoming significant programming challenges. We first sought to use the JAGS software [50] to implement Metropolis-Hastings, but it did not allow us to model the observed trinary outcomes Y_{ic} as the trichotomized maximum of a Gaussian process. Although workarounds may have been possible, we ultimately programmed the MH sampler ourselves in R and C++ with the `Rcpp` package [51], taking pains to efficiently calculate the quantities derived in Section 3.4. Using R alone, which we tried, proved unworkably slow, but with a C++ back-end we find that the code is efficient enough for practical use.

CHAPTER 4

Phase I Clinical Trial Method for Flexible Toxicity and Efficacy Curves

4.1 Introduction

Classic phase I methodology, such as the 3+3 design [2] and continual reassessment method (CRM) [4], assume that the probabilities of toxicity (T) and efficacy (E) increase strictly with dose. These assumptions are generally accepted for cytotoxic agents and imply that, among doses with acceptable T, the highest dose has the highest probability of E. Thus the trials may select the optimal dose based solely on T. However, for the newer class of molecularly-targeted agents (MTA), the T and E curves may plateau (increase and then level off) or peak (increase and then decrease) [27, 28], as depicted in Figure 4.1. As a result, a phase I method for an MTA should (1) incorporate both T and E, since the maximum tolerated dose is not necessarily the most likely to be efficacious; and (2) allow the curves to vary flexibly with dose.

Previous authors have developed methods to model both T and E for a single agent, including Thall and Russell [19], Gooley et al. [20], Braun [21], Thall and Cook [22], Wages and Tait [23], Thall and Nguyen [24], and Li et al. [25]. All these methods are novel and have appealing properties, but we restrict our attention to two methods that are representative of the literature: Thall and Cook, and Li et al.

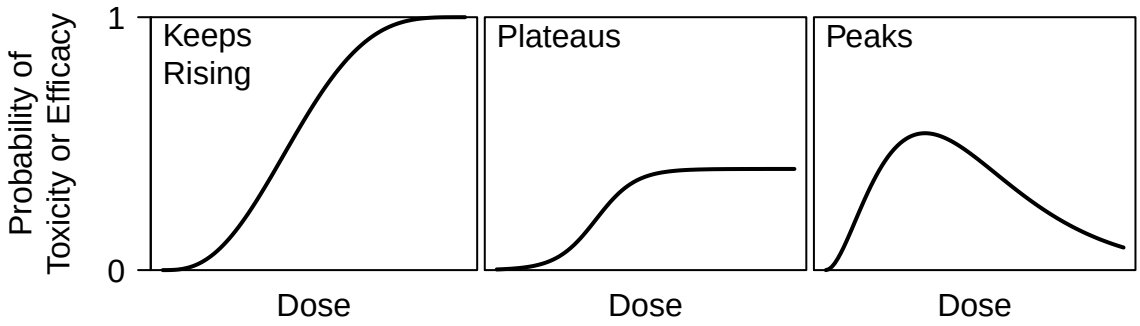

Figure 4.1. Three possible dose-response curves, for either toxicity or efficacy.

The "EffTox" method of Thall and Cook uses a parametric model in which the log-odds of T varies linearly with dose, while the log-odds of E varies quadratically. The quadratic form can capture both plateaus and peaks. An additional parameter captures the association between T and E. To rank the doses, EffTox uses a utility metric which reduces the T and E rates to a single-dimensional number, essentially measuring the distance from a point (π^E, π^T), where π^E and π^T are probabilities of E and T, to the optimal point $(1, 0)$ where there is always E and never T.

The "TEPI" (toxicity and efficacy probability interval) method of Li et al. is an interval-based method, in which the $[0, 1] \times [0, 1]$ grid, with π^E on the abscissa and π^T on the ordinate, is partitioned into rectangular subregions. Before the trial starts, a clinician assigns a rule to each region: either de-escalate, stay, or escalate the dose for the next cohort. A simple statistical model, in which the probabilities of E and T are all independent across doses, is used to estimate the rates for the current dose. These estimates will place the current dose in a subregion of the grid, thus determining the next dose to select. At the end of the trial, a utility metric is used to select the optimal dose.

Both the EffTox and TEPI methods are appealing and effective, but they have some drawbacks. In particular, the parametric forms of EffTox may be too restrictive, correlating response rates too strongly across doses, while the completely independent rates in TEPI may not be restrictive enough. To address these issues, we propose

a method in which we borrow information across doses without imposing a strict parametric form. Separately for the two outcomes T and E, each dose level has its own parameter for the probability of the outcome, and these probabilities for all doses are linked via a covariance matrix that borrows information across doses. In particular, we transform the probabilities to be unbounded, and then link them with a conditional autoregressive (CAR) model, which has been used in geospatial analysis of lattice data [29–31]. In essence, our model treats a set of candidate dose levels as forming a path graph, with connections between neighboring dose levels. This CAR model is applied separately for both T and E data. To select an optimal dose, we adapt the utility-based algorithm of EffTox, but we introduce an element of randomness to encourage exploration among the doses.

Below, we introduce our method, consisting of statistical model and dose-finding algorithm, then we evaluate its characteristics in a simulation study, and we conclude with a discussion.

4.2 Model and Methods

4.2.1 Defining the Model

Consider J candidate dose levels, ordered by increasing dose from 1 to J. With n_j patients given dose j, we collect E and T counts $y_j^E, y_j^T \in \{0, 1, \ldots, n_j\}$, where, e.g., y_j^T is the number of patients with a dose-limiting toxicity at dose j. We write the model below, where the * superscript can be replaced by either E for efficacy or T for toxicity. The model has the same form for each outcome:

$$\begin{aligned} y_j^*|\pi_j^* &\sim \text{Bin}\big(n_j, \pi_j^*\big), j = 1, \ldots, J \\ g(\boldsymbol{\pi}^*) &\sim \mathcal{N}\big(g(\boldsymbol{\pi}_0^*), \boldsymbol{\Sigma}^* = \sigma^{*2}(\mathbf{I}_J - \lambda^*\mathbf{W})^{-1}\mathbf{T}\big) \\ \boldsymbol{\pi}^* &= [\pi_1^* \ \ldots \ \pi_J^*], \quad \boldsymbol{\pi}_0^* = [\pi_{01}^* \ \ldots \ \pi_{0J}^*] = \text{skeleton}. \end{aligned} \tag{4.1}$$

Specifically, y_j^* is binomial with probability parameter π_j^*. The transformed probabilities $g(\boldsymbol{\pi}^*) = [g(\pi_1^*) \ldots g(\pi_J^*)]$ are given a CAR covariance structure, where $g : (0,1) \to \mathbb{R}$ is a known link function, e.g., the logit function. The function g should be monotonic increasing, in which case the prior median of the event probabilities equal the skeleton, i.e. median$(\pi_j^*) = \pi_{0j}^*$. The $J \times J$ matrix $\mathbf{W}$ is a weighted adjacency matrix for our dose levels. To define $\mathbf{W}$, we first define the regular $J \times J$ adjacency matrix $\mathbf{W}^\dagger$ with ijth element $w_{ij}^\dagger$:

$$w_{ij}^\dagger = \begin{cases} 1 & \text{if } |i-j| = 1 \\ 0 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Thus $\mathbf{W}^\dagger$ has ones just off the main diagonal, and zeros everywhere else, representing the adjacency structure of a path graph, as illustrated in Figure 4.2. The matrix $\mathbf{W}$ is defined as the row-normalized (that is, each row sums to one) version of $\mathbf{W}^\dagger$. In notation, the ijth element of $\mathbf{W}$ is $w_{ij} = w_{ij}^\dagger / w_{i+}^\dagger$. The diagonal matrix $\mathbf{T} = \text{diag}(1/w_{1+}^\dagger, \ldots, 1/w_{J+}^\dagger)$ scales the variance of each $g(\pi_j^*)$ according to how many neighbors dose j has (either 1 or 2), and furthermore $\mathbf{T}$ is needed to ensure that the covariance matrix $\boldsymbol{\Sigma}^* = \text{cov}(g(\boldsymbol{\pi}^*))$ is positive definite. The matrix $\boldsymbol{\Sigma}^*$ is a function of two tuning parameters: σ^* controls the prior standard deviation of each $g(\pi_j^*)$ and thus of each π_j^*; and λ^* controls the amount of information borrowed across neighboring doses. The ijth element of $\boldsymbol{\Sigma}^*$, which may be useful for model calibration, is

$$\begin{aligned} \boldsymbol{\Sigma}_{ij}^* =& (\sigma^*)^2 (\lambda^*)^{|i-j|} \\ & \times \frac{\left[(1+x)^{\min(i,j)-1} + (1-x)^{\min(i,j)-1}\right]\left[(1+x)^{J-\max(i,j)} + (1-x)^{J-\max(i,j)}\right]}{2x[(1+x)^{J-1} - (1-x)^{J-1}]} \end{aligned} \tag{4.2}$$

where $x = \sqrt{1-(\lambda^*)^2}$. See the Appendix for a derivation of this formula.

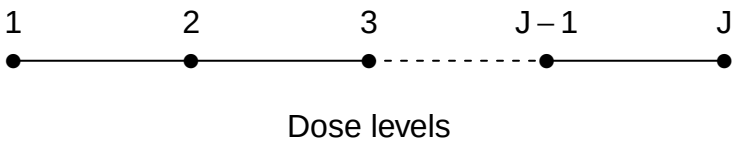

Figure 4.2. Path graph structure for J dose levels. The dose levels are ordered by increasing dosage. Dots represent doses. Lines connect adjacent/neighboring doses.

There is an additional requirement to ensure that $\boldsymbol{\Sigma}^*$ is positive definite: we need λ^* to satisfy both $\lambda^* w_{\min} < 1$ and $\lambda^* w_{\max} < 1$, where $w_{\min}$ and $w_{\max}$ are the minimum and maximum eigenvalues of $\mathbf{W}$ [31]. Given our construction of $\mathbf{W}$ (alternatives are possible for a CAR model; for example we did not have to normalize the row sums), we have $w_{\min} = -1$ and $w_{\max} = 1$ for any $J \geq 2$ (see the Appendix for a proof of this claim). Thus we require $\lambda^* \in (-1, 1)$. However, noting from Equation (4.2) that negative values of λ^* correspond to negative correlations among all first-order neighbors $g(\pi_j^*)$ and $g(\pi_{j+1}^*)$, and positive values correspond to positive correlations, we further restrict the parameter space to $\lambda^* \in (0, 1)$.

Our model treats E and T as independent outcomes, and no data are shared between them. While we could introduce a new association parameter between the outcomes, we refrain based on the recommendation of Cunanan and Koopmeiners [52]. They investigated various copula models to link E and T and compared with an independent model, finding that the correlation parameters are difficult to estimate in copula models given the small sample sizes in phase I trials, and further that the independent model performs just as well.

As a final comment about the construction of our model, we note that it is called a CAR model because the multivariate expression $g(\boldsymbol{\pi}^*) \sim \mathcal{N}(g(\boldsymbol{\pi}_0^*), \sigma^{*2}(\mathbf{I}_J - \lambda^* \mathbf{W})^{-1}\mathbf{T})$ can be equivalently written [31] as the conditional univariate expression

$$g(\pi_j^*) | g(\boldsymbol{\pi}_{(j)}^*) \sim \mathcal{N}\left(g(\pi_{0j}^*) + \lambda^* \sum_{k=1}^{J} w_{jk}[g(\pi_k^*) - g(\pi_{0k}^*)],\ \sigma^{*2}/w_{j+}^{\dagger} \right) \tag{4.3}$$

for $j = 1, \ldots, J$, where $\boldsymbol{\pi}_{(j)}^* = [\pi_1^* \ \ldots \ \pi_{j-1}^* \ \pi_{j+1}^* \ \ldots \ \pi_J^*]$. In this form, we see that

each $g(\pi_j^*)$, given its first-order neighbors, is conditionally independent of the other $g(\pi_k^*)$. This is because $w_{jk} = 0$ in (4.3) whenever dose k is not a first-order neighbor of dose j, so the sum on the right-hand side has only one or two non-zero terms. This conditional form of the model also more clearly shows the roles λ^* and σ^* play, and in particular we see precisely how λ^* controls borrowing from neighboring doses. Lastly, we see the consequence of choosing to use a row-normalized $\mathbf{W}$: it turns the sum in (4.3) into an average, since if dose j has m first-order neighbors, then w_{jk} will be $1/m$ for those neighbors and zero for all other doses.

4.2.2 Estimation of Model Parameters

In this section, for convenience we drop the superscript * denoting either toxicity or efficacy. The estimation procedure we describe will be carried out twice, once for each outcome. To estimate the parameters $\boldsymbol{\pi} = [\pi_1 \ \ldots \ \pi_J]$, we use the Metropolis-Hastings (MH) algorithm to generate a sample from the posterior distribution $p(g(\boldsymbol{\pi})|\mathbf{y})$, and then transform it via g^{-1} into a sample from $p(\boldsymbol{\pi}|\mathbf{y})$. First, let $\eta_j = g(\pi_j)$ and $\boldsymbol{\eta} = [\eta_1 \ \ldots \ \eta_J]$. Then

$$\begin{aligned} p(\eta_j|\mathbf{y}, \boldsymbol{\eta}_{(j)}) &\propto p(\mathbf{y}, \eta_j, \boldsymbol{\eta}_{(j)}) \\ &\propto p(\mathbf{y}|\eta_j, \boldsymbol{\eta}_{(j)}) p(\eta_j|\boldsymbol{\eta}_{(j)}) \\ &\propto p(y_j|\eta_j) p(\eta_j|\boldsymbol{\eta}_{(j)}). \end{aligned} \tag{4.4}$$

Thus to generate a sample from $p(\boldsymbol{\eta}|\mathbf{y})$ with the MH algorithm, we must calculate $p(y_j|\eta_j)p(\eta_j|\boldsymbol{\eta}_{(j)}), j = 1, \ldots, J$. This is straightforward, since $p(y_j|\eta_j)$ is binomial from (4.1) and $p(\eta_j|\boldsymbol{\eta}_{(j)})$ is univariate normal from (4.3). We implement the MH algorithm using the JAGS software [50]. With JAGS, the user must specify three sample sizes: one for adaptation (during which JAGS adjusts proposal distributions to achieve optimal mixing), one for burn-in, and finally one from the posterior distribution. In our experience, it is more than sufficient to specify these sample sizes as 1,000, 5,000, and 20,000, respectively, and generating this chain takes less than two seconds.

4.2.3 Estimation of DLT Rates

The MH algorithm provides us with a sample of size S from $p(\boldsymbol{\eta}^*|\mathbf{y})$, where $\boldsymbol{\eta}^* = g(\boldsymbol{\pi}^*)$. We denote this sample as $\boldsymbol{\eta}^{*(1)}, \ldots, \boldsymbol{\eta}^{*(S)}$. Our posterior estimates of the event rates are then $\hat{\boldsymbol{\pi}}^* = \hat{\mathrm{E}}[\boldsymbol{\pi}^*|\mathbf{y}] = \hat{\mathrm{E}}[g^{-1}(\boldsymbol{\eta}^*)|\mathbf{y}] = \sum_{i=1}^{S} g^{-1}(\boldsymbol{\eta}^{*(i)})/S$. This estimator can be interpreted as the posterior probability of the event (toxicity or efficacy) for a new patient assigned to dose j. That is, $\Pr(y_j^{\text{new}} = 1|\mathbf{y}) = \mathrm{E}[\Pr(y_j^{\text{new}} = 1|\pi_j^*, \mathbf{y})|\mathbf{y}] = \mathrm{E}[\Pr(y_j^{\text{new}} = 1|\pi_j^*)|\mathbf{y}] = \mathrm{E}[\pi_j^*|\mathbf{y}]$, where y_j^{new} is the binary outcome for the new patient. With our sample, we can also easily calculate posterior credible intervals for the π_j^* and other posterior quantities of interest.

4.2.4 Dose-Finding Algorithm

Running a phase I trial requires an algorithm for selecting which dose to assign to the next patient. The algorithm will proceed by iteratively moving toward the "best" dose and updating our estimates of $\boldsymbol{\pi}^E$ and $\boldsymbol{\pi}^T$. Since we have two outcomes, we need a metric that combines efficacy and toxicity rates into a one-dimensional number for each dose, allowing us to rank the pairs $\hat{\boldsymbol{\pi}}_j = (\hat{\pi}_j^E, \hat{\pi}_j^T)$. We adopt the metric developed by Thall and Cook [22], called desirability and denoted $\delta(\boldsymbol{\pi})$ for $\boldsymbol{\pi} \in [0,1]^2$.

To calculate desirability, first pick three points, $\boldsymbol{\pi}_1^\dagger, \boldsymbol{\pi}_2^\dagger, \boldsymbol{\pi}_3^\dagger \in [0,1]^2$ that are equally desirable. Fit a parametric curve of the form $\pi^T = f(\pi^E) = a + b/\pi^E + c/(\pi^E)^2$ through these three points and call this contour $\mathcal{C}$. For any point $\boldsymbol{\pi} \in [0,1]^2$, we can draw a line that connects it to the point $(1,0)$ and that intersects $\mathcal{C}$; denote this point of intersection $\mathbf{q}$. Note that the point $(1,0)$ has optimal desirability, since the probabilities of efficacy and toxicity are 1 and 0, respectively. The desirability of the point $\boldsymbol{\pi}$ is then

$$\delta(\boldsymbol{\pi}) = \frac{||\mathbf{q} - (1,0)||}{||\boldsymbol{\pi} - (1,0)||} - 1 \tag{4.5}$$

where $||\cdot||$ denotes Euclidean distance, so that higher values of δ are better. The

process of calculating $\delta(\boldsymbol{\pi})$ is depicted in Figure 4.3.

With this desirability metric, we propose the following algorithm, adapted from Thall and Cook with a slight modification. First, we define a set of acceptable doses. Given data $\mathcal{D}$, a dose j is acceptable if both

$$\Pr(\pi_j^E > \underline{\pi}^E | \mathcal{D}) > p^E \quad \text{and}$$
$$\Pr(\pi_j^T < \bar{\pi}^T | \mathcal{D}) > p^T$$

where $\underline{\pi}^E$ and $\bar{\pi}^T$ are lower and upper bounds, respectively, for efficacy and toxicity,

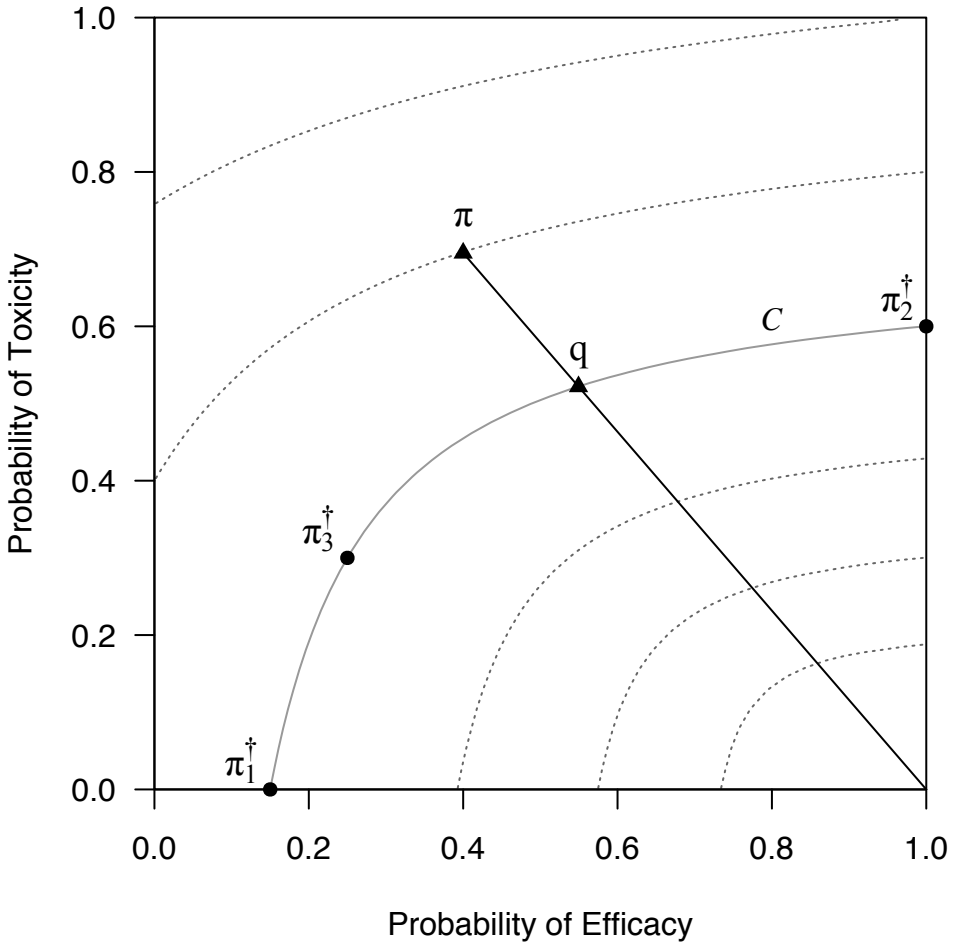

Figure 4.3. Efficacy-toxicity trade-off contours. All points on a contour have the same desirability. The reference contour $\mathcal{C}$ is formed by connecting the clinician-chosen points $\boldsymbol{\pi}_1^\dagger, \boldsymbol{\pi}_2^\dagger, \boldsymbol{\pi}_3^\dagger$. The desirability of a point $\boldsymbol{\pi}$ is inversely proportional to its distance from the point $(1, 0)$, according to Equation (4.5). This figure is based on a similar one in Thall and Cook [22].

and p^E and p^T are probability cutoffs. These four quantities are fixed at the beginning of the trial. If no dose meets these acceptability criteria, then the trial is stopped early. Otherwise, calculate $\delta(\hat{\boldsymbol{\pi}}_j)$ for the current dose j and its first-order neighbors. We then randomly select the next dose as one of these two or three dose levels, where the probability of picking dose k is proportional to $\exp(\delta(\hat{\boldsymbol{\pi}}_k))$. The selected dose is assigned to a new cohort of m patients, where m is typically 1 or 3. When the trial has finished following the desired number of patients N, the optimal dose is the one with the highest estimated desirability.

4.2.5 Calibrating the Method

In addition to choosing the J doses, the number of patients N, and the cohort size m, our method requires setting the following parameters: the probability skeletons π_{0j}^E and π_{0j}^T, $j = 1, \ldots, J$; the prior standard deviations σ^E and σ^T and the prior correlations λ^E and λ^T; the minimum acceptable E rate $\underline{\pi}^E$ and the maximum acceptable T rate $\bar{\pi}^T$; the probability cutoffs p^E and p^T; and the equal-desirability points $\boldsymbol{\pi}_1^\dagger, \boldsymbol{\pi}_2^\dagger, \boldsymbol{\pi}_3^\dagger$. The skeletons, minimum/maximum acceptable rates, and equal-desirability points may be chosen by a clinician. The parameters σ^E, σ^T, λ^E, λ^T, p^E, and p^T should then be set by computer simulation. To reduce the dimension of the search space, we recommend fixing $\sigma^E = \sigma^T$, $\lambda^E = \lambda^T$, and $p^E = p^T$.

4.3 Simulation Study

We conduct a simulation study, simulating adaptive trials and evaluating the performance of our CAR method. We also compare our method with three others: the CRM using only toxicity data, the EffTox method of Thall and Cook, and the TEPI method of Li et al. The true model that we use for simulating data is very simple. For each dose level there is a true T and E rate $\pi_j^{T,\text{true}}$ and $\pi_j^{E,\text{true}}$, and we simulate

an individual patient's outcomes from Bern($\pi_j^{T,\text{true}}$) and Bern($\pi_j^{E,\text{true}}$). Note that, according to our simulation scheme, a patient's E and T outcomes are not correlated.

4.3.1 Simulation Scenarios

We adopt the simulation scenarios from Thall and Cook (referred to as TC) and Li et al. (referred to as Li).

In TC, there are $J = 5$ doses, $N = 72$ patients enrolled in cohorts of size m, and the true E and T rates always increase with dose across six scenarios. Our method is not optimized for strictly increasing E and T rates, but we are nonetheless interested in how our method performs in this context. As we are borrowing the scenarios from TC, we do not need to implement their EffTox method, and we simply re-report their results with $m = 3$. Additionally, since our dose-finding algorithm is adopted from EffTox, we can also use their tuning parameter values: $\underline{\pi}^E = 0.5$, $\bar{\pi}^T = 0.1$, $p^E = p^T = 0.1$, $\boldsymbol{\pi}_1^\dagger = (0.45, 0)$, $\boldsymbol{\pi}_2^\dagger = (0.55, 0.1)$, $\boldsymbol{\pi}_3^\dagger = (0.84, 0.16)$. We use cohorts of both $m = 3$ and $m = 1$ for the CAR method. For the skeletons, we set π_{0j}^* to the average of $\pi_j^{*,\text{true}}$ across the six TC scenarios, giving $\boldsymbol{\pi}^E = [0.26, 0.40, 0.52, 0.62, 0.71]$ and $\boldsymbol{\pi}^T = [0.05, 0.08, 0.11, 0.15, 0.21]$. We calibrated the prior standard deviation and correlation parameters via simulation to be $\sigma^E = \sigma^T = 0.7$ and $\lambda^E = \lambda^T = 0.5$.

In Li, $J = 4$, $N = 27$, $m = 3$, and the T rate always increases while the E rate may increase, plateau, or peak in six scenarios. In scenario 3, we note that the E rate has a dramatic peak, increasing from 0.1 to 0.7, then falling to 0.2 across the first three dose levels. The Li paper compares its TEPI method with both EffTox and the CRM, so we re-report these results and add in those from our CAR method. Since Li also implemented EffTox, we use their tuning parameter values: $\underline{\pi}^E = 0.4$, $\bar{\pi}^T = 0.4$, $p^E = 0.3$, $p^T = 0.05$, $\boldsymbol{\pi}_1^\dagger = (0.2, 0)$, $\boldsymbol{\pi}_2^\dagger = (1, 0.6)$, $\boldsymbol{\pi}_3^\dagger = (0.5, 0.5)$. We set the skeletons to $\boldsymbol{\pi}^E = [0.4, 0.4, 0.4, 0.4]$ and $\boldsymbol{\pi}^T = [0.1, 0.2, 0.3, 0.4]$. We calibrated the remaining parameters to be $\sigma^E = \sigma^T = 0.7$ and $\lambda^E = \lambda^T = 0.5$.

In each scenario, we simulate 1,000 adaptive trials with our CAR method. When looking at the TC scenarios, in addition to evaluating our method with both $m = 3$ and $m = 1$, we also vary the dose-finding algorithm. That is, in one version we randomize the next dose selection based on estimated desirability as described previously, while in another version we remove the randomization, always picking the dose with highest estimated desirability, subject to a constraint preventing the algorithm from moving by more than one dose. However, we only calibrate the model parameters once, with $m = 3$ and using randomization.

4.3.2 Simulation Results

The comparison of CAR with EffTox in the TC scenarios is in Table 4.1. The comparison with TEPI, CRM, and EffTox from Li is in Table 4.2. Focusing first on the TC scenarios, we note a general trend that the CAR method tends to prefer lower doses than EffTox. In scenarios 1 and 4, where the highest dose is optimal, EffTox selects the right dose more often than CAR. In scenarios 2 and 3, where the optimal dose is not the highest, CAR with randomization selects the optimal dose more often than EffTox, and further EffTox too often selects doses above the optimal which leads to excessive risk of toxicity. We also observe, comparing the different versions of CAR, that adding randomness to the dose-finding algorithm can lead to big gains in selecting the true optimal dose, while reducing the cohort size from 3 to 1 leads to modest gains. We call attention in particular to scenario 1, where the gains from randomization are quite large. Finally, in scenarios 5 and 6, all the doses are undesirable, either too toxic or not sufficiently efficacious, and both EffTox and CAR do similarly well at stopping the trial early and not recommending any dose for further testing.

In the Li scenarios in Table 4.2, CAR typically performs around as well, and sometimes better than TEPI; and mostly better than both EffTox and CRM. The

Table 4.1. Simulation results comparing CAR method with EffTox, and using Thall and Cook's scenarios. In these scenarios, the true rates of E and T always increase with dose. The CAR method uses either 24 cohorts of 3, or 72 cohorts of 1. The dose-finding algorithm is either deterministic or random with selection probabilities proportional to exp(estimated desirability). In the 4 CAR columns, the label in parentheses indicates the cohort size (1 or 3) and the type of dose-finding algorithm (D for deterministic or R for random). For EffTox, there are 24 cohorts of 3 patients, and the dose-finding algorithm is not randomized. In each scenario, the highlighted row shows the desired outcome, which is either the optimal dose or stopping early.

	Dose	True Probability			Selection Percentage (Number of Patients Treated)				
Scenario	Level	Tox	Eff	Desir.	EffTox	CAR (3D)	CAR (3R)	CAR (1D)	CAR (1R)
1	1	0.01	0.05	-0.43	0.0 (3.0)	0.0 (3.0)	0.0 (5.5)	0.0 (1.0)	0.0 (2.2)
	2	0.02	0.20	-0.32	0.0 (4.0)	0.4 (3.3)	0.0 (7.2)	0.0 (1.6)	0.0 (5.0)
	3	0.03	0.35	-0.18	0.7 (5.1)	2.0 (12.3)	1.2 (16.1)	3.6 (9.8)	0.4 (15.6)
	4	0.04	0.60	0.27	5.8 (7.1)	94.8 (52.4)	30.0 (25.1)	92.8 (57.9)	22.0 (27.5)
	5	0.05	0.80	1.17	92.8 (52.2)	1.6 (0.7)	68.0 (17.8)	2.4 (1.0)	77.6 (21.7)
	Early Stop				0.7	1.2	0.8	1.2	0
2	**1**	0.01	0.57	0.26	0.1 (5.4)	22.0 (17.4)	31.2 (19.0)	2.8 (2.9)	22.0 (12.4)
	2	0.03	0.58	0.24	20.5 (21.6)	32.4 (23.7)	32.8 (23.4)	21.6 (15.4)	28.0 (21.8)
	3	0.06	0.60	0.22	61.9 (29.5)	44.0 (28.3)	33.6 (18.2)	71.6 (48.4)	45.6 (22.7)
	4	0.20	0.62	-0.24	16.1 (11.6)	1.2 (2.4)	2.4 (8.9)	2.8 (4.4)	2.8 (11.5)
	5	0.32	0.64	-0.50	0.9 (3.6)	0.0 (0.0)	0.0 (2.4)	0.0 (0.0)	0.4 (3.2)
	Early Stop				0.5	0.4	0	1.2	1.2
3	1	0.02	0.20	-0.32	0.0 (3.4)	0.0 (3.1)	0.0 (7.2)	0.0 (1.0)	0.0 (3.5)
	2	0.03	0.40	-0.12	1.6 (8.8)	2.0 (6.4)	2.0 (14.2)	0.4 (3.2)	0.8 (12.6)
	3	0.04	0.60	0.27	32.2 (20.8)	64.4 (43.4)	39.2 (22.4)	60.0 (42.3)	38.0 (25.0)
	4	0.06	0.68	0.46	49.4 (22.3)	33.6 (18.8)	54.8 (20.5)	38.8 (25.1)	53.2 (22.9)
	5	0.20	0.74	-0.20	15.7 (16.0)	0.0 (0.2)	4.0 (7.7)	0.0 (0.0)	7.6 (7.9)
	Early Stop				1	0	0	0.8	0.4
4	1	0.01	0.52	0.12	0.0 (3.5)	12.4 (12.4)	3.2 (14.1)	0.4 (1.3)	2.0 (7.9)
	2	0.015	0.62	0.41	0.1 (4.3)	31.6 (23.8)	12.0 (19.6)	11.2 (8.9)	7.6 (15.4)
	3	0.02	0.71	0.80	1.1 (5.3)	51.6 (33.2)	36.8 (18.7)	76.4 (53.9)	23.6 (19.8)
	4	0.025	0.79	1.39	4.6 (6.6)	4.4 (2.6)	30.8 (13.2)	12.0 (8.0)	41.6 (18.9)
	5	0.03	0.86	2.24	94.0 (52.2)	0.0 (0.0)	17.2 (6.4)	0.0 (0.0)	25.2 (10.0)
	Early Stop				0.1	0	0	0	0
5	1	0.18	0.05	-0.51	0.1 (3.4)	0.0 (3.1)	0.0 (5.1)	0.0 (1.1)	0.0 (2.1)
	2	0.22	0.20	-0.47	0.9 (8.3)	0.0 (5.3)	0.0 (7.2)	0.4 (4.4)	0.0 (4.7)
	3	0.26	0.35	-0.46	1.6 (3.6)	0.4 (12.6)	1.6 (12.9)	0.4 (10.6)	0.8 (11.0)
	4	0.30	0.47	-0.49	1.4 (0.8)	0.8 (7.6)	0.8 (9.1)	0.0 (6.2)	0.8 (8.1)
	5	0.33	0.58	-0.52	0.2 (0.3)	0.4 (2.9)	0.8 (3.2)	0.0 (2.0)	0.8 (3.2)
	Early Stop				97.3	98.4	96.8	99.2	97.6
6	1	0.08	0.15	-0.40	0.4 (5.3)	0.0 (3.6)	0.8 (6.6)	0.0 (1.6)	0.0 (3.2)
	2	0.18	0.38	-0.33	11.4 (20.1)	4.0 (14.5)	12.0 (15.6)	5.2 (11.9)	5.6 (13.1)
	3	0.25	0.52	-0.39	1.3 (4.5)	4.0 (16.7)	7.2 (18.2)	2.0 (16.1)	6.0 (15.9)
	4	0.30	0.59	-0.47	0.0 (1.1)	2.8 (8.4)	3.6 (9.5)	2.0 (6.6)	2.0 (8.0)
	5	0.33	0.62	-0.54	0.0 (0.4)	2.0 (2.0)	1.6 (3.1)	0.4 (1.6)	0.8 (2.6)
	Early Stop				86.9	87.2	74.8	90.4	85.6

Table 4.2. Simulation results comparing CAR method with TEPI, CRM, and EffTox, and using Li et al.'s scenarios. In these scenarios, the T rate always increases with dose while the E rate may increase, plateau, or peak. There are 4 dose levels, and we enroll up to 27 patients. In each scenario, the highlighted row shows the desired outcome, which is either the optimal dose or stopping early.

	Dose	True Probability			Selection Percentage (Number of Patients Treated)			
Scenario	Level	Tox	Eff	Desir.	TEPI	CRM	EffTox	CAR
1	1	0.16	0.05	-0.19	22.1 (6.0)	7.9 (7.9)	0 (3.1)	0.8 (6.2)
	2	0.2	0.1	-0.16	17.9 (5.7)	23.6 (7.5)	0 (3.1)	0.0 (6.6)
	3	0.25	0.15	-0.13	17.2 (5.1)	31.9 (6.3)	0 (3.8)	7.2 (4.3)
	4	0.3	0.18	-0.12	7.5 (4.3)	36.6 (5.2)	8 (5.3)	12.4 (4.2)
	Early Stop				35.3	0	92	79.6
2	**1**	0.15	0.8	1.96	83.9 (9.1)	6.3 (7.5)	66 (17.8)	82.4 (13.6)
	2	0.2	0.8	1.50	13.6 (8.5)	23.6 (7.6)	30 (8.4)	13.6 (8.3)
	3	0.25	0.8	1.12	2.1 (5.6)	32.7 (6.5)	3 (0.7)	3.1 (3.7)
	4	0.3	0.8	0.82	0.3 (3.8)	37.4 (5.3)	1 (0.1)	0.9 (1.5)
	Early Stop				0.1	0	0	0
3	1	0.1	0.1	-0.13	7.2 (4.4)	2.1 (5.5)	3 (3.7)	0.4 (7.0)
	2	0.2	0.7	1.08	88 (12.3)	32.1 (8.9)	42 (11.8)	88.0 (10.7)
	3	0.3	0.2	-0.10	0.3 (7.1)	60.6 (10)	3 (4.3)	0.8 (6.5)
	4	0.7	0.1	-0.35	0.1 (2.3)	5.2 (2.6)	2 (2.0)	0.4 (1.8)
	Early Stop				4.4	0	50	10.4
4	1	0.15	0.43	0.32	53.9 (6.1)	7.7 (7.7)	19 (7.3)	41.6 (10.4)
	2	0.2	0.52	0.48	41.3 (9.6)	43.8 (9.7)	49 (12.4)	44.4 (9.1)
	3	0.4	0.5	0.15	3.6 (9.0)	41.7 (7.6)	22 (5.3)	10.9 (5.4)
	4	0.5	0.6	0.06	1.2 (2.1)	6.8 (2.0)	5 (1.1)	2.7 (2.1)
	Early Stop				1.2	0	6	0.4
5	1	0.1	0.2	-0.03	16.4 (4.7)	3.1 (5.4)	1 (3.6)	2.4 (8.1)
	2	0.2	0.6	0.71	65.4 (8.6)	25.0 (8.2)	48 (12.4)	68.8 (10.0)
	3	0.3	0.6	0.48	13.8 (7.2)	45.4 (8.6)	38 (8.5)	22.3 (6.2)
	4	0.4	0.6	0.25	1.0 (4.9)	26.5 (4.8)	10 (2.0)	6.5 (2.7)
	Early Stop				3.4	0	3	0
6	1	0.5	0.4	-0.07	33.9 (14.9)	99.3 (25.8)	16 (7.9)	19.2 (9.8)
	2	0.6	0.5	-0.12	0.3 (1.8)	0.7 (1.1)	13 (6.6)	9.1 (6.0)
	3	0.7	0.6	-0.20	0.0 (0.1)	0.0 (0.1)	2 (1.2)	1.6 (0.8)
	4	0.8	0.8	-0.27	0.0 (0.0)	0.0 (0.0)	0 (0.2)	0.0 (0.6)
	Early Stop				65.8	0	69	70.1

CRM in particular does quite poorly across scenarios, as it does not make use of efficacy data. In scenarios 2-4, the true optimal dose is one of the two lowest, and we suspect that this contributes to CAR beating EffTox, as again, EffTox appears to want to select higher doses. In scenario 1, CAR beats TEPI in stopping early more often. In scenario 5, TEPI and CAR select the true optimal dose around as often, but TEPI more often selects the lowest, non-efficacious dose, while CAR more often escalates to more toxic doses even though efficacy has plateaued.

4.4 Discussion

We have proposed a method for phase I clinical trials in which we share information across doses without restricting the shapes of the dose-toxicity and dose-efficacy curves. Our method is applicable to settings in which the true curves cannot be assumed strictly increasing. Thus, while we do not recommend our CAR method for evaluation of cytotoxic chemotherapies, it may be appropriate for MTAs. The mechanism of action of MTAs is such that, beyond a threshold, further dose escalation may not increase toxicity or efficacy rates, and in some cases may lead to decreasing rates [27, 28].

In evaluating our method, we tried varying both cohort size and whether the dose-finding algorithm uses randomization in selecting a dose for the next cohort, as opposed to always picking the dose with highest estimated desirability. In our experience, clinicians often prefer cohorts of three, perhaps due to familiarity with the 3+3 method. Simulation results suggest, however, that the CAR method benefits from cohorts of one, leading to modest increases in selecting the true optimal dose. Randomization produces even larger increases, and additionally encourages greater exploration of the doses, leading to more information on toxicity and efficacy across the range of doses. However, our use of exp(estimated desirability) for the randomiza-

tion probabilities (after sum-to-one standardization) was based on convenience, and alternatives should be investigated.

As with many phase I methods, our method has a number of tuning parameters to specify before starting the trial. In addition to the skeletons and equal-desirability points, there are parameters which may be considered less intuitively meaningful and thus harder to tune. For this latter set of parameters, we find it convenient to equate corresponding parameters between the efficacy and toxicity models, thereby reducing the number of free parameters. While this likely prevents us from optimally tuning the method, it makes the tuning far more feasible, and our simulation study suggests the results are still satisfactory.

Finally, we note that the CAR method may be readily extended to study combination therapies. A first-order neighbor of a combination dose may be defined as any other combination dose reachable by changing exactly one component dose by one level, either up or down. After defining this neighborhood structure, many details will carry through from our single-therapy CAR method.

4.5 Appendix

4.5.1 Eigenvalues of the Row-Normalized Adjacency Matrix

Claim. The minimum and maximum eigenvalues of the row-normalized $J \times J$ adjacency matrix $\mathbf{W}$ are -1 and 1, respectively.

Proof. First we show that all eigenvalues of $\mathbf{W}$ must belong to the closed interval $[-1, 1]$. By the Gershgorin circle theorem [53, 54], all eigenvalues of $\mathbf{W}$ must lie in at least one closed interval of the form $[w_{jj} - w_{j+}, w_{jj} + w_{j+}]$, for $j = 1, \ldots, J$ where w_{jj} is the jth diagonal element of $\mathbf{W}$ and w_{j+} is the jth row sum of $\mathbf{W}$. In our case, $w_{jj} = 0$ and $w_{j+} = 1$ for all j, so the J intervals are all the same: $[-1, 1]$. Thus all eigenvalues lie in $[-1, 1]$.

Next we show that both -1 and 1 are indeed eigenvalues of $\mathbf{W}$, and are therefore, by Gershgorin, the minimum and maximum eigenvalues. Since the row-sums of $\mathbf{W}$ are 1, we have $\mathbf{W}\mathbf{1}_J = \mathbf{1}_J$, where $\mathbf{1}_J$ is a vector of ones. Thus 1 is an eigenvalue of $\mathbf{W}$. Our argument for -1 being an eigenvalue is more involved. Recall that we are treating the J dose levels as forming a path in the graph theoretic sense. A path is a bipartite graph [54], meaning the nodes (doses) can be partitioned into two sets such that the connections are always *between* sets, and never within sets. Figure 4.4 illustrates this for the case $J = 6$.

Given this graph structure, the nodes can be re-ordered (as numbered in Figure 4.4) so that $\mathbf{W}$ is a block matrix with 0-matrices on the diagonal,

$$\mathbf{W} = \begin{bmatrix} \mathbf{0}_{n\times n} & \mathbf{W}_1 \\ \mathbf{W}_2 & \mathbf{0}_{m\times m} \end{bmatrix},$$

where $n = \lceil J/2 \rceil$ and $m = \lfloor J/2 \rfloor$, and the two non-zero submatrices $\mathbf{W}_1$ $(n \times m)$ and $\mathbf{W}_2$ $(m \times n)$ both have have row-sums equal to 1. Note that the labeling of nodes, and thus their order in $\mathbf{W}$ is arbitrary, and in particular changing the order does not

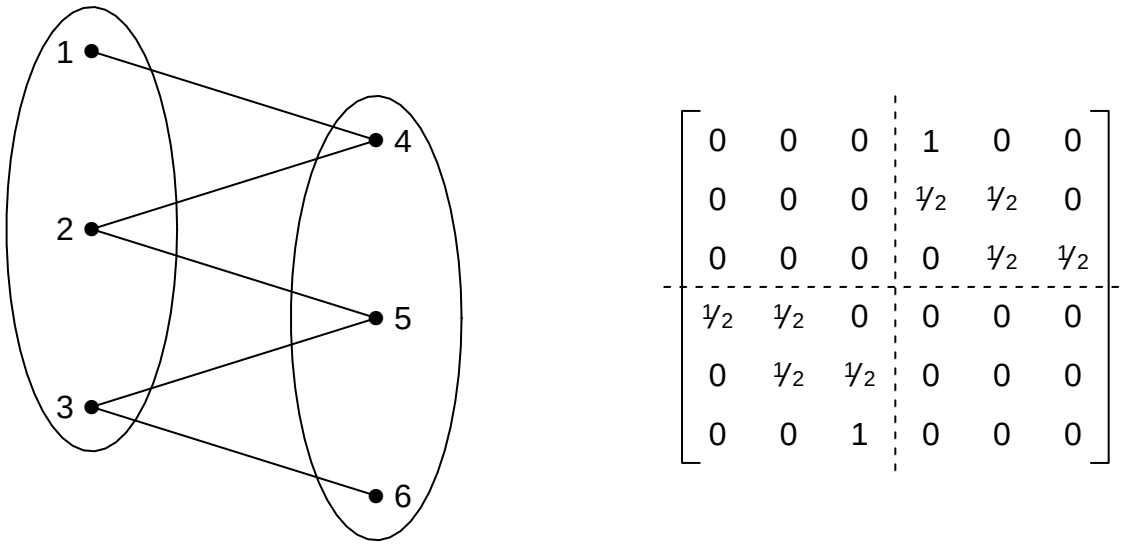

Figure 4.4. Bipartite graph structure and associated row-normalized adjacency matrix, with six dose levels. The doses have been relabeled corresponding to the bipartite structure.

change the eigenvalues. Now consider the column vector $\mathbf{v} = [\mathbf{1}_n, -\mathbf{1}_m]^T$. We have

$$\mathbf{W}\mathbf{v} = \begin{bmatrix} \mathbf{0}_{n\times n} & \mathbf{W}_1 \\ \mathbf{W}_2 & \mathbf{0}_{m\times m} \end{bmatrix} \begin{bmatrix} \mathbf{1}_n \\ -\mathbf{1}_m \end{bmatrix} = \begin{bmatrix} -\mathbf{W}_1\mathbf{1}_m \\ \mathbf{W}_2\mathbf{1}_n \end{bmatrix} = \begin{bmatrix} -\mathbf{1}_n \\ \mathbf{1}_m \end{bmatrix} = -\mathbf{v}.$$

Thus -1 is an eigenvalue of $\mathbf{W}$, and this completes the proof. □

4.5.2 Derivation of the CAR Covariance Matrix

We derive a closed-form expression for the ijth element of the covariance matrix $\mathbf{\Sigma}^* = (\sigma^*)^2(\mathbf{I}_J - \lambda^*\mathbf{W})^{-1}\mathbf{T}$. For convenience, we drop the superscript $*$ notation. First we note that $\mathbf{A} = \sigma^2\mathbf{\Sigma}^{-1} = \mathbf{T}^{-1}(\mathbf{I}_J - \lambda\mathbf{W}) = \mathbf{T}^{-1} - \lambda\mathbf{W}^\dagger$, where $\mathbf{T}^{-1} = \text{diag}(w_{1+}^\dagger, \ldots, w_{J+}^\dagger)$. Thus the matrix $\mathbf{A}$ is tridiagonal, with the form

$$\mathbf{A} = \begin{bmatrix} 1 & -\lambda & & & & & \\ -\lambda & 2 & -\lambda & & & \mathbf{0} & \\ & -\lambda & 2 & \ddots & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & \ddots & 2 & -\lambda & \\ & \mathbf{0} & & & -\lambda & 2 & -\lambda \\ & & & & & -\lambda & 1 \end{bmatrix}.$$

Using the formula of Usmani [55], the ijth element of the inverse of $\mathbf{A}$ is

$$\begin{aligned}(\mathbf{A}^{-1})_{ij} &= \begin{cases} (-1)^{i+j}(-\lambda)^{j-i}\theta_{i-1}\phi_{j+1}/\theta_J & \text{if } i \le j \\ (-1)^{i+j}(-\lambda)^{i-j}\theta_{j-1}\phi_{i+1}/\theta_J & \text{if } i > j \end{cases} \\ &= \lambda^{|i-j|}\theta_{\min(i,j)-1}\phi_{\max(i,j)+1}/\theta_J\end{aligned}$$

where θ_i and ϕ_i satisfy the recurrence relations

$$\begin{aligned}\theta_i &= (\mathbf{T}^{-1})_{ii}\theta_{i-1} - \lambda^2\theta_{i-2} \quad \text{for } i = 2, 3, \ldots, J; \\ \phi_i &= (\mathbf{T}^{-1})_{ii}\phi_{i+1} - \lambda^2\phi_{i+2} \quad \text{for } i = J-1, \ldots, 1;\end{aligned} \tag{4.6}$$

with initial conditions $\theta_0 = 1$, $\theta_1 = (\mathbf{T}^{-1})_{11} = 1$, $\phi_{J+1} = 1$, and $\phi_J = (\mathbf{T}^{-1})_{JJ} = 1$. For our matrix $\mathbf{A}$, which is symmetric and persymmetric, we see that the sequence θ_i is just the sequence ϕ_i in reverse, that is, $\phi_i = \theta_{J-i+1}$. Thus our expression for

$(\mathbf{A}^{-1})_{ij}$ simplifies to

$$(\mathbf{A}^{-1})_{ij} = \lambda^{|i-j|}\theta_{\min(i,j)-1}\theta_{J-\max(i,j)}/\theta_J. \tag{4.7}$$

Our task now is to find an explicit formula for θ_i given the recurrence relation in Equation (4.6). This relation can be broken up into two parts, noting that $(\mathbf{T}^{-1})_{ii} = 2$ for $i = 2, \ldots, J-1$ and $(\mathbf{T}^{-1})_{JJ} = 1$:

$$\theta_i = 2\theta_{i-1} - \lambda^2\theta_{i-2} \quad \text{for } i = 2, 3, \ldots, J-1, \text{ and} \tag{4.8a}$$

$$\theta_J = \theta_{J-1} - \lambda^2\theta_{J-2}. \tag{4.8b}$$

Equation (4.8a) can be solved with a generating function [56], giving $\theta_i = \frac{1}{2}[(1+x)^i + (1-x)^i]$ for $i = 2, \ldots, J-1$, where $x = \sqrt{1-\lambda^2}$. Plugging this expression for θ_i into Equation (4.8b) gives $\theta_J = \frac{1}{2}x[(1+x)^{J-1} - (1-x)^{J-1}]$. We now have a closed-form expression for θ_i for $i = 0, \ldots, J$, which we plug into Equation (4.7) for our final result:

$$\begin{aligned}\mathbf{\Sigma}_{ij} &= \sigma^2(\mathbf{A}^{-1})_{ij} \\ &= \sigma^2\lambda^{|i-j|}\frac{\left[(1+x)^{\min(i,j)-1} + (1-x)^{\min(i,j)-1}\right]\left[(1+x)^{J-\max(i,j)} + (1-x)^{J-\max(i,j)}\right]}{2x[(1+x)^{J-1} - (1-x)^{J-1}]}.\end{aligned}$$

CHAPTER 5

Conclusion

In this dissertation, we identified three settings in which commonly available phase I data could be better leveraged to estimate an optimal dose for a single agent. In Chapter 2, for a cytotoxic agent for which one dose administration is sufficient for determining the DLT rate, we proposed counting multiple DLTs and LLTs per patient, and we found that our method often improved the estimation of the probability of having at least one DLT, relative to the CRM. Other methods accounting for multiple events have been developed [11–14], but ours does so in a way which does not require the redefinition of the maximum tolerated dose to work with a new endpoint. Our method could be extended to handle partial patient follow-up in a manner similar to the time-to-event CRM [5], allowing a new patient to be enrolled before the previous patient has reached the end of the observation period. This could speed up a trial and reduce operational costs, but care must be taken to not lose the benefits of sequential adaptation. Future research may explore acceptable wait-times between patients, and how to down-weight patients with partial follow-up.

In Chapter 3, for a cytotoxic agent administered multiple times, and for which toxicities may be delayed and cumulative dose is important, we developed a model inspired by pharmacokinetics to estimate the probability of DLT over all cycles of administration. Our model builds on previous work [15–18] by incorporating four elements we consider important: allowing for a cumulative dose effect, allowing the

dose to change between cycles, accounting for within-patient correlation of responses across cycles, and allowing the occurrence of LLT in one cycle to alter the probability of DLT in a subsequent cycle. We encountered challenges developing a model that includes all of these desired features and remains estimable from datasets of limited size. We determined that some parameters needed to be fixed, rather than estimated, by eliciting values from experts. As a final concession, we presented our model for use in retrospectively analyzing data from a completed trial, rather than prospectively guiding a trial through its adaptive dose-finding phase. As a retrospective method, however, we found promising results: if a trial is run with the CRM using only cycle 1 responses, but data from all cycles are collected, then our model can reanalyze the data at the end of the trial and better estimate the MTD.

In Chapter 4, for a molecularly targeted agent, for which both toxicity and efficacy responses must be accounted for in selecting an optimal dose, we adopted the conditional autoregressive model to estimate response rates. Previously developed methods have incorporated toxicity and efficacy [19–25], but the contribution of our method is to borrow response information across dose levels without imposing any functional form on the dose-response curves. Simulations showed that our method could adapt to a variety of dose-toxicity and dose-efficacy patterns, and often performed at least as well as competing methods. As we described above for the methods of Chapter 2, our CAR method could similarly be extended to handle partial follow-up. Such an extension may be particularly compelling when modeling both toxicity and efficacy, as the timescale for observing the two responses may greatly differ, with efficacy often taking longer to determine. If efficacy, for example, takes half a year to assess, then requiring full follow-up for 30 patients would take 15 years, which is likely unacceptable, particularly if toxicity can be assessed within only one month. Allowing partial follow-up could make such a trial feasible.

We also note that our CAR method can be naturally extended to study com-

bination therapies. With k agents, the neighborhood graph for the CAR method would be a k-dimensional lattice, and two combination dose levels would be first-order neighbors if exactly one of the k agents differ by one level. With this neighborhood structure, the associated adjacency matrix can be easily defined and many details of our single-agent method carry through. Further research may explore the utility of this method in both phase I and phase II trials.

In all of the work presented in this dissertation, we found untapped information in data that are already routinely collected, and we developed methods to exploit this information to improve the conduct of phase I trials. Although many phase I trials still use the 3+3 design or a variant for both cytotoxic agents and MTAs [1, 32], it is incumbent upon statisticians to continue developing and arguing for feasible, more efficient alternatives. The safety and well-being of patients depends on it.

BIBLIOGRAPHY

[1] Le Tourneau C., Lee J. J., and Siu L. L. Dose escalation methods in phase I cancer clinical trials. *Journal of the National Cancer Institute*, 101(10):708–720, 2009.

[2] Storer B. E. Design and analysis of phase I clinical trials. *Biometrics*, 45(3):925–937, 1989.

[3] Dixon W. J. and Mood A. M. A method for obtaining and analyzing sensitivity data. *Journal of the American Statistical Association*, 43(241):109–126, 1948.

[4] O'Quigley J., Pepe M., and Fisher L. Continual reassessment method: A practical design for phase 1 clinical trials in cancer. *Biometrics*, 46(1):33–48, 1990.

[5] Cheung Y. K. and Chappell R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics*, 56(4):1177–1182, 2000.

[6] Faries D. Practical modifications of the continual reassessment method for phase I cancer clinical trials. *Journal of Biopharmaceutical Statistics*, 4(2):147–164, 1994.

[7] Goodman S. N., Zahurak M. L., and Piantadosi S. Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine*, 14(11):1149–1161, 1995.

[8] O'Quigley J. and Shen L. Z. Continual reassessment method: A likelihood approach. *Biometrics*, 52(2):673–684, 1996.

[9] Ivanova A. Escalation, group and A + B designs for dose-finding trials. *Statistics in Medicine*, 25(21):3668–3678, 2006.

[10] Paoletti X. and Kramar A. A comparison of model choices for the continual reassessment method in phase I cancer trials. *Statistics in Medicine*, 28(24):3012–3028, 2009.

[11] Bekele B. N. and Thall P. F. Dose-finding based on multiple toxicities in a soft tissue sarcoma trial. *Journal of the American Statistical Association*, 99(1):26–35, 2004.

[12] Yuan Z., Chappell R., and Bailey H. The continual reassessment method for multiple toxicity grades: A Bayesian quasi-likelihood approach. *Biometrics*, 63(1):173–179, 2007.

[13] Chen Z., Krailo M. D., Azen S. P., and Tighiouart M. A novel toxicity scoring system treating toxicity response as a quasi-continuous variable in phase I clinical trials. *Contemporary Clinical Trials*, 31(5):473–482, 2010.

[14] Van Meter E. M., Garrett-Mayer E., and Bandyopadhyay D. Proportional odds model for dose-finding clinical trial designs with ordinal toxicity grading. *Statistics in Medicine*, 30(17):2070–2080, 2011.

[15] Legedza A. T. and Ibrahim J. G. Longitudinal design for phase I clinical trials using the continual reassessment method. *Controlled Clinical Trials*, 21(6):574–588, 2000.

[16] Doussau A., Thibaut R., and Paoletti X. Dose-finding design using mixed-effect proportional odds model for longitudinal graded toxicity data in phase I oncology clinical trials. *Statistics in Medicine*, 32(30):5430–5447, 2013.

[17] Zhang J. and Braun T. M. A phase I Bayesian adaptive design to simultaneously optimize dose and schedule assignments both between and within patients. *Journal of the American Statistical Association*, 108(503), 2013.

[18] Fernandes L. L., Taylor J. M. G., and Murray S. Adaptive phase I clinical trial design using Markov models for conditional probability of toxicity. *Journal of Biopharmaceutical Statistics*, 26(3):475–498, 2016.

[19] Thall P. F. and Russell K. E. A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics*, 54(1):251–264, 1998.

[20] Gooley T. A., Martin P. J., Fisher L. D., and Pettinger M. Simulation as a design tool for phase I/II clinical trials: An example from bone marrow transplantation. *Controlled Clinical Trials*, 15(6):450–462, 1994.

[21] Braun T. M. The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials*, 23(3):240–256, 2002.

[22] Thall P. F. and Cook J. D. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*, 60(3):684–693, 2004.

[23] Wages N. A. and Tait C. Seamless phase I/II adaptive design for oncology trials of molecularly targeted agents. *Journal of Biopharmaceutical Statistics*, 25(5):903–920, 2015.

[24] Thall P. F. and Nguyen H. Q. Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *Journal of Biopharmaceutical Statistics*, 22(4):785–801, 2012.

[25] Li D. H., Whitmore J. B., Guo W., and Ji Y. Toxicity and efficacy probability interval design for phase I adoptive cell therapy dose-finding clinical trials. *Clinical Cancer Research*, 2016.

[26] Postel-Vinay S., Gomez-Roca C., Molife L. R., Anghan B., Levy A., Judson I., De Bono J., Soria J.-C., Kaye S., and Paoletti X. Phase I trials of molecularly targeted agents: Should we pay more attention to late toxicities? *Journal of Clinical Oncology*, 29(13):1728–1735, 2011.

[27] Postel-Vinay S., Arkenau H.-T., Olmos D., Ang J., Barriuso J., Ashley S., Banerji U., De-Bono J., Judson I., and Kaye S. Clinical benefit in phase-I trials of novel molecularly targeted agents: does dose matter? *British Journal of Cancer*, 100(9):1373–1378, 2009.

[28] Jain R. K., Lee J. J., Hong D., Markman M., Gong J., Naing A., Wheler J., and Kurzrock R. Phase I oncology studies: evidence that in the era of targeted therapies patients on lower doses do not fare worse. *Clinical Cancer Research: An Official Journal of the American Association for Cancer Research*, 16(4):1289–1297, 2010.

[29] Besag J. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2):192–236, 1974.

[30] Cressie N. A. C. *Statistics for Spatial Data*, part II: Lattice Data, pages 381–573. John Wiley & Sons, Inc., 2015.

[31] Wall M. M. A close look at the spatial structure implied by the CAR and SAR models. *Journal of Statistical Planning and Inference*, 121:311–324, 2004.

[32] Hansen A. R., Graham D. M., Pond G. R., and Siu L. L. Phase 1 trial design: Is 3 + 3 the best? *Cancer Control*, 21(3):200–208, 2014.

[33] Rogatko A., Schoeneck D., Jonas W., Tighiouart M., Khuri F. R., and Porter A. Translation of innovative designs into phase I trials. *Journal of Clinical Oncology*, 25(31):4982–6, 2007.

[34] CTCAE. Cancer Therapy Evaluation Program. Common Terminology Criteria for Adverse Events, Version 3.0, DCTD, NCI, NIH, DHHS. https://ctep.cancer.gov, 2003.

[35] Ezzalfani M., Zohar S., Qin R., Mandrekar S. J., and Deley M.-C. L. Dose-finding designs using a novel quasi-continuous endpoint for multiple toxicities. *Statistics in Medicine*, 32(16):2728–2746, 2013.

[36] Wang Y. and Ivanova A. Dose finding with continuous outcome in phase I oncology trials. *Pharmaceutical Statistics*, 14(2):102–107, 2015.

[37] Lee S. M., Cheng B., and Cheung Y. K. Continual reassessment method with multiple toxicity constraints. *Biostatistics*, 12(2):386–398, 2011.

[38] Postel-Vinay S., Collette L., Paoletti X., Rizzo E., Massard C., Olmos D., Fowst C., Levy B., Mancini P., Lacombe D., Ivy P., Seymour L., Le Tourneau C., Siu L. L., Kaye S. B., Verweij J., and Soria J.-C. Towards new methods for the determination of dose limiting toxicities and the assessment of the recommended dose for further studies of molecularly targeted agents–Dose-limiting toxicity and toxicity assessment recommendation group for early trials of targeted therapies, an European organisation for research and treatment of cancer-led study. *European Journal of Cancer*, 50(12):2040–2049, 2014.

[39] Andrews G., Askey R., and Roy R. *Special Functions*, chapter 2. Cambridge University Press, Cambridge, 1999.

[40] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016.

[41] Hankin R. Special functions in R: Introducing the gsl package. *R News*, 6(4), 2006.

[42] Hoff P. D. *A First Course in Bayesian Statistical Methods*. Springer Texts in Statistics. Springer New York, New York, NY, 2009.

[43] Yong K., Cavet J., Johnson P., Morgan G., Williams C., Nakashima D., Akinaga S., Oakervee H., and Cavenagh J. Phase I study of KW-2478, a novel Hsp90 inhibitor, in patients with B-cell malignancies. *British Journal of Cancer*, 114(1):7–13, 2016.

[44] Simon R., Freidlin B., Rubinstein L., Arbuck S. G., Collins J., and Christian M. C. Accelerated titration designs for phase I clinical trials in oncology. *Journal of the National Cancer Institute*, 89(15):1138–1147, 1997.

[45] Fernandes L. L. *Adaptive Phase I and II Clinical Trial Designs in Oncology with Repeated Measures using Markov Models for the Conditional Probability of Toxicity*. PhD thesis, University of Michigan, 2014.

[46] Yin J., Qin R., Ezzalfani M., Sargent D. J., and Mandrekar S. J. A Bayesian dose-finding design incorporating toxicity data from multiple treatment cycles. *Statistics in Medicine*, 36(1):67–80, 2017.

[47] Uhlenbeck G. E. and Ornstein L. S. On the theory of the Brownian motion. *Physical Review*, 36(5):823–841, 1930.

[48] Hastings W. K. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

[49] Cheung K. Y. *dfcrm: Dose-finding by the continual reassessment method*, 2013. R package version 0.2-2.

[50] Plummer M. JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*, 2003.

[51] Eddelbuettel D. and François R. Rcpp: Seamless R and C++ integration. *Journal of Statistical Software*, 40(8):1–18, 2011.

[52] Cunanan K. and Koopmeiners J. S. Evaluating the performance of copula models in phase I-II clinical trials under model misspecification. *BMC Medical Research Methodology*, 14:51, 2014.

[53] Gerschgorin S. Über die Abgrenzung der Eigenwerte einer Matrix. *Bulletin de l'Académie des Sciences de l'URSS*, 6:749–754, 1931.

[54] Brouwer A. E. and Haemers W. H. *Spectra of Graphs*. Universitext. Springer New York, New York, NY, 2012.

[55] Usmani R. A. Inversion of a tridiagonal Jacobi matrix. *Linear Algebra and its Applications*, 212:413–414, 1994.

[56] Wilf H. S. *Generatingfunctionology*. Academic Press, Boston, MA, 1994.