Adaptive Phase I and II Clinical Trial Designs in Oncology with Repeated Measures using Markov Models for the Conditional Probability of Toxicity

by

Laura Levette Fernandes

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biostatistics)
in The University of Michigan
2014

Doctoral Committee:

Professor Jeremy M.G. Taylor, Co-chair
Associate Professor Susan Murray, Co-chair
Assistant Professor Rashmi Chugh
Professor Thomas Braun

For my parents George and Tina, siblings Leona, Jonathan and Joel,

grandmother Maria and husband Lars.

Thank you for everything, I love you.

ACKNOWLEDGMENTS

Successfully completing my thesis and graduating with a PhD in biostatistics has become a reality due to the concerted efforts, prayers and interactions of many individuals. I would like to thank Jeremy and Susan, my co-advisors, for all the hours they spent in meetings and disseminating their knowledge to me. Thank you Susan for your guidance in helping me write and polish my thesis. Thank you Jeremy for introducing me to various aspects of clinical research in oncology and for taking me on as a graduate student research assistant (GSRA). Thank you to Tom and Rashmi, for agreeing to be on my committee and for all the helpful comments you have given me.

I would like to thank all the other faculty and staff in the department of biostatistics and especially Professor Timothy Johnson who funded me as a GSRA during my masters. I would also like to thank all the collaborators and staff at the Cancer Center for a rich working experience in cancer research.

I am thankful to my office mates, past and present, for all the helpful suggestions and advice to survive graduate school; Jungeun Lee, Yong-Seok Park, Yong Zhang,

Lyrica Liu, Shelby Huang, Ludi Fan, Maria Larkina, Yi-Ann Ko, Su Chen and Jared Foster have been amazing pals. Special thanks to Nabihah Tayob, Azarius Reda, Quang Duong, Miku Kawakami, and Naseem Jauhar for being the best friends since our first semester at UM.

I am grateful to all the staff, parishioners and 'pew-peeps' at St. Marys Student Parish on campus for being my second home, for the many friendships and for being nurtured and nourished in this place.

My parents, George and Tina, my sister Leona, brothers Jonathan and Joel and my dearest grandmother Maria P. Dias, I love you dearly and am very grateful for your presence in my life. Thank you for your unconditional love, support and belief in my dream and for reminding me about it even when my own resolve wavered.

My husband Lars Daldorff, I had not planned to marry before graduating! Thank you for coming into my life and for your love and patience in reading the numerous drafts of my thesis, looking forward to growing old together. Thank you also for your warm and welcoming family.

Finally I am most thankful to God, for guiding me all along in my life and giving me nothing but the best in life. I am thankful for the gift of faith to believe and to trust that all things are possible through God who alone suffices in the end!

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

ABSTRACT

Adaptive Phase I and II Clinical Trial Designs in Oncology with Repeated Measures using Markov Models for the Conditional Probability of Toxicity

by

Laura Levette Fernandes

Co-chairs: Professor Jeremy M.G. Taylor

Professor Susan Murray

We consider models for the dose toxicity relationship in early clinical trials in oncology where different dose levels of a study drug are being tested over multiple cycles in the same patient and an assessment of toxicity is made for each cycle. We propose three models using conditional probability of toxicity in specifying the dose-toxicity relationship in patients receiving repeated doses assuming that they did not have any dose limiting toxicities (DLTs) on past cycles. We first develop the conditional Markov model in a phase I settings where the patients are allowed to escalate/de-escalate dose levels, from a choice of five possibilities, over six cycles. In the second setting the conditional Markov model is applied to a completed phase II clinical trial

in sarcoma patients from the paper by Worden et al. (2005) where two dose levels of the study drug, ifosamide, were tested over four cycles. The model adequately fits the dose-toxicity relationship at each of the cycles and demonstrates flexibility offered in including additional covariate terms to describe the relationship. Finally the conditional Markov model is extended to the ordinal case where patient responses are classified as severe, mild or none and might prove beneficial in assigning future doses closer to the patient's actual frailty. Bayesian estimation of the parameters is formulated and evaluated through simulations in all the three methods. Methods for utilizing the dichotomous and ordinal outcome method to conduct a phase I study, including choices for selecting doses for the next cycle for each patient, are developed and designs of clinical trials using the models in simulation settings are presented. Comparison of the dichotomous and ordinal outcome Markov models are also presented exploring the potential benefits of using ordinal outcomes in conducting a trial.

CHAPTER 1

Introduction

Early phase I clinical trials of a new agent in oncology are conducted as dose-finding experiments with a focus on estimating the maximum tolerated dose (MTD) and understanding the dose-toxicity relationship. The designs of such trials typically explore the toxicity at a predefined set of possible dose levels of the agent. Since the MTD will be used in subsequent studies of the agent it is important that it be established with some level of confidence from the phase I trial. The trials are typically small with less than 30 patients, non-randomized and sequential in nature so that during the trial patients are assigned the maximum dose considered safe and tolerable based on available information at that point. A key question in the conduct of these studies is what dose should be assigned to the next patient who is about to enroll in the study. There are many different approaches that can be taken, some are algorithmic, such as the commonly used '3+3' design, others are based on a statistical model such as the continual reassessment method (CRM) [O'Quigley

et al., 1990] and variations of it such as the escalation with overdose control (EWOC) [Babb et al., 1998] design. Model based designs are based on statistical principles and use information from all the patients in the trial to make decisions on dose assignment for new patients. Much research has shown that model based designs are better at estimating the MTD and in treating patients closer to the therapeutic dose level than the '3+3' design [O'Quigley and Chevret, 1991, Thall and Lee, 2003]. In model based designs an explicit target toxicity rate of say 30% is specified, and a statistical model is posited for the relationship between the dose and the toxicity. At the time the new patient is about to enroll the model is fit to the data, then the dose that would give at or just below the expected target toxicity rate is selected for the new patient. The form of the statistical model for the relationship between dose and probability of a dose limiting toxicity (DLT) would usually be simple and have a smooth, sigmoid, monotonic shape. As the data accumulates during the trial the model is refit, leading to possibly a different dose assignment for the next patient.

The initial patients in single-dose trials are started off in their first cycle of treatment at low dose levels and even if they continue to receive multiple doses on additional cycles only the data from the first cycle is used when deciding the dose level for the next patient. A clinical drawback of considering the outcome measure to be based on just the first one or two cycles, is what if there is a DLT at a later cycle, it would probably be important to take that into consideration in recommending a dose to use in the future. [Postel-Vinay et al., 2011] showed that DLT's do frequently

occur in later cycles. From an ethical standpoint this design could be improved by allowing patients to receive the highest dose level that is the most safest and by using the data from all the patients in the trial when making dose assignment decisions. Trials that allows multiple doses per patient, impose restrictions on the dose assignment choices available to the patients wherein patients are administered the same dose level on all the cycles. Such restrictions prevent patients from escalating to a higher dose level and receiving more of the study drug when other patients in the trial are performing well and vice versa the patients are prevented from de-escalating to a lower dose level in the event that many toxicities are observed in the trial from the other patients.

The benefit of accelerated titration designs was recognized by [Simon et al., 1997] who provided the rationale for allowing patients to vary doses across cycles. [Simon et al., 1997] considered a random effects models to simulate data with separate toxicities measures for each cycle. This model was used to simulate data for the evaluation of the accelerated titration method, but the model was not used for data analysis and trial conduct. Motivated by considerations of pharmacokinetics [Legedza and Ibrahim, 2000] developed a model for repeated toxicity measures for each patient. Their model included a random effect to allow for different levels of frailty for a person, giving within subject correlation, and also included a term to represent cumulative effects of toxicity. However, they had considerable computational difficulties in fitting their model and eventually a much simplified version of the model

was able to be fit without estimating the random effects. More recently [Doussau et al., 2013] presented models incorporating ordinal outcomes from patients receiving multiple cycles of doses. One of the major drawbacks of these models is that they only apply to situations in which the patients receive the same dose level on all the cycles thereby taking away the advantages of intra-patient dose escalation or de-escalation described earlier.

If a patient does experience a toxicity on any cycle they are typically taken off the study and they would not provide further data for the assessment of toxicity. Denoting 0 to represent no dose limiting toxicity (DLT) and 1 to represent a DLT from a dosing cycle, where the National Cancer Institute [NCI, 2003] criteria of grading toxicities defines grades higher than 3 as a DLT. The data for each patient would either consist of a series of zeroes (for example 000000) or a series of zeroes followed by a one (for example 001). While a subject-specific random effect is an appealing way to incorporate concepts of frailty, it is clear that fitting models with random effects to the above type of data is going to be very challenging.

This dissertation presents a new approach of using conditional probability of toxicity to model the dose-toxicity relationship in patients with multiple cycles of the study drug assuming that further drugs are given only if the patient had no DLTs in the previous cycles. The use of conditional Markov models is the novel unifying idea in the three chapters of the dissertation.

In Chapter 2 the conditional Markov model in presented in a phase I setting.

We develop a two-state Markov model, with the states being 0 and 1. State 1 is considered a terminating state occurring when a patient experiences a DLT. The model is presented for five dose levels assuming that patients in the study would receive doses until completion of six cycles without a DLT. We adopt a Bayesian approach to estimation and to provide improved small sample performance of the estimates we utilize informative priors that can be solicited from experts prior to the trial. Parameter estimation by allowing patients to vary dose levels over the course of the trial will be demonstrated. Additional simulations to demonstrate the potential benefits of analyzing data from all the patients over all the cycles as opposed to reducing it to a single binary outcome per patient will be presented. Finally the use of the model in designing and executing a sequential trial will be presented.

Chapter 3 focuses on the applicability and extensions of the conditional Markov model in modeling the data generated from a completed phase II clinical trial. Data from the oncology trial conducted by [Chugh et al., 2007] is used as an example. In this randomized phase II clinical trial two dose levels of the study drug ifosamide were tested over four cycles in patients having soft tissue sarcoma. Various models incorporating covariates are proposed to correctly specify the dose-toxicity relationship. Priors are developed for the parameters in the model and the flexibility offered in including additional covariate terms is demonstrated both via simulations and through the example dataset.

Chapter 4 provides extensions to the concept of modeling the conditional proba-

bilities in a phase I setting with three ordinal outcomes; severe, mild or none toxicity in the response. The dichotomized conditional Markov model is modified to account for the mild toxicity responses in the past and its use is demonstrated in the conduct of an adaptive clinical trial. The benefits of using the ordinal outcomes are presented via simulations when compared to the initial two-state Markov model presented in Chapter 2. To conclude Chapter 5 presents an overall discussion of the proposed conditional Markov models and considers a number of potential extensions and modifications of these models in other settings.

CHAPTER 2

Adaptive Phase I Clinical Trial Design Using Markov Models for Conditional Probability of Toxicity

2.1 Introduction

A key question in the conduct of dose-finding phase I trials in oncology is what dose should be assigned to the next patient who is about to enroll in the study. The algorithmic approach could use '3+3' design [Storer, 1989] while a statistical model based approach could use the continual reassessment method (CRM) [O'Quigley et al., 1990] and variations of it such as the escalation with overdose control (EWOC) design [Babb et al., 1998]. Much research [O'Quigley and Chevret, 1991, Thall and Lee, 2003] has shown that model based designs are better in estimating the maximum tolerated dose (MTD) and in treating patients closer to the therapeutic dose level than the '3+3' design. In such model based designs a smooth, sigmoid, monotonic

shape is posited for the relationship between the dose and the probability of toxicity and an explicit target toxicity rate of say 30% is specified. When a new patient is about to enroll the model is fit to the data, and the dose at the acceptable expected target toxicity rate is selected for the new patient. As the data accumulates during the trial the model is refit, leading to possibly a different dose assignment for the next patient.

Because the trials typically start at a cautiously low dose level, some of the patients, especially early on in the study, are treated at a low dose level and hence probably receive limited benefit from the treatment. [Simon et al., 1997] provided the rationale for the accelerated titration design, where patients were allowed to receive different doses on each cycle. A random effects model was used to simulate the toxicities that could occur on different cycles. Motivated by pharmacokinetic considerations [Legedza and Ibrahim, 2000] developed models for repeated toxicity measures for each patient by including a random effect to account for patient correlation and a term for capturing the cumulative effect of toxicity. Due to considerable computational difficulties in fitting the model, owing to the nature of the data, they were only able to fit a much simplified version.

If a patient does experience a dose limiting toxicity (DLT) on any cycle they are typically taken off the study not providing further data for the assessment of toxicity. If 0 represents no toxicity from a cycle and 1 represents a DLT then the data for each patient would either consist of a series of zeroes (for example 000000) or a series of

zeroes followed by a one (for example 001). While a subject-specific random effect is an appealing way to incorporate concepts of frailty, fitting models with random effects to such example data is challenging. In this chapter as an alternative we develop a two-state Markov model, with the states being 0 and 1. Because 1 is a terminating state, we only need to consider the transition probabilities out of state 0. We explicitly model conditional probabilities of toxicity in a cycle given that the patient is toxicity-free to date. At the first cycle the probability of toxicity depends just on the dose, at later cycles the conditional probability of toxicity can depend on additional covariates such as the cumulative dose and the maximum of the past doses.

The model includes a number of parameters, which need to be estimated from the data. Since we envision that the model would be fit during the conduct of the trial, an estimation method is necessary that can be used even for small sample sizes, as would be the situation early in the trial. We adopt a Bayesian approach and to provide improved small sample performance of the estimates we utilize informative priors that can be solicited from experts prior to the trial.

Once the parameter values are known the form of the model allows a number of different calculations to be made. For example, the probability of toxicity on the next cycle as a function of dose can be calculated. Also the probability of toxicity on any future cycle can be calculated, and this will be a function of the sequence of doses that will be given on each of the future cycles. This raises an interesting question

as to how to select the next dose. Should it only be influenced by the probability of toxicity on the next cycle, or should a more long term horizon be taken into account and the probability of toxicity at any future time be considered. In selecting the dose for the next cycle it may be beneficial to think not just about the dose for that cycle, but also the doses for future cycles.

Since the design allows for intra-patient dose changes during the conduct of the trial, the recommendation at the end of the trial could also be a sequence of doses, which vary from one cycle to the next. Wild between-cycle variations in the recommended dose level are unlikely to be clinically acceptable, however a modest variation, such as dose level 3 for the first 2 cycles, then dose level 4 for the last four cycles, could be envisioned. Allowing for intra-patient dose variation also presents another practical concern. At the end of the trial a recommended schedule of doses will be provided, yet no patient in the trial may have exactly followed this regimen. Thus another consideration in deciding the next dose for each patient, is that the schedule of doses for that patient should be one that could be recommended at the end of the trial, or at least close to one that it is conceivable to recommend.

In Section 2.2 we describe the Markov model providing intuition on model features. The Bayesian estimation method is described, with consideration given to the selection of the prior distributions. In Section 2.3 we evaluate properties of the estimation method in a static situation of a small and a moderate sample size. In Section 2.4 we consider using the model in the design and conduct of a trial and

consider optimality criteria for choosing the next dose for a patient. We evaluate the designs and compare them with some simple alternatives. We end with a discussion in Section 2.5.

2.2 Methodology

2.2.1 Notation and data structure

We assume that there are five increasing dose levels of an experimental study drug represented by $d_g, g = 1, \ldots 5$, that will be studied in $i = 1 \ldots N$ patients. Each patient i completes $K_i \leq 6$ cycles, where K_i may be less than six if a patient experiences a DLT or if the patient drops out for other reasons. On each cycle $k = 1, \ldots, K_i$, patient i receives a dose $d_{i,k}$ equal to one of the five values of d_g. A patient's cumulative dose prior to cycle $k = 1, \ldots, K_i$ is $D_{i,k} = \sum_{j=1}^{k} d_{i,j-1}$, so that $D_{i,1} = 0$. We also use the notation that $d_{i,k-1} = 0$ for cycle $k = 1$ and $d_{i,k}^{\ddagger} = max(d_{i,1}, \ldots, d_{i,k-1})$, the maximum of doses assigned to patient i until current cycle k.

The occurrence of a DLT for patient i on cycle k is $Y_{i,k}$, with $Y_{i,k} = 0$ indicating no DLT and $Y_{i,k} = 1$ for a DLT. Patients stop receiving the drug if they experience a DLT thus the possible patterns of $Y_{i,k}$ values for a patient are a sequence of zeroes or a sequence of zeroes followed by one. The observed data after n patients have enrolled in the trial is $\{(Y_{i,1}, \ldots, Y_{i,K_i}, d_{i,1}, \ldots, d_{i,K_i}), i = 1, \ldots, n\}$.

2.2.2 Proposed Markov model

We propose Model 2.1 for the conditional probability of toxicity, $p_{i,k} = P(Y_{i,k} = 1|Y_{i,k-1} = 0, \ldots, Y_{i,1} = 0)$, for patient i on cycle k given that patient i has experienced no previous DLTs as,

$$\ln(1 - p_{i,k}) = -\alpha \left(d_{i,k} - \rho d^{\ddagger}_{i,k}\right)^{+} - \beta D_{i,k} d_{i,k} \tag{2.1}$$

or equivalently,

$$p_{i,k} = 1 - \exp\left\{-\alpha \left(d_{i,k} - \rho d^{\ddagger}_{i,k}\right)^{+} - \beta D_{i,k} d_{i,k}\right\}.$$

The term $(d_{i,k} - \rho d^{\ddagger}_{i,k})^{+}$ is, equal to $(d_{i,k} - \rho d^{\ddagger}_{i,k})$ if $d_{i,k} > \rho d^{\ddagger}_{i,k}$, and is zero otherwise. Intuition behind Model 2.1 can be appreciated by starting with the first cycle, $k = 1$, when $p_{i,1} = 1 - \exp(-\alpha d_{i,1})$ and only α comes into play in explaining the dose related toxicity. To obtain valid probability estimates, $\alpha \in [0, \infty]$ so that $p_{i,1} \in [0, 1]$ is an increasing function of $d_{i,1}$. Note that we do not need to develop a model for $P(Y_{i,k} = 1|Y_{i,k-1} = 1)$ since once a patient develops a DLT at cycle $k - 1$ no further dose is administered to the patient.

On subsequent cycles we have two different terms to account for the conditional probability of toxicity. The first term $(d_{i,k} - \rho d^{\ddagger}_{i,k})^{+}$ accounts for difference between the current assigned dose $d_{i,k}$ and a factor (ρ) of the maximum of the previous doses

$d^{\ddagger}_{i,k}$, while the second term tries to capture the effect of the cumulative dose $D_{i,k}$.

The parameter ρ can be thought of as reflecting the amount of memory about whether a dose was tolerable. When $\rho = 1$, $(d_{i,k} - \rho d^{\ddagger}_{i,k})^+$ reduces to $(d_{i,k} - d^{\ddagger}_{i,k})^+$ as the difference between the current assigned dose and the maximum of the previous doses. If the current dose is less or equal to the maximum of the previous doses, the difference will be zero and will not contribute towards the probability estimate i.e., there is a strong memory that a dose equal to or higher than the current dose was tolerable hence the current dose is more likely to be tolerable. When $\rho = 0$, the term $(d_{i,k} - \rho d^{\ddagger}_{i,k})^+$ reduces to $d_{i,k}$ and implies that there is no memory of the previous doses that had been tolerated. Intermediate values of ρ between zero and one have intermediate amount of memory. Thus this term tries to capture the within-patient correlation between dose cycles.

The term $\beta d_{i,k} D_{i,k}, \beta \geq 0$ is designed to capture the idea that there may be "damage" accumulated from prior doses and the amount of this "damage" plays a role in determining the probability of toxicity when a new dose is administered. This term is constructed so that the contribution of the cumulative dose is in proportion to the current dose $d_{i,k}$ and will not be relevant if $d_{i,k} = 0$.

Figure 2.1 plots the conditional probability of toxicity for different instances of α, β and ρ and aids in understanding the working properties of the Markov Model 2.1. The solid line with open circles shows the probability of toxicity for the first cycle at each of the five dose groups. The curve is the same in all the nine panels

since the probability of toxicity on the first cycle is influenced only by α which is the same in all the instances. The dashed line with crosses corresponds to the conditional probability of toxicity on the second cycle assuming that patients have received dose level three with no DLTs on cycle 1 and any one of the five dose levels on the second cycle. In the top left panel with $\beta = 0, \rho = 0$ cycle 2 gives probabilities equivalent to those seen in cycle 1. This is because there is no cumulative effect of dose ($\beta = 0$) and patients surviving the first cycle are treated as though they are similar to patients on cycle 1 with respect to chance of toxicity since ($\rho = 0$) i.e., no memory. The first row from left to right indicates that increasing β gives increasing probabilities of toxicity on cycle 2 even when $\rho = 0$. Panels in the first column from top to bottom indicate that when there is no cumulative effect of dose ($\beta = 0$) on cycle 2 increases in ρ make patients less likely to experience a toxicity. For instance $\rho = 1$ suggests that all patients who would have experienced toxicity at dose level three ($d_{i,1} = d_3$) were eliminated from the trial during cycle 1 resulting in probability of toxicity equal to zero until $d_{i,2} > d_3$ in the lower left panel. Hence toxicities in cycle 2 are both a function of patient selection in subsequent cycles as influenced by ρ as well as cumulative dose effects as influenced by β

2.2.2.1 Comparison to existing models

This section provides a brief comparison of the dose-toxicity relationship captured by the Markov Model 2.1 compared to alternative models. [Simon et al., 1997]

modeled a latent continuous toxic response $W_{i,j}$ for person i at time period j as,

$$W_{i,j} = \beta_i^S + \epsilon_{i,j} + \log(d_{i,j}^S + \alpha^S D_{i,j}^S) \tag{2.2}$$

where $d_{i,j}^S$ is the dose for person i at time j, $D_{i,j}^S$ is the ith person's cumulative dose prior to time j and the two random effect terms, $\beta_i^S \sim N(\mu_{\beta^S}, \sigma_{\beta^S}^2)$ accounting for inter-patient variability or frailty and $\epsilon_{i,j} \sim N(0, \sigma_\epsilon^2)$ representing the intra-patient variability. This model was used to model data generated from a trial using a pre-defined escalation plan with the continuous toxicity response categorized into different levels using pre-defined thresholds. On the first cycle there is no cumulative dose and there is no parameter to capture the contribution of the current dose $d_{i,j}^S$ which is simply reduced to a log transformed term, β_i^S and $\epsilon_{i,j}$ are the only terms that help in explaining the effect of the first dose. On subsequent cycles α^S captures the effect of the cumulative dose and can be likened to the β parameter in the Markov model. The effect of the current dose is not captured by any parameter but is tied in with the σ_ϵ^2 which also tries to capture the intra-patient dose dependency. Hence in broad terms β_i^S can be likened to the α term and σ_ϵ^2 to the ρ term in the Markov model. The Markov Model 2.1 has one less parameter to estimate and yet provides a similar explanation of the dose toxicity profile.

[Legedza and Ibrahim, 2000] proposed the use of clearance rate λ^L for cumulative effects. The form of their model is as follows,

$$\text{logit}(p^L_{i,j}) = \epsilon^L + \beta^L \log(d^L_{i,j} + D^L_{i,j}\exp(-\lambda^L)) \tag{2.3}$$

where $p^L_{i,j}$ and $d^L_{i,j}$ is the probability of toxicity and the dose for person i at time j respectively, $D^L_{i,j}$ is the ith person's cumulative dose prior to time j. The ϵ^L term is not patient or cycle dependent. A simpler model which excluded the ϵ^L term was also considered. Due to the in-feasibility in estimating λ^L with small sample sizes the clearance rate was assumed to be a constant, $\lambda^L = \log(2)$. The model has a fixed intercept and the effect of the current dose is captured by β^L while λ^L captures the effect of the cumulative dose. Legedza's model needs the estimation of only two parameters and in the absence of the ϵ^L term a single parameter β^L, however it does not capture any dependency between the response of a patient on different cycles.

Notice that Legedza's model has increasing probability of toxicity with dose on subsequent cycles. There is no concession given to the patient for surviving a higher dose level on the first cycle. In comparison the Markov model 2.1 allows the toxicity on the second cycle to be higher or lower than that on the first cycle. Depending upon the cumulative effect of the dose, patients surviving a higher dose on the first cycle are less likely to have a toxicity on the second cycle and the model adequately accounts for this.

A third approach using a cure rate model for estimating the cumulative effect of multiple administrations of the study drug is provided by [Zhang and Braun, 2013]. This approach considers multiple dose levels administered to patients at fixed time points with a goal to select the optimal dose level and schedule (regimen) at the end of the study. Individual hazard contributions from the doses are summed up in estimating the cumulative effect of the multiple administrations. A cure rate model is used to describe the hazard function. The hazard of a DLT at time t following administration of dose d_1 at t_1 is given by the formula, $h(t) = \theta_{i,1}F(\nu_{i,1}|\phi)$. Thus the probability of having a DLT in an interval (t_k, t_{k+1}) can be calculated by solving the integral, $p_k = exp(-\int_{t_k}^{t_{k+1}} h(u)du)$. Zhang et al assumes that the time to DLT from any of the individual administrations is independent in contrast to the Markov model that considers the dependency between patient responses through the concept of frailty in ρ. This also leads to the second difference where the probability of toxicity (defined via the hazard) is assumed to increase with dose administrations while the Markov model allows flexibility for a decrease in probability on the second cycle. The Markov Model 2.1 assumes that the observed response in a cycle is due to the drug administered during that cycle while in Zhang's method, doses are administered based on a schedule until the toxic response is observed allowing for delayed toxicities which are not allowed in the Markov model. Although the Markov model is presented for six cycles the final recommendation of the regimen could be made for cycles less than six that provide an acceptable overall probability of toxicity

which is similar to the idea used by Zhang in schedule selection.

[Pye and Whitehead, 2012] presented at a conference a Bayesian designs for phase I clinical trials in cancer by assuming the the observations arose as interval-censored from a survival model. They used a generalized linear model to represent the relationship between the probability of experiencing a DLT during cycle j of the treatment conditional on there being no DLT prior to cycle j on dose level k in a time to event (survival) setting. The probability of observing a toxicity on cycle j assuming dose level k has been administered is given by $p^P_{k,j}$ and estimated using the model:

$$p^P_{k,j} = 1 - exp(-exp(\gamma_j + \beta^P \log(d_k))) \tag{2.4}$$

A Beta prior is assigned on $p^P_{k,j}$ by incorporating prior beliefs through pseudo-data. Parameters estimates are found using GLM software. Their model would have to estimate $j+1$ parameters, corresponding to the j dosing cycles and the effect of the dose. For the particular case with six cycles a total of seven parameters would have to be estimated, in contrast the Markov model estimates only three parameters. Their model also assumes proportional hazard across the dose levels and requires that the same dose level be given on all the cycles. There is no parameter to account for the cumulative effect of the dose. The conditional aspect of the model is similar to our Markov model and shares the same feature of estimating the probability of toxicity on all the cycles.

More recently [Doussau et al., 2013] provided a mixed effects proportional odds model to incorporate ordinal outcomes in a phase I setting to describe the probability of a severe toxicity and the trend in the risk of toxicity with time. This method does not explicitly model the tendency to discontinue cycles for patients who have demonstrated previous DLT, although the resulting estimated toxicity rates may be conditional in nature. In addition the cumulative effect of the dose is not captured and patients are not allowed to escalate or de-escalate doses. The details of this method will be discussed further in Chapter 4 when the ordinal Markov model is presented.

2.2.3 The likelihood, prior and posterior distributions

2.2.3.1 Probability Skeleton

The dose levels to be studied are transformed to d_g via pre-specified skeleton probabilities denoted by q_g. The skeleton probabilities incorporate prior beliefs about the dose-toxicity relationship and correspond to the probability of observing a toxicity on the first cycle for each of the dose levels. In our set up of the Markov model the probability of toxicity on the first cycle is given by $\ln(1 - p_{i,1}) = -\alpha d_{i,1}$ and does not depend on ρ and β. The doses d_g are obtained by transforming q_g via $d_g = -\ln(1 - q_g)$ and thereby setting the prior mean on $\alpha = 1$. Similar transformations are described by [Lee and Cheung, 2009] in the context of the CRM and have been used by many other authors in other contexts [Lee et al., 2011, Cheung

and Elkind, 2010]. The probability skeleton information can be elicited from prior animal studies or from the clinicians. In the absence of such information, [Lee and Cheung, 2009] suggest sensitivity analysis across different skeleton choices.

2.2.3.2 Prior selection and posterior distribution

Based on the study design, patients contribute to the likelihood until they experience a DLT or the final K^{th} cycle is completed. That is, a person with toxicity on cycle K_i gives data $(Y_{i,1} = 0, Y_{i,K_i-1} = 0, \ldots, Y_{i,K_i} = 1, d_{i,1} \ldots d_{i,K_i})$ and the contribution to the likelihood is,

$$P(Y_{i,1} = 0, \ldots, Y_{i,K_i-1} = 0, Y_{i,K_i} = 1) \quad = \quad p_{i,K_i} \prod_{j=1}^{K_i-1} (1 - p_{i,j}). \tag{2.5}$$

And a person completing K cycles without toxicity gives data $(Y_{i,1} = 0, \ldots, Y_{i,K-1} = 0, Y_{i,K} = 0, d_{i,1} \ldots d_{i,K})$ with likelihood contribution as

$$P(Y_{i,1} = 0, \ldots, Y_{i,K-1} = 0, Y_{i,K} = 0) \quad = \quad \prod_{j=1}^{K} (1 - p_{i,j}). \tag{2.6}$$

In general, subject i on cycle k contributes $L_{i,k}(Y_{i,k}|\alpha, \beta, \rho) = (p_{i,k})^{Y_{i,k}}(1 - p_{i,k})^{1-Y_{i,k}}$ to the likelihood, with $p_{i,k}$ parameterized as in Model 2.1 and interpreted as the probability of toxicity on cycle k conditional on having no prior DLTs in previous

cycles. The resulting likelihood for the entire study population is given by,

$$L(Y|\alpha,\beta,\rho) = \prod_{i=1}^{N}\prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha,\beta,\rho).$$

Our goal lies in estimating the posterior distribution of $p_{i,k},\ k = 1,\ldots,K$ in terms of the posterior distributions of parameters α, β and ρ. Prior distributions on these parameters should reflect any auxiliary knowledge of the toxicity profile for the drug/agents being used in the trial, with a large prior variance when this knowledge is limited. In setting the prior on α, the positive real axis is the permitted range of values and a lognormal (μ,σ^2) is used as a suitable prior having the form

$$\pi(\alpha|\mu,\sigma) = \frac{1}{\sqrt{2\pi\sigma^2}}\frac{\exp(-(\log\alpha-\mu)^2/2\sigma^2)}{\alpha}$$

Specifying the prior mean for α as 1 and the prior variance as 4, providing a coefficient of variance (CV) of 2, μ,σ are estimated using the expressions for the mean and variance of the lognormal density, $E(\alpha|\mu,\sigma) = exp(\mu+\sigma^2/2)$ and $Var(\alpha|\mu,\sigma) = exp\{2(\mu+\sigma^2)\} - exp(2\mu+\sigma^2)$.

On cycles $k > 1$ we have multiple dose administrations and need to assign priors on β and ρ. As mentioned earlier in Section 2.2.2 $\rho \in [0,1]$ and captures the correlation within patients receiving multiple doses, with values near zero indicating that the toxicity outcome is not influenced by previously administered doses and a value

near one indicating a lower chance of toxicity from a previously administered dose. A $Beta(a, b)$ prior is used on ρ having density of the form

$$\pi(\rho|a, b) = (\rho)^{a-1}(1-\rho)^{1-b}$$

The hyperparameters are set to $a = 5$ and $b = 1$ and using the expressions for the mean $a/\{a+b\}$ and variance $ab/\{(a+b)^2(a+b+1)\}$ the prior on ρ has a mean of 0.833 and variance of 0.02.

The lognormal density is used as the prior on $\beta > 0$. In setting the prior mean for β two approaches could be considered. Based on the construction of the $\beta D_{i,k} d_{i,k}$ term its contribution is likely to be much smaller than that of the $\alpha(d_{i,k} - \rho d_{i,k}^{\ddagger})$ term. Arbitrarily set the ratio of these two terms to be 0.2 for patients receiving the third dose level ($d_g = d_3$) on the fourth ($k = 4$) cycle. Setting $\rho = 0.80$ and solving for β provides the mean of the prior on β. The standard deviation (SD) of the prior is set to two times the mean to provide a coefficient of variation of two. A second method for setting the prior involves eliciting another skeleton, the probabilities of completing the entire regimen of $K = 6$ cycles with no toxicities assuming that the dose was the same on all the cycles. By setting $\alpha = 1$ and $\rho = 0.80$, five different values of β corresponding to the dose levels d_g are obtained. The prior mean is set

to the mean of these five values of β and the variance is set to either the SD of these five values or to two times the mean to obtain a CV of two.

The posterior distribution for α, β and ρ given the observed data Y is then

$$f(\alpha, \beta, \rho|Y) = \frac{\prod_{i=1}^{N} \prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha, \beta, \rho)\pi_\beta(\beta)\pi_\alpha(\alpha)\pi_\rho(\rho)}{\int_0^1 \int_0^\infty \int_0^\infty \prod_{i=1}^{N} \prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha, \beta, \rho)\pi_\beta(\beta)\pi_\alpha(\alpha)\pi_\rho(\rho) d\beta d\alpha d\rho}.$$

The posterior distribution of α, β and ρ from Model 2.1 can be estimated via Markov Chain Monte Carlo (MCMC) methods [Robert and Casella, 1999] using just another Gibbs sampler (`JAGS`) *rjags* [Plummer, 2011] package through [`R` Development Core Team, 2011]. `JAGS` includes several algorithms for sampling from the posterior distributions produced from the MCMC iterations, for instance the standard Gibbs sampler is available for this purpose. Details of setting up the MCMC simulations are given in Section 2.2.4.

2.2.4 MCMC sampling procedure

MCMC is a general method based on drawing values of α, β and ρ from approximate distributions and then correcting the draws to better approximate the target posterior distribution $f(\alpha, \beta, \rho|Y)$ [Gelfand and Smith, 1990, Gilks et al., 1993]. New samples are drawn based on the current value (the Markov property) and often from two chains starting at disparate initial values. The goal is to have the simulated draws trace a path throughout the parameter space of α, β and ρ, this is achieved by

running the simulations for a large number of draws and monitoring the convergence through diagnostic tests. The Gibbs sampler is the most frequently used algorithm for drawing samples in a multivariate set up. At each iteration of the Gibbs sampler, samples are drawn for each of the parameters conditional on the values of the other parameters. In practice JAGS uses different samplers for each of the parameters depending on the best choice i.e., ease in simulation and simplicity. Inference is based on the posterior samples which need to be assessed for convergence. The early simulation runs known as the burn-in period are discarded, to allow the model to cover most of the sample space values before drawing values for the posterior distribution.

Assessing the dependence of iterations in each sequence through correlation plots, helps in determining the need for thinning. If samples are found to have a high degree of correlation between samples they defy the assumption that subsequent draws from the posterior are independent. To remedy this issue a thinning factor is used to discard the samples in the sequence and retain only a subset of the samples. Monitoring the convergence based on multiple sequences with over disparate starting or initial values gives rise to mixing of several chains and provides the calculation of the R statistic or the potential scale reduction factor [Gelman and Rubin, 1992]. The idea is that the distributions generated from the two separate initial values should converge to the same target distribution confirmed through the Gelman-Rubin plots and R statistic which measures whether there is a significant difference between the variance within several chains and the variance between several chains by scale

reduction factors. Values lower than 1.1 are considered to be acceptable indications of convergence. Samples are discarded and additional samples are iteratively generated until acceptable convergence diagnostics are obtained. Posterior means and other quantities are then estimated from the final chosen sample.

2.2.4.1 Implementation in JAGS

The `JAGS` MCMC approach runs in three stages. In the first compilation stage, the data likelihood and the density definitions of the priors are specified in a model file saved under a `.bug` extension. The model file and the data are passed into the `JAGS` for compilation along with the list of parameters , α, β and ρ, that have to be monitored. The number of parallel chains to be run by `JAGS` are also defined at the compilation stage, where each parallel chain produces independent samples from the posterior distribution. At this stage the compiled model also contains the initial values for all the parameters that are monitored in each of the chains. The `JAGS` code is provided in Appendix 2.6.1. In the second adaptive stage, samplers are automatically assigned by `JAGS` after a pre-specified adaptive phase for each of the parameters based on the likelihood definition of the model. In the third burn-in stage, 10K samples are discarded and finally posterior samples of 100K (thinned by 20) are used in simulations presented in later sections. Before using the samples from the two chains for reporting they are monitored and assessed through diagnostic tests. The correlation between samples generated at each iteration of the MCMC

chain for each of the parameters needs to be sufficiently low. The posterior means, $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$ of the three parameters, are used to calculate the various probabilities of interest.

2.3 Operating characteristics/Results

All simulation results presented in this section demonstrate the estimation of parameters assuming that all the patients have completed the trial. This section studies the following model properties when used in estimating the conditional probabilities 1) the effect of the priors on parameter estimation, 2) the efficiency gains obtained in the parameter estimates when patients are allowed to have dose escalation and/or de-escalation over multiple cycles and lastly 3) demonstration of the benefits in using the Markov Model 2.1 in comparison to two different models with single binary endpoints.

2.3.1 Effect of priors on estimation

The effect of the degree of informativeness as defined by the SD of the priors in estimating the parameters is explored in this section via simulations. In addition the robustness of the estimation process to prior misspecification when the mean of the prior does not coincide with the parameter true values used in generating the data is also studied. A total of 500 datasets were generated under four different cases and using the skeleton probabilities $q_g = (0.02, 0.05, 0.10, 0.15, 0.23)$. In the first and

second case, true values of the parameters were $\alpha = 1$, $\beta = 0.5$ and $\rho = 0.8$ which were changed to $\alpha = 0.8$, $\beta = 0.5$ and $\rho = 0.8$ in the third case and $\alpha = 1$, $\beta = 0.8$ and $\rho = 0.8$ in the fourth case. The $N = 30$ patients were distributed equally to receive one of the five dose levels on all the cycles until completion of $K = 6$ cycles or occurrence of a DLT.

The degree of informativeness in the priors differed at the estimation stage. The prior means $E(\alpha) = 1$ and $E(\beta) = 0.5$ were the same in all the four cases and matched the true value in Cases 1 and 2 but differed from the true values in Cases 3 and 4. The SD was set to two times the mean, $SD(\alpha) = 2$ and $SD(\beta) = 1$ in Cases 1, 3 and 4. In the second case the prior standard deviations were set to five times the mean, $SD(\alpha) = 5$ and $SD(\beta) = 2.5$. In all the four cases $Beta(5, 1)$ prior was used on ρ.

The parameter estimates from the 500 simulated datasets are presented in Table 2.1 with the rows grouped by the four cases. The four columns report (1) the true value of the parameters, (2) the mean and SD of the prior, (3) the mean of the estimated values from 500 datasets and the mean bias from the true value in parenthesis, (4) the mean SD (MSD) of the estimates from 500 datasets, (5) the empirical SD (ESD) of the 500 estimates, (6) the coverage rate of the 95% credible interval across the 500 datasets. Results indicate that the bias in parameter estimates is low except for β in Case 4. The mean SD is slightly higher than the ESD giving slightly conservative estimates of variability that lead to higher coverage rates. The

MSD of the estimates is lower than the prior SD for α and β but comparable in the case of ρ indicating minimal information in estimating this parameter.

The probability estimates obtained from the 500 simulated datasets are presented in Table 2.2 grouped by the four cases and each of the rows corresponding to one of the five dose levels. The columns indicate the mean estimate of the conditional probability of toxicity on the first, the second, the sixth cycle and the overall probability of toxicity on any of the cycles along with the bias from the true values in parenthesis. The results suggest that the model performs suitably, even with prior misspecification, in estimating the true values. It was decided to use the prior from the first case for the simulation results presented hereafter.

2.3.2 Properties of parameter estimates with intra-patient dose variability

The simulation results presented in Section 2.3.1 assumed that the patients were assigned to receive the same dose level on each of the six cycles. This section explores the effects on estimation in the presence of dose heterogeneity within each patient i.e., allowing patients to have dose escalation and de-escalation across the six cycles assuming that there are a total of $N = 30$ patients in the trial. In actual trial conduct we would have dose combinations that are sensible and that do not vary at every cycle. For instance in a regimen of six cycles we might expect to have the first three cycles on d_1 and then switch to d_2 on the subsequent cycles, implying that

dose escalation happened on the fourth cycle. A typical combination of de-escalation might include higher doses d_3, d_4 or d_5 on the first three cycles and the lower dose on the next three cycles. Given that we have five dose levels it is possible to have $P_2^5 = 20$ different combinations of two doses at a time with either escalation or de-escalation. In general we do not allow patients to skip dose levels and ignoring such dose combinations results in eight assignable combinations, $d_1d_2, d_2d_3, d_3d_4, d_4d_5$ and $d_2d_1, d_3d_2, d_4d_3, d_5d_4$ where patients change their dose level on the fourth cycle. Additional dose combinations include three dose levels with changes on cycle three and cycle five of the form, $d_1d_2d_3, d_2d_3d_4, d_3d_4d_5$ and $d_3d_2d_1, d_4d_3d_2, d_5d_4d_3$, providing another six combinations. In total there are 19 possible treatment courses including the five without dose variation listed in Table 2.3.

In the simulation results presented earlier with $N = 30$ patients there was an equal distribution of patients over the five dose levels d_g, with each of the six patients having the same dose level on all the six cycles. This gives 36 assigned cycles for every d_g i.e., d_1 is assigned 36 times, d_2 is assigned 36 times etc.. To obtain a fair comparison to the current setting, the $N = 30$ patients were assigned to each of the 19 combinations while ensuring that there were 36 cycles of each of the dose levels. The dosing profile of these 30 patients is given in Table 2.25 in the Appendix 2.6.3.

The skeleton probability used for the dose transformations is d_g is $q_g = (0.02, 0.05, 0.10, 0.15, 0.23)$ similar to the one used in earlier simulations. A total of 500 datasets were simulated with patients having regimens as listed in Table 2.25. Conditional

probability of toxicity, $p_{i,k}$, for each patient i at each cycle k was calculated using the Markov Model 2.1 for fixed values of $\alpha = 1$ and $\beta = 0.5$ and $\rho = 0.8$ and their probabilities were used to simulate toxicities. Priors on the parameters used were similar to those in the previous analyses.

The results of the parameter estimates from the simulations are presented in Table 2.4 under Section 2.3.2. The corresponding parameter estimates from Case 1 in Table 2.1 are placed under Section 2.3.1 for easy comparison. We notice that the bias and mean SD of α is slightly lower in Section 2.3.2 but that of β is slightly higher.

Table 2.5 presents the corresponding probability estimates. The columns present the probability of toxicity estimates with the bias from the true value in parenthesis and the empirical SD (ESD) of the estimates from the 500 replicates. The results indicate that the bias is comparable but the variability across simulation goes down slightly when patients have the same dose. We conclude that there were no problems in fitting the model by allowing patients to have dose variability and that the results do not have major deviations with regard to the bias and efficiency.

2.3.3 Comparison with models for a single binary endpoint

This section explores the potential gains in using all the data from the six cycles in estimating the probability of toxicity on the first cycle or on any cycle using the Markov Model 2.1 versus models with a single binary end point per patient. We

consider the special case with no dose variation across cycles.

Simulation results are presented based on 500 datasets each having either $N = 10$ and $N = 30$ patients, distributed equally to receive one of the five doses d_g for a maximum of $K = 6$ cycles. The probability skeleton used for the doses is $q_g = (0.02, 0.05, 0.10, 0.15, 0.23)$. For every patient i assigned to dose d_g on cycle k the probability of a toxic response $p_{i,k}$ is calculated using the Markov Model 2.1 and known values of $\alpha = 1, \beta = 0.5$ and $\rho = 0.8$. A DLT response $Y_{i,k}$ is assigned based on a Bernouli($p_{i,k}$) random draw. A patient i continues to receive the same dose on cycle $k + 1$ until $Y_{i,k} = 1, k < 6$ or $k = 6$.

As mentioned earlier in Chapter 1 existing methods for analyzing trials with multiple cycles for a single patient either consider the data only from the first cycle in estimating the probability of toxicity ignoring the toxicities that happen on later cycles or consider an overall toxic response that might have occurred on any of the cycles. In either of the two cases the data for each patient is reduced to a single binary outcome.

Continuing with the notation from the Markov Model 2.1, in the first instance the data is reduced to a single binary outcome by defining $\grave{Y}_i = 1$ for patient i if $Y_{i,1} = 1$ and $\grave{Y}_i = 0$ if $Y_{i,1} = 0$. The probability of toxicity, $\grave{p}_i$, on the first cycle is given by Model 2.7 as follows,

$$\ln(1 - \grave{p}_i) = -\gamma d_i. \tag{2.7}$$

The prior on γ is similar to that used on α, a lognormal density with mean one and variance four.

In the second instance the data is reduced to a single binary outcome $Y_i' = 1$, across all of the cycles for each patient i if $Y_{i,j} = 1$, for any $j \leq 6$ and $Y_i' = 0$ if $Y_{i,6} = 0$. The Model 2.8 used in estimating the probability of toxicity on any cycle, p_i' in this case is,

$$\ln(1 - p_i') = -\delta d_i'. \tag{2.8}$$

Where the probability of toxicity on any of the six cycles (m_j) corresponds to Markov Model 2.1 via

$$m_j = 1 - \prod_{j=1}^{K}(1 - p_{i,j}), \tag{2.9}$$

with $p_{i,j}$ as defined in equation 2.6. The doses d_i' are based on using a probability skeleton $(m_1, m_2, m_3, m_4, m_5)$ corresponding to having a toxicity on any of the six cycles. We then assume that δ has a lognormal prior distribution with mean of one and variance four.

Results in Table 2.6 indicates adequate model fit for the parameters from Markov Model 2.1 and Models 2.7 and 2.8. Table 2.7 presents the simulation results from comparing the Markov Model 2.1 to the two alternatives, Model 2.7 and Model 2.8. The rows are grouped based on the comparison with Model 2.7 or Model 2.8. The columns are grouped by $N = 10$ and $N = 30$ patients and present the probability

of toxicity estimates with the bias in parenthesis and the Empirical SD (ESD) of the 500 estimates. Comparing the results from $N = 10$ and $N = 30$ patients we notice that there is a gain in efficiency and decrease in the bias for all the three models for the larger sample size. The efficiency is slightly higher with comparable bias in the estimates from the Markov Model 2.1 in comparison to both the simpler models. We conclude that there is mild gain in efficiency especially when using $N = 10$ patients and no harm is done is fitting a larger model. The slight gain in efficiency in comparison to Model 2.8 could be attributed to fact the Markov Model 2.1 incorporates the cycle specific information in the process of estimating the parameters and hence provides better overall estimates of the probability of toxicity.

2.3.4 Comparison with models for a single binary endpoint with unequal subjects at the dose levels

In the previous Section 2.3.3 the comparison of the Markov Model 2.1 with the alternative two models was presented when the patients were distributed equally over all the dose levels. In practice there is unequal distribution of patients in a trial at the various dose levels. Simulations results in this section explore the differences in the estimation when there are 3, 3, 10, 10 and 4 patients assigned to each of the five dose levels in a trial with a total of $N = 30$ patients. In the case with $N = 10$ patients the distribution of the patients was 1, 2, 3, 3 and 1 among the five dose levels. Keeping all other features of the data generation unchanged from the equal patient per dose

level case described in Section 2.3.3 a total of 500 datasets were simulated. Table 2.8 presents the results of the parameter estimates for the three models while Table 2.9 presents the probability estimates from comparing the Markov Model 2.1 to the two alternatives, Model 2.7 and Model 2.8.

Comparing the results from $N = 10$ and $N = 30$ patients in Table 2.9 we notice that there is a gain in efficiency and decrease in the bias for all the three models. The bias is lower and the efficiency is higher in the estimates from the Markov Model 2.1 in comparison to the Binary Model 2.7. In the case of comparison to the Binary Model 2.8, either the bias or the empirical SD of the estimates is lower in Markov Model 2.1 if not both simultaneously.

2.4 Implementation of a clinical trial

This section describes the application of the Markov Model 2.1 in designing a sequential clinical trial. The safety criteria for dose assignment, two possible plans in conducting the trial and the evaluation of the trial properties are considered.

2.4.1 Safety Criteria

We begin by defining the safety criteria rules for dose assignment in carrying out a trial with dose escalation and/or de-escalation. Define $r_{g,k} = g, g = 1 \dots 5$ as one of the five dose levels on cycle k corresponding to $d_g, g = 1 \dots 5$ the transformed doses using the probability skeleton. Let $r^{max}_{g,k+1}$ denote the maximum allowed dose

that could be assigned on cycle $k+1$. The following commonly used dose escalation rules will be followed in defining the safety criteria to be used while carrying out an adaptive clinical trial based on Markov Model 2.1.

- The first and the second patient on the trial will be assigned the second lowest dose level, $r_{g,1} = 2$, on cycle 1, allowing the lowest dose level to be eligible for future patients if DLTs are seen in the first few patients on study. For the first patient if there is no DLT, the same dose level is assigned on the second cycle. For subsequent patients and cycles the following rules will be effective.

- Patients are allowed to escalate by one dose level from their previous dose, i.e., a patient tolerating dose level $r_{g,k}$ on cycle k can be assigned doses no higher than $min(r_{g,k}+1, 5)$ on cycle $k+1$.

- A patient can experience a maximum of three dose levels in a dosing regimen, unless de-escalation to a lower dose is required. I.e.,a patient tolerating dose level $r_{g,1}$ on cycle 1 can possibly receive $r_{g,1}+2$, as its highest dose level in the dosing regimen. In combination with the previous rule a patient tolerating dose level $r_{g,k}$ on cycle k can be assigned doses no higher than $r^{max}_{g,k+1} = min(r_{g,1} + 2, r_{g,k}+1, 5)$ on the cycle $k+1$.

- For each new patient being assigned a dose level on cycle 1, the maximum dose level choice would be limited to $r^{max}_{g,1} = max(r^{\ddagger}_{g,1}+1, r^{\ddagger}_{g,k})$, where $r^{\ddagger}_{g,1}$ is the maximum of all the past dose levels assigned to the patients on cycle $k=1$

and $r^{\ddagger}_{g,k}$ is the maximum of all the past dose levels assigned to the patients in the study on cycles $k > 1$. This ensures that the new patient may only jump one dose level from previously assigned cycle 1 doses and may not exceed doses experienced on the trial otherwise.

- The study will conclude when none of the dose levels are included in the tolerable range as determined by the safety criteria defined below or the N^{th} patient has completed the trial.

2.4.2 Defining the eligible regimen set, $R^{regimen}_{i,k}$

Typically in single dose, single cycle trials one assumes a toxicity bound of, say, 30%. Then the current estimate of the probability of toxicity for each dose is compared with this bound to decide on the next dose. Defining bounds is more complex when patients can receive multiple doses on multiple cycles. We will consider the probability of toxicity for the next dose, for the whole sequence of doses and for the sequence of future doses. Let $\mathring{P}(A) = \{\mathring{P}(A_1), \ldots, \mathring{P}(A_K)\}$ be a vector of acceptable toxicity limits for each cycle $1, \ldots, K$. It is convenient to restrict limits for cycles $2, \ldots, K$ to be equivalent and equal to $\mathring{P}(A_2)$, rather than justify different acceptable toxicity levels at each cycle. Define $\mathring{P}(C)$ as the upper limit of the acceptable probability of toxicity across all K cycles, and for patients who have already completed at least one cycle let $\mathring{P}(B)$ be the acceptable probability of toxicity limit on all the remaining cycles.

In general, for the bounds to be consistent with one another, we require $\{1 - \mathring{P}(C)\} \leq \prod_{k=1}^{K}\{1-\mathring{P}(A_k)\}$ and $\{1-\mathring{P}(B)\} \leq \prod_{k=2}^{K}\{1-\mathring{P}(A_k)\}$; these further reduce to $\{1-\mathring{P}(C)\} \leq \{1-\mathring{P}(A_1)\} \times \{(1-\mathring{P}(A_2)\}^{K-1}$ and $\{1-\mathring{P}(B)\} \leq \{(1-\mathring{P}(A_2)\}^{K-1}$ when we assume the limit $\mathring{P}(A_2)$ for cycles $2, \ldots, K$. In practice, one selects bounds for $\mathring{P}(A_1)$ and $\mathring{P}(C)$, and this automatically places restrictions on $\mathring{P}(A_2)$ and $\mathring{P}(B)$. For instance with $\mathring{P}(A_1)$ in the range of $0 - 0.2$ and setting $\mathring{P}(C)$ as either 0.30 or 0.40, legitimate values for $\mathring{P}(A_2)$ are presented in Table 2.10. So for $\mathring{P}(A_1) = 0.20$, $\mathring{P}(C) = 0.30$, we find that $\mathring{P}(A_2)$ can be no larger than 2.64% and $\mathring{P}(B)$ can be no larger than 12.5% so that conditional probabilities of toxicity on later cycles $2, \ldots, K$ are very small.

Monitoring the safety of the patients is ensured by assigning doses that satisfy the set of safety criteria defined in Section 2.4.1 as well as satisfy the bounds $\mathring{P}(A_1), \mathring{P}(A_2), \mathring{P}(B)$ and $\mathring{P}(C)$ defined above.

In general denote dosing regimens by the vector of doses across the K cycles $(r_{g,1}, \ldots, r_{g,K})$. As each patient progresses through cycles $k = 1 \ldots K$, members m of the set of eligible regimens denoted by $R_{i,k}^{regimen}$ change over time as experience on the study matures. For instance, on cycle k, potential members m, of $R_{i,k}^{regimen}$ for patient i take the form $(o_{i,1}, \ldots, o_{i,k-1}, r_{g,k}, r_{g,6})$ where $o_{i,k}$ denote previously tolerated doses for patient i on cycle k and future assigned doses $(r_{g,k}, \ldots, r_{g,6})$ must not exceed $r_{g,l}^{max}$ for $l = k, \ldots, 6$ and must not conflict with bounds defined by $\mathring{P}(A_1), \mathring{P}(A_2), \mathring{P}(B)$ and $\mathring{P}(C)$. For a patient on cycle 1, it is convenient to limit members of $R_{i,1}^{regimen}$ to

reduce computation. Table 2.3 lists a set of desirable regimens that can be used to construct a limited version of $R_{i,1}^{regimen}$ satisfying safety constraints.

The following random variables are useful to collect and statistically summarize immediate and accumulated toxicities during the conduct of the trial. Define $A_{i,k,j}$ as the event of toxicity on cycle k for patient i at dose level j given that there were no DLTs in the past. Hence $\mathrm{P}(\mathrm{A}_{i,k,j}) = p_{i,k}$, where $p_{i,k}$ is calculated using the Markov Model 2.1 for dose level j, and current estimates of $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$ can be used to define its corresponding estimate, $\hat{P}(A_{\hat{i},k,j}) = \hat{p}_{i,k}$. Define $B_{i,k,m}$ as the event of having a toxicity on any remaining cycle k until K for a member m of the regimen set $R_{i,k}^{regimen}$, where $\mathrm{P}(\mathrm{B}_{i,k,m}) = 1 - \prod_{l=k}^{K}(1 - p_{i,l})$ and $\hat{P}(B_{i,k,m}) = 1 - \prod_{l=k}^{K}(1 - \hat{p}_{i,l})$. Also define $C_{i,k,m}$ as the event of toxicity for a future patient assigned to regimen m from person $i's$ regimen set $R_{i,k}^{regimen}$ i.e., $C_{i,k,m} = B_{i,1,m}$ with $\hat{P}(C_{i,k,m}) = \hat{P}(B_{i,1,m})$. During the course of the trial $\hat{P}(B_{i,k,m})$ estimates the current best guess of patient toxicity probability on the remaining cycles while $\hat{P}(C_{i,k,m})$ estimates the best guess of the toxicity probability profile for future patients undergoing regimen m.

2.4.3 Expected dose

A higher planned dose might not be attractive if fewer cycles can be completed at that dose level due to DLTs. During the course of the trial, Markov Model 2.1 can be used to estimate the expected total dose for members m of the eligible regimen set $R_{i,k+1}^{regimen}$ and potentially use this information as part of selecting the current best

regimen for patient i. Using the expression presented in equations 2.5 and 2.6 the expected total dose for a new patient i is,

$$= d_{i,1}p_{i,1} + \sum_{k=2}^{K-1}\left\{\left(\sum_{j=1}^{k} d_{i,j}\right) p_{i,k}\prod_{j=1}^{k-1}(1-p_{i,j})\right\} + \sum_{j=1}^{K} d_{i,j}\prod_{j=1}^{K-1}(1-p_{i,j}).$$

For a continuing patient i in the study who is ready for dose administration on cycle k in the trial the expression for the expected dose is,

$$= \sum_{j=1}^{k-1} d_{i,j} + d_{i,k}p_{i,k} + \sum_{m=k+1}^{K-1}\left\{\left(\sum_{j=1}^{m} d_{i,j}\right) p_{i,m}\prod_{j=k}^{m-1}(1-p_{i,j})\right\} + \sum_{j=1}^{K} d_{i,j}\prod_{j=k}^{K-1}(1-p_{i,j}).$$

2.4.4 Running the trial

Dosing decisions are governed by Markov Model 2.1 and safety criteria laid out in section 2.4.1. In practice, this requires having current information on all patients in the trial so that new and continuing patients have the most up-to-date information as dose recommendations are made. In particular, each time a dose is recommended we should have current estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$, a defined set of eligible regimens $R_{i,k}^{regimen}$ for patient i being dosed on cycle k and estimates of $\hat{P}(A_{i,k,j}), \hat{P}(B_{i,k,m})$ and $\hat{P}(C_{i,k,m})$. Hence, development of an automated procedure is recommended for this trial design.

At the start of the trial assign two patients to the second lowest dose level, $r_{g,1} = 2$. Patients completing a cycle without a DLT will usually either stay at the same dose level or escalate to a higher dose level, although a de-escalation recom-

mendation is possible if additional data on other patients is trending toward lower dose recommendations. On the first cycle a patient i on the study has $r_{g,1}^{max}$ possible choices for dose level with corresponding estimated conditional probabilities of toxicity $\hat{P}(A_{i,1,j}), j = 1 \ldots r_{g,1}^{max}$. Eligible choices for dose level j must satisfy $\hat{P}(A_{i,1,j}) \leq \mathring{P}(A_1)$ on cycle one and $\hat{P}(A_{i,k,j}) \leq \mathring{P}(A_2)$ on cycles, $k > 1$. If no eligible doses are identified, then the model fit indicates that the trial has no remaining safe dose levels. More often, multiple dose levels satisfy the $\mathring{P}(A)$ safety criterion, and we consider not just the subsequent dose, but all remaining doses in making a recommendation. I.e., we must consider the estimates of $\hat{P}(B_{i,k,m})$ and $\hat{P}(C_{i,k,m})$ of all possible members m of $R_{i,k}^{regimen}$.

Two plans for choosing a dose level in this case include 1) maximizing the expected dose over the entire regimen for a patient or 2) selecting a dose that aligns (matches) with one of the dosing schemes from a list of desirable regimens presented in Table 2.3. These two plans are outlined below.

2.4.4.1 Maximizing the expected dose

For a patient being dosed on cycle k, the expected dose is calculated using expressions in Section 2.4.3 for each of the regimens in $R_{i,k}^{regimen}$. The dose level that maximizes the expected dose and satisfies $\hat{P}(B_{i,k,m}) \leq \mathring{P}(B)$ and $\hat{P}(C_{i,k,m}) \leq \mathring{P}(C)$ is selected for cycle k.

2.4.4.2 Matching a regimen

Cycle 1 patients still base their first dose on maximizing their expected dose using current data and subject to eligibility of regimens as previously described. Otherwise, to gain the most possible experience with regimens that would be recommended at the end of the trial, one may favor members m of $R_{i,k}^{regimen}$ that "nudge" the current patient's regimen toward one of the $q = 1, \ldots, 19$ suggested dosing regimens in Table 2.3, some subset of these or completely different user defined regimens not included in this table.

Define distance between a regimen vector m in $R_{i,k}^{regimen}$ and a regimen vector q from Table 2.3 as $R_l = \sum_{c=1}^{K} |m_c - q_c|$. We may wish to select the regimen m that minimizes this distance across all m and q. Alternatively, one might select the subset of regimens from $R_{i,k}^{regimen}$ that satisfy some prespecified distance limit $R_l < a$ and then choose the regimen that assigns the maximum allowable dose from that subset on the next cycle.

2.4.5 Recommending a regimen

At the conclusion of the study, the estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$ are used to calculate the overall probability of toxicity $\hat{P}(C_j)$ for all of the $j = 1 \ldots 19$ regimens listed in Table 2.3. During the conduct of the trial the target probability bounds used were $\mathring{P}(A)$, $\mathring{P}(B)$ and $\mathring{P}(C)$. These bounds especially, $\mathring{P}(C)$, are usually set at higher than acceptable values in practice and when selecting the final regimen we

would use $P^r(A_1)$ and $P^r(C)$ which might be lower than or equal to $\mathring{P}(A_1)$ and $\mathring{P}(C)$ respectively. For example in running the trial $\mathring{P}(C) = 0.40$ which implies an overall toxicity of 40% but in practice 30% toxicities are what we would want to see in the trials. The final recommended (R) regimen can be selected using both $P^r(A_1)$ and $P^r(C)$ or using $P^r(C)$ alone. In the first option the regimen satisfying $\hat{P}(A_1) \leq P^r(A_1)$, $\hat{P}(C_j) \leq P^r(C)$ and maximizing the expected dose is selected as the recommended dosing regimen while in the second option only $\hat{P}(C_j) \leq P^r(C)$ and maximizing the expected dose condition is used in regimen selection.

A corresponding target (T) regimen is also selected and used as a reference for gauging the properties of the completed trial. The target regimen is selected by calculating $P(A_1)$ and $P(C)$ based on the true values of α, β and ρ and is also selected from one of the 19 regimens presented in Table 2.3. For example corresponding to the 19 regimens, the columns in Table 2.11 list the probability of toxicity on the first cycle $P(A_1)$, the probability of toxicity on any cycle $P(C)$, the expected dose on the entire regimen, a flag set to one if the regimen qualifies when using only $P^r(C) = 0.30$ and three other flags set to one when using $P^r(C) = 0.30$ and $P^r(A_1) = (0.05, 0.10, 0.20)$. The regimen 12, 444333, has the maximum expected dose and is selected when only option one $P^r(C) = 0.30$ is used. If option 2 is used for the selection of the target regimen then regimen 15, 223344 is selected when $P(A_1) = 0.05$ or 0.10 and regimen 12, 444333, when $P(A_1) = 0.20$. In this instance with $\alpha = 1, \beta = 0.2, \rho = 0.8$ regimens that offered dose variation were recommended since they offered a higher

expected dose.

2.4.6 Algorithmic form of the two plans

Plan 1 - Maximize the expected dose on the study for each patient i.

- For new patient $\hat{i}$ on cycle 1

    1. Estimate $\hat{P}(A_{\hat{i},1,j})$ using current estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$ at each of the $j = 1 \dots r_{g,1}^{max}$ dose levels. Where $r_{g,1}^{max} = max(r_{g,1}^{\ddagger} + 1, r_{g,k}^{\ddagger})$, $r_{g,1}^{\ddagger}$ is the maximum of all the past dose levels assigned to the patients on cycle $k = 1$ and $r_{g,k}^{\ddagger}$ is the maximum of all the past dose levels assigned to the patients in the study on cycles $k > 1$.
    2. Subset the dose levels that satisfy $\hat{P}(A_{\hat{i},1,j}) \leq \mathring{P}(A_1)$ over all dose levels.
    3. For the dose levels satisfying $\hat{P}(A_{\hat{i},1,j}) \leq \mathring{P}(A_1)$ subset the list of possible regimens from Table 2.3 and calculate the overall probability of toxicity $\hat{P}(C_{\hat{i},1,j})$.
    4. Select the dose level that has an overall probability of toxicity $\hat{P}(C_{\hat{i},1,j}) \leq \mathring{P}(C)$ and maximum expected dose.
    5. If none of the doses satisfy $\hat{P}(A_{\hat{i},1,j}) \leq \mathring{P}(A_1)$ and if there are continuing patients in the study then wait until updated estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$ allow doses to be assigned else the study is terminated.

- For continuing patient i on cycle $k > 1$,

1. List doses $\hat{P}(A_{i,k+1,j}) \leq \mathring{P}(A_2)$ from $r^{max}_{g,k+1} = min(r_{g,1} + 2, r_{g,k} + 1, 5)$ possible choices.

2. If there is more than one satisfying dose level then list the possible dose regimen set $R^{regimen}_{i,k+1}$.

3. Calculate the probability of toxicity $\hat{P}(B_{\hat{i},k+1,m})$ on the remainder of the cycles for each of the regimens m in $R^{regimen}_{i,k+1}$ and the corresponding expected dose using the current estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$.

4. Select the dose level that has probability of toxicity $\hat{P}(B_{\hat{i},k+1,m}) \leq \mathring{P}(B)$ and maximizes the expected dose.

5. If none of the doses satisfy $\hat{P}(A_{\hat{i},k+1,j}) \leq \mathring{P}(A_2)$ and if there are continuing patients in the study then wait until updated estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$ allow doses to be assigned else the study is terminated.

Plan 2 - Observing a favorable dosing regimen by the end of the study

- For new patient $\hat{i}$ on cycle 1

1. Estimate $\hat{P}(A_{\hat{i},1,j})$ at each of the $j = 1 \ldots r^{max}_{g,1}$ dose levels. Where $r^{max}_{g,1} = max(r^{\ddagger}_{g,1} + 1, r^{\ddagger}_{g,k})$, $r^{\ddagger}_{g,1}$ is the maximum of all the past dose levels assigned to the patients on cycle $k = 1$ and $r^{\ddagger}_{g,k}$ is the maximum of all the past dose levels assigned to the patients in the study on cycles $k > 1$.

2. Subset the dose levels that satisfy $\hat{P}(A_{\hat{i},1,j}) \leq \mathring{P}(A_1)$ over all dose levels.

3. For the dose levels satisfying $\hat{P}(A_{\hat{i},1,j}) \leq \mathring{P}(A_1)$ subset the list of possible regimens from Table 2.3 and calculate the overall probability of toxicity $\hat{P}(C_{\hat{i},1,j})$ using the current estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$
4. Select the dose level that has an overall probability of toxicity $\hat{P}(C_{\hat{i},1,j}) \leq \mathring{P}(C)$. Select the highest dose if more than one satisfying dose level.
5. If none of the doses satisfy $\hat{P}(A_{\hat{i},1,j}) \leq \mathring{P}(A_1)$ and if there are continuing patients in the study then wait until updated estimates of $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$ allow doses to be assigned else the study is terminated.

- For continuing patient i on cycle $k > 1$,

1. List doses $\hat{P}(A_{i,k+1,j}) \leq \mathring{P}(A_2)$ from $r^{max}_{g,k+1} = min(r_{g,1} + 2, r_{g,k} + 1, 5)$ possible choices.
2. If there is more than one satisfying dose level then list the possible dose regimen $R^{regimen}_{i,k+1}$ set.
3. Subset the regimens satisfying probability of toxicity $\hat{P}(B_{i,k+1,m}) \leq \mathring{P}(B)$ from the $R^{regimen}_{i,k+1}$ set.
4. For each of the regimens l in the subset calculate the distance R_l from the favorable dosing regimens based on Table 2.3.
5. Select the highest dose level that has $R_l \leq$ a pre-specified value. If none of the regimens satisfy the R_l condition then select the dose level that maximizes the total expected dose.

6. If none of the doses satisfy $\hat{P}(A_{i,k+1,j}) \leq \mathring{P}(A_2)$ and if there are continuing patients in the study then wait until updated estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$ allow doses to be assigned else the study is terminated.

2.4.7 Properties of the design

For the purposes of evaluating the properties of the simulation of clinical trials over multiple replications and comparing the various plans and properties of the target probabilities various test statistics will be calculated that can be grouped into 1) trial conduct or patient characteristics and 2) regimen characteristics as explained below.

Patient characteristics

1. Mean dose per patient over all the replicates. In each of the replicates the total dose given to all the patients will be tracked and then averaged across the number of the patients in that trial. A higher mean dose is desirable implying that the patients in the trial were able to receive as much of the study drug as possible.

2. Mean number of toxicities per study across all the replicates. Also noting the percentage of trials stopping early. Lower values of the mean toxicities are desirable and lower number of trials of stopping early indicate that all patients were assigned to dose levels within the framework of the safety criteria.

3. At the conclusion of every trial the proportion of patients receiving a cumulative dose greater than the expected dose of the recommended regimen is averaged over all the iterations. Higher values are desirable indicating that patients in the trial had experience of the dose quantity recommended at the end of the trial.

4. The proportion of patients receiving dose greater than the expected dose based on the target regimen target (T) regimen is also considered. A higher value is desirable indicating patients in the trial received dose quantities considered safe by the target regimen.

5. Mean number of patients whose regimen matches exactly with the recommended regimen. The distance from the recommended regimen is calculated for each of the patients in the study, and the proportion of patients having distance $= 0$ are averaged all the replications. Higher values are desirable indicating that patients in the trial had experience with the recommended regimen.

6. Mean number of patients having a regimen that matches the target regimen. Higher values are desirable indicating that patients in the trial had dosing regimens matching the target regimen.

7. Mean number of patients having distance ≤ 2 from the recommended (R) regimen. A slightly less stricter rule checking patients with regimens differing by two dose levels. Higher values are desirable.

8. Mean number of the patients having distance ≤ 2 from the target (T) regimen. Higher values are desirable.

Regimen characteristics

1. The mean distance between the target (T) and the recommended (R) regimen over all the replicates. Lower values are desirable indicating a match between the target and recommended regimen.

2. The proportion of trials with an exact match between the target and the recommended regimen. Higher values are desirable.

3. Mean of expected dose based on the target regimen. Higher values are desirable.

4. Mean of expected dose based on the recommended regimen. Higher values matching the target regimen are desirable.

5. The mean of target regimen toxicities. The probability of observing a toxicity on any cycle given the target regimen and the true parameter values is calculated. Low values are desirable.

6. The mean of recommended regimen toxicities. Low values are desirable.

2.4.8 Simulation Design and Results

A clinical trial recruiting a maximum of N=30 patients and each patient having a maximum of six cycles was conducted over 500 replicates. The values of

$\alpha = 1, \beta = 0.2, \rho = 0.8$ were used as the true values in generating the patient response. The skeleton probability used for the doses was $0.02, 0.05, 0.10, 0.16, 0.23$. It is assumed that a new patient is ready to be assigned a dose on the first cycle when the continuing patients in the trial have completed their dosing cycles. This assumption tallies with simulation settings in [Pye and Whitehead, 2012] and simplifies the number of iterations performed during the estimation process. An alternative to this assumption would be to induce an arrival time model via an Exponential distribution and has been excluded to present simplified results.

The following factors and questions were studied via simulation studies:

1. Differences in trial properties when using Plan 1 - Maximize the expected dose or Plan 2 - Match a regimen.

2. Selection of distance restriction R_j in executing Plan 2.

3. Whether to use both $\mathring{P}(B)$ and $\mathring{P}(C)$ or only $\mathring{P}(C)$ in trial conduct?

4. What values to use for $\mathring{P}(A)$, $\mathring{P}(B)$ and $\mathring{P}(C)$?

5. Whether to use $P^r(C) = 0.30$ or $P^r(A_1) = \mathring{P}(A_1)$ and $P^r(C) = 0.30$ in regimen selection at trial conclusion for recommended regimen selection?

6. Acceptable values for $P^r(C)$ and $P^r(A_1)$.

For different combinations of $\mathring{P}(A)$, $\mathring{P}(B)$ and $\mathring{P}(C)$ a total of 500 trials were simulated. Table 2.12 lists the results for the patient characteristics with the rows

grouped by the option used to select the regimen (option 1 or 2) at the conclusion of the trials executed using Plan 1, maximizing the expected dose. The rows are further grouped by $\mathring{P}(B) = (0.30, 0.40)$ within which each row correspond to cases of $\mathring{P}(A_1) = (0.05, 0.10, 0.20)$. The columns list the eight criteria presented in Section 2.4.7 for the patient characteristics. Tables 2.13 through 2.15 list the patient characteristics of trials executed when Plan 2 was in effect with $R_l \leq 3$, $R_l \leq 2$ and R_l set to the minimum distance possible. The results for the regimen characteristics are in Tables 2.16 through 2.19. The columns list the six criteria presented in Section 2.4.7 for the regimen characteristics. Results within each table are grouped by the two options for regimen selection and $\mathring{P}(B) = (0.30, 0.40)$ with the rows corresponding to $\mathring{P}(A_1) = (0.05, 0.10, 0.20)$. Some remarks are presented on the various aspects of the trial results.

Remark 1: Contrast using $\mathring{P}(A_1) = (0.05, 0.10, 0.20)$. Varying the cap on the first cycle affects the trial conduct in terms of which dose level is eligible on the first cycle. Lower values of $\mathring{P}(A_1) = (0.05, 0.10, 0.20)$ imply stringent rules for selecting doses with higher toxicities but also results in a less stringent rules on subsequent cycles due to the relationship between $\mathring{P}(A_1), \mathring{P}(A_2)$ and $\mathring{P}(C)$. Based on Tables 2.12 through 2.15 we notice that as the values of $\mathring{P}(A_1)$ increase the mean dose received by the patients in the trial also increases. Secondly the observed mean toxicities tend to be low when $\mathring{P}(A_1)$ is high, this could be because the toxicities on subsequent cycles

are averted due to the low $\mathring{P}(A_2)$ value and also because $\rho = 0.8$ in the simulations once the patients receive a high dose and survive it without a DLT, they are less likely to have a DLT on subsequent cycles. No clear trend is observed in the number of patients receiving cumulative dose greater than the recommended expected dose. The number of patients receiving the recommended regimen increases with increase in $\mathring{P}(A_1)$. The number of trials stopping early also decreases with increase in $\mathring{P}(A_1)$ suggesting that higher values of $\mathring{P}(A_1)$ allow patients to have dose assignment. In the case of the regimen characteristics the mean distance between the target (T) and the recommended (R) regimen increases as $\mathring{P}(A_1)$ increases. In general it seems that having a high value of $\mathring{P}(A_1) = 0.20$ gives rise to properties that are favored in terms of patients receiving a higher mean dose and recommended dose matching the target regimen selection more often.

Remark 2: Contrast the effect of distance in the regimen matching plan. Comparing Tables 2.13, 2.14 and 2.15 for differences in the patient characteristics results corresponding to distance less than 3, distance less than 2 and distance equal to the minimum possible value some differences are observed. In the first two instances when none of the regimens satisfy the distance criterion the algorithm switches to Plan 1 - maximizing the dose, while in the third instance the algorithm selects the dose that offers the the minimum distance from one of the 19 regimens. At low values of $\mathring{P}(A_1) = 0.05$ the difference in the number of patients matching the recommended regimen exactly increases drastically to 12.14 from 0.61 but not much of a differ-

ence when $\mathring{P}(A_1) = 0.20$. The Tables 2.17 and 2.19 do have major differences since they deal mostly with the regimen selection at the conclusion of the trial. We can conclude that patients having an exact match with the recommended regimen differs when using Plan 2 for low values of $\mathring{P}(A_1)$ and that the differences are minimal when $\mathring{P}(A_1)$ increases.

Remark 3: Contrast using $P^r(A_1)$ and $P^r(C)$ versus only $P^r(C)$ in selection of the regimens. These flags are concerned with the selection of the recommended and the target regimen at the end of the trial. Consider results within Table 2.12 and 2.13 for contrasting the effects of the regimen selection option. The number of patients having total dose higher than the recommended or the target expected dose is higher in option 1 when both the flags $P^r(A_1)$ and $P^r(C)$ are considered. The results in both the options are comparable in the instance when $\mathring{P}(A_1) = 0.20$ (the third row). Comparing the proportion of trials having an exact match between the target and the recommended regimen within Table 2.16 and 2.17 the proportions drop for restrictive $P^r(A_1) = 0.05, 0.10$ case while remain unchanged for the $P^r(A_1) = 0.20$ case. The expected dose of the regimens selected based on the single condition is higher than that when using both the conditions. The number of trials that cannot offer a recommended regimen reduces when using a single condition for the regimen selection. Having a restrictive condition on the first cycle plays a huge role in running of the trial and definitely affects the regimen selection.

Remark 4: Contrast the two plans of maximizing the dose and regimen matching. Comparing 2.12 and 2.15 for differences in the patient characteristics. The mean dose received is higher in Plan 1 when $\mathring{P}(A_1) = 0.05, 0.10$ and not much different from Plan 2 when $\mathring{P}(A_1) = 0.20$. The proportion of toxicities is slightly higher in Plan 1. The proportion of expected dose received in patients similar in both the plans but differs only when $\mathring{P}(A_1) = 0.20$ being much higher in plan 2. The number of patients having an exact match with the recommended regimen is higher in Plan 2. The mean of the distance between the true and recommended regimens is mostly higher in Plan 1 based on Tables 2.16 and 2.19. The results depict differences in trial properties between the two plans. Although Plan 1 offers higher mean dose the number of toxicities are also slightly higher. Plan 2 on the other hand offers a higher number of patients matching the recommended regimen and lower toxicities but tends to have lower mean dose per patient.

Remark 5: Is there a difference in using different values of $\mathring{P}(B) = (0.30, 0.40)$. Within each of the Tables 2.12 - 2.15 comparing the results of the mean dose and toxicity of the trials there does not seem to be huge differences in the results. It is not so obvious how the use of $\mathring{P}(B)$ affects the patient trial properties. In reference to the regimen characteristic Tables 2.16 through 2.19 there seem to be some slight differences in Plan 1 but no difference in the tables for regimen matching Plan 2. There does not seem to be any effect of using $\mathring{P}(B)$ during the trial conduct. To

further understand the effect of using both $\mathring{P}(C)$ and $\mathring{P}(B)$ versus only $\mathring{P}(C)$ in conducting the trial an additional set of simulations were carried out by assigning doses on cycles after the first cycle by calculating the probability of toxicity on the entire regimen and selecting the regimens that satisfy the $\mathring{P}(C)$ condition. In the case of Plan 2 R_l was set to the minimum distance possible. Simulations were carried out for a single setting of $\mathring{P}(A_1) = 0.20$ and $\mathring{P}(C) = 0.40$ and are presented in Table 2.20 which should be compared to the case when $\mathring{P}(A_1) = 0.20$ from results in Table 2.12 and Table 2.15. The mean dose and toxicities are comparable in both Plan 1 and Plan 2. The numbers are different in the case of patients having dose greater than the recommended dose, the patients with distance ≤ 2 is higher.

Remark 6: Contrast the use of differing values of $\mathring{P}(C) = (0.30, 0.40)$ *for fixed* $\mathring{P}(A_1) = 0.20$ *and absence of* $\mathring{P}(B)$. Based on Table 2.20 it was seen that having higher value of $\mathring{P}(C)$ has better trial properties in terms of higher mean dose. The mean number of toxicities increase but are still lower than 30%. The number of patients matching the recommended regimen are also higher when the $\mathring{P}(C)$ is higher.

2.5 Discussion

The Markov Model 2.1 presented in this chapter is simple in the sense that it allows for estimation of only three parameters and yet is capable of modeling the complex repeated data structure by accounting for the within-patient dose dependency through ρ. In the instance with larger amounts of data, more terms could be

added to the model that account for patient characteristics like gender or age but in the setting of small number of patients it might not be feasible to estimate the parameters especially in the early stages of the trial.

Besides the conditional nature, the major feature of the Markov Model 2.1 is its ability to allow patients surviving previous dose levels to have a lower probability of toxicity on subsequent cycles. The extension of the Markov Model 2.1 in carrying out a trial within the framework of safety criteria provides an excellent model based approach in designing adaptive clinical trials. The dose level selection considered in this chapter is complicated because of repeated measures aspect of the data and the choices that have to be made regarding skeleton probabilities, prior probabilities, escalation rules and safety criteria. The model presented and the simulations performed represent a framework for considering these issues. The results obtained apply to the specific situation that is being considered in the simulations for the probability skeleton, particular values of α,β and ρ. Further simulations need to be done to understand the working of the model in a broader framework. Thus while the specific choices we made are, we believe, reasonable, we do not claim they are optimal or necessarily appropriate in every conceivable context. But we do think that the framework and ideas are adequate and adaptable to match other contexts.

Through simulations we have demonstrated that the Markov model is able to estimate the conditional probabilities adequately in both small ($N = 10$) and larger

($N = 30$) sample sizes. The model also performs better in comparison to using a single binary endpoint and shows that that there are gains in estimation through use of all the data from all the patients in the trial. Simulation results have also shown that there are benefits in allowing patients to escalate or de-escalate dose levels to both the patients and the estimation of the parameters.

The simulation results for carrying out an adaptive clinical trial are affected by multiple factors and further investigation of how the parameters, safety rules, target probabilities and the regimen selection criteria on the selection of the doses for each of the patients in needed. The model could be extended to any number of dose levels and cycles although we have presented the model for five dose levels and six cycles.

The Markov model presented in this chapter is most relevant to clinical trials involving cytotoxic drugs where the toxicity is assumed to increase with the cumulative effect. Having non-delayed outcomes is also essential to the study design so that the DLT could be assigned at the end of the cycle to the appropriate dose level for the patient.

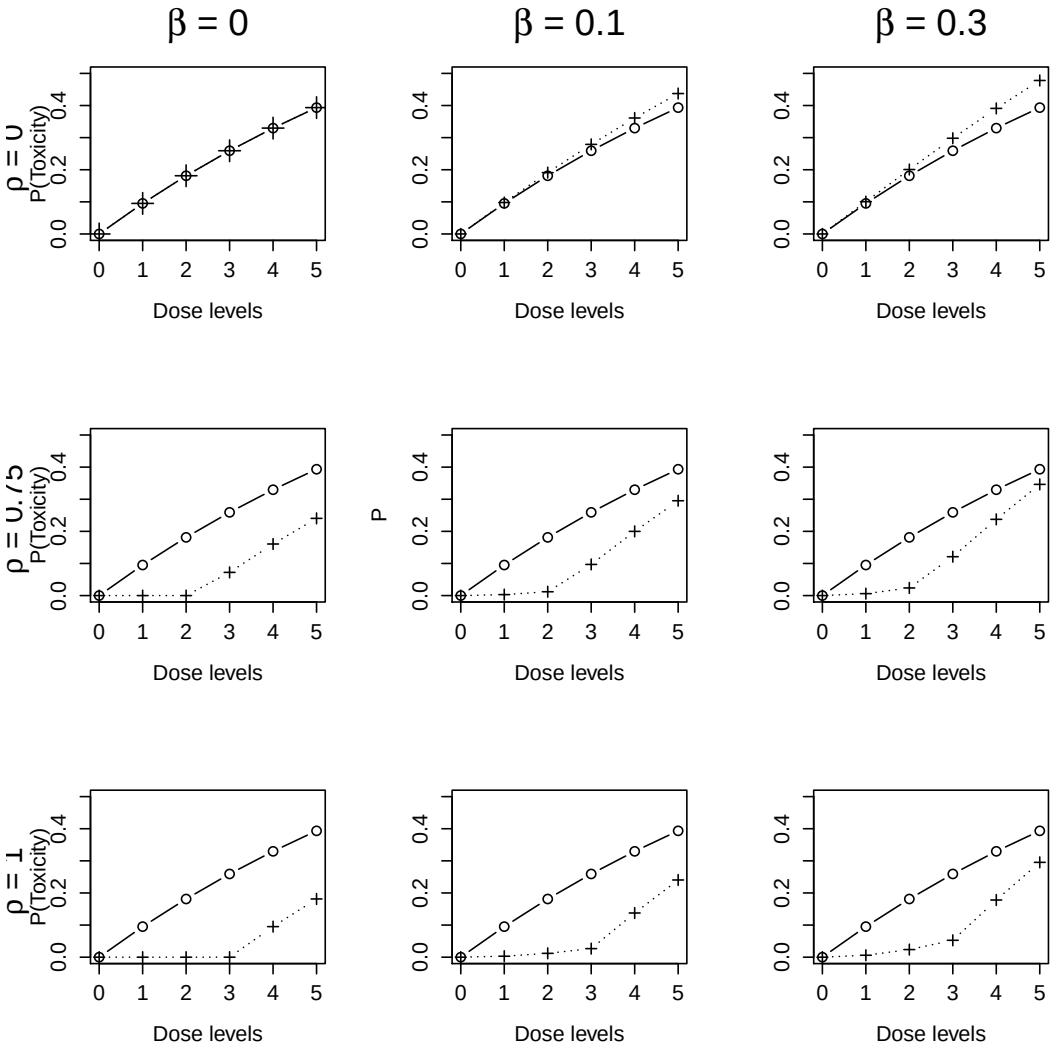

Figure 2.1: Conditional P(toxicity) with cycle 1 as the reference based on Model 2.1. Open circles depict probabilities for cycle 1 with $\alpha = 0.5$ across the five dose levels; crosses depict conditional probabilities of toxicity on cycle 2, assuming dose level three was administered on cycle 1 and one of five dose levels on the second cycle. Probabilities on cycle 2 are arranged by increasing β shown left to right and increasing ρ shown from top to bottom

Table 2.1: Table comparing the effect of priors in estimating the parameters using Model 2.1 from 500 simulated datasets containing $N = 30$ patients each receiving one of five dose groups $d_1 \ldots d_5$.

	True Value	Prior Mean (SD)	Estimate(Bias)	MSD[1]	ESD[2]	Coverage
Case 1						
α	1	1 (2)	0.957 (-0.043)	0.458	0.417	96.6
β	0.50	0.50 (1)	0.511 (0.011)	0.338	0.296	97.8
ρ	0.80	0.83 (0.14)	0.806 (0.006)	0.132	0.056	100
Case 2						
α	1	1 (5)	0.969 (-0.031)	0.477	0.455	94.6
β	0.50	0.50 (2.5)	0.464 (-0.036)	0.353	0.338	95.0
ρ	0.80	0.83 (0.14)	0.793 (-0.007)	0.136	0.057	100
Case 3						
α	0.80	1 (2)	0.795 (-0.005)	0.407	0.360	96.0
β	0.50	0.50 (1)	0.500 (< 0.001)	0.318	0.274	97.8
ρ	0.80	0.83 (0.14)	0.809 (0.009)	0.134	0.053	100
Case 4						
α	1	1 (2)	1.013 (0.013)	0.477	0.450	95.4
β	0.80	0.50 (1)	0.719 (-0.081)	0.450	0.373	95.2
ρ	0.80	0.83 (0.14)	0.793 (-0.007)	0.140	0.055	100

[1] MSD is mean of the SD from 500 estimates

[2] ESD is empirical SD of the 500 estimates

Table 2.2: Table comparing the effect of priors in estimating the conditional probability of toxicity (with bias from the true value in parenthesis) using Model 2.1 from 500 simulated datasets containing $N = 30$ patients each receiving one of five dose groups $d_1 \dots d_5$. Case 1 and 2 use $\alpha = 1, \beta = 0.5, \rho = 0.8$, while Case 3 use $\alpha = 0.8, \beta = 0.5, \rho = 0.8$ and Case 4 uses $\alpha = 1, \beta = 0.8, \rho = 0.8$ to generate the datasets. The priors have means of $\alpha = 1, \beta = 0.5, \rho = 0.8$ and $SD = 2$ in Cases 1, 3 and 4 and $SD = 5$ in Case 2.

	Cycle 1	Cycle 2	Cycle 6	Any Cycle
	Est(bias)	Est(bias)	Est(bias)	Est(bias)
Case 1				
d_1	0.019 (-0.001)	0.004 (< 0.001)	0.005 (< 0.001)	0.041 (-0.002)
d_2	0.048 (-0.002)	0.011 (< 0.001)	0.016 (< 0.001)	0.111 (-0.004)
d_3	0.095 (-0.005)	0.025 (-0.001)	0.047 (< 0.001)	0.246 (-0.008)
d_4	0.152 (-0.008)	0.048 (-0.001)	0.104 (-0.001)	0.425 (-0.014)
d_5	0.217 (-0.013)	0.081 (-0.002)	0.197 (-0.002)	0.623 (-0.022)
Case 2				
d_1	0.019 (-0.001)	0.004 (< 0.001)	0.005 (< 0.001)	0.041 (-0.002)
d_2	0.048 (-0.002)	0.011 (< 0.001)	0.016 (-0.001)	0.110 (-0.005)
d_3	0.096 (-0.004)	0.025 (-0.001)	0.045 (-0.003)	0.242 (-0.013)
d_4	0.153 (-0.007)	0.047 (-0.002)	0.098 (-0.007)	0.413 (-0.025)
d_5	0.218 (-0.012)	0.079 (-0.004)	0.184 (-0.016)	0.604 (-0.040)
Case 3				
d_1	0.016 (< 0.001)	0.003 (< 0.001)	0.004 (< 0.001)	0.034 (-0.001)
d_2	0.040 (< 0.001)	0.009 (< 0.001)	0.014 (< 0.001)	0.094 (-0.003)
d_3	0.080 (-0.001)	0.021 (-0.001)	0.043 (-0.001)	0.216 (-0.007)
d_4	0.128 (-0.002)	0.041 (-0.002)	0.096 (-0.002)	0.385 (-0.013)
d_5	0.184 (-0.004)	0.071 (-0.003)	0.186 (-0.005)	0.583 (-0.023)
Case 4				
d_1	0.020 (< 0.001)	0.005 (< 0.001)	0.006 (< 0.001)	0.045 (0.001)
d_2	0.050 (< 0.001)	0.013 (< 0.001)	0.020 (-0.001)	0.125 (-0.001)
d_3	0.100 (< 0.001)	0.03 (< 0.001)	0.060 (-0.003)	0.282 (-0.009)
d_4	0.159 (-0.001)	0.057 (-0.001)	0.134 (-0.010)	0.487 (-0.023)
d_5	0.227 (-0.003)	0.098 (-0.003)	0.254 (-0.024)	0.700 (-0.039)

Table 2.3: Table with the 19 favorable dose regimen combinations over the six cycles. Each of the row regimens indicate the dose level assigned on corresponding cycle.

Cycle	1	2	3	4	5	6
Regimen 1	1	1	1	1	1	1
Regimen 2	2	2	2	2	2	2
Regimen 3	3	3	3	3	3	3
Regimen 4	4	4	4	4	4	4
Regimen 5	5	5	5	5	5	5
Regimen 6	1	1	1	2	2	2
Regimen 7	2	2	2	3	3	3
Regimen 8	3	3	3	4	4	4
Regimen 9	4	4	4	5	5	5
Regimen 10	2	2	2	1	1	1
Regimen 11	3	3	3	2	2	2
Regimen 12	4	4	4	3	3	3
Regimen 13	5	5	5	4	4	4
Regimen 14	1	1	2	2	3	3
Regimen 15	2	2	3	3	4	4
Regimen 16	3	3	4	4	5	5
Regimen 17	5	5	4	4	3	3
Regimen 18	4	4	3	3	2	2
Regimen 19	3	3	2	2	1	1

Table 2.4: Parameter estimates from two different patient profiles, based on Model 2.1 from 500 simulated datasets containing $N = 30$ patients receiving one of five dose groups $d_1 \dots d_5$ over six cycles. Results presented for $\alpha = 1, \beta = 0.5$ and $\rho = 0.8$

	Estimate(Bias)	MSD[1]	ESD [2]	Coverage
Section 2.3.2				
α	0.989 (-0.011)	0.446	0.425	95.2
β	0.561 (0.061)	0.361	0.322	98
ρ	0.808 (0.008)	0.134	0.056	100
Section 2.3.1				
α	0.957 (-0.043)	0.458	0.417	96.6
β	0.511 (0.011)	0.338	0.296	97.8
ρ	0.806 (0.006)	0.132	0.056	100

[1] MSD is mean of the SD from 500 estimates
[2] ESD is empirical SD of the 500 estimates

Table 2.5: Estimates with bias from the true value of conditional probability of toxicity on each of the six cycles estimated using Model 2.1 from 500 simulated datasets containing $N = 30$ patients with dose variability in Section 2.3.2 and same dose level in Section 2.3.1, each receiving one of five dose groups $d_1 \dots d_5$. Results presented for $\alpha = 1, \beta = 0.5$ and $\rho = 0.8$

	Cycle 1		Cycle 2		Cycle 6		Any Cycle	
Section 2.3.2	Est(bias)	ESD[1]	Est(bias)	ESD[1]	Est(bias)	ESD[1]	Est(bias)	ESD[1]
d_1	0.020 (< 0.001)	0.008	0.004 (< 0.001)	0.002	0.005 (< 0.001)	0.002	0.041 (-0.001)	0.016
d_2	0.049 (-0.001)	0.021	0.011 (-0.001)	0.005	0.017 (< 0.001)	0.006	0.113 (-0.002)	0.037
d_3	0.098 (-0.002)	0.04	0.026 (-0.001)	0.010	0.049 (0.002)	0.018	0.253 (-0.001)	0.067
d_4	0.156 (-0.004)	0.06	0.048 (-0.001)	0.017	0.110 (0.005)	0.042	0.437 (-0.001)	0.098
d_5	0.223 (-0.007)	0.082	0.083 (< 0.001)	0.027	0.209 (0.009)	0.080	0.640 (-0.005)	0.113
Section 2.3.1								
d_1	0.019 (-0.001)	0.008	0.004 (< 0.001)	0.002	0.005 (< 0.001)	0.002	0.041 (-0.002)	0.016
d_2	0.048 (-0.002)	0.020	0.011 (< 0.001)	0.004	0.016 (< 0.001)	0.005	0.111 (-0.004)	0.036
d_3	0.095 (-0.005)	0.039	0.025 (-0.001)	0.008	0.047 (< 0.001)	0.016	0.246 (-0.008)	0.065
d_4	0.152 (-0.008)	0.060	0.048 (-0.001)	0.014	0.104 (-0.001)	0.038	0.425 (-0.014)	0.093
d_5	0.217 (-0.013)	0.081	0.081 (-0.002)	0.023	0.197 (-0.002)	0.073	0.623 (-0.022)	0.108

[1] ESD is empirical SD of the 500 estimates

Table 2.6: Parameter estimates from 500 simulated datasets containing $N = 10$ and $N = 30$ patients receiving one of five dose groups $d_1 \ldots d_5$ over six cycles comparing comparing estimates based on Model 2.1 to estimates based on Models 2.7 and 2.8 and having **equal** patients on all the dose levels.

	Estimate(Bias)	MSD[1]	ESD[2]	Coverage
N=10 patients				
α	0.945 (0.055)	0.682	0.585	99.6
β	0.515 (-0.015)	0.464	0.308	100
ρ	0.820 (-0.020)	0.137	0.040	100
γ	0.916 (0.084)	0.703	0.630	99.8
δ	0.976 (0.024)	0.548	0.536	92.6
N=30 patients				
α	0.989 (0.011)	0.466	0.408	96.0
β	0.507 (-0.007)	0.343	0.285	98.8
ρ	0.802 (-0.002)	0.133	0.058	99.8
γ	0.940 (0.060)	0.474	0.428	96.6
δ	0.998 (0.002)	0.333	0.331	94

[1] MSD is mean of the SD from 500 estimates

[2] ESD is empirical SD of the 500 estimates

Table 2.7: Probability of toxicity estimates based on 500 simulated datasets containing $N = 10$ and $N = 30$ patients receiving one of five dose groups $d_1 \dots d_5$ over six cycles with α, β, ρ based on Model 2.1, γ from Model 2.7, δ from Model 2.8. Results presented for $\alpha = 1, \beta = 0.5$ and $\rho = 0.8$ and with **equal** patients on the five dose levels.

	$N = 10$		$N = 30$	
	Estimate(Bias)	ESD[1]	Estimate(Bias)	ESD[1]
First cycle				
Markov Model 2.1				
d_1	0.019 (-0.001)	0.011	0.020 (< 0.001)	0.008
d_2	0.047 (-0.003)	0.028	0.049 (-0.001)	0.020
d_3	0.093 (-0.007)	0.053	0.098 (-0.002)	0.038
d_4	0.148 (-0.012)	0.081	0.156 (-0.004)	0.059
d_5	0.21 (-0.02)	0.109	0.223 (-0.007)	0.081
Binary Model 2.7				
d_1	0.018 (-0.002)	0.012	0.019 (-0.001)	0.008
d_2	0.045 (-0.005)	0.030	0.047 (-0.003)	0.021
d_3	0.090 (-0.010)	0.057	0.093 (-0.007)	0.040
d_4	0.143 (-0.017)	0.086	0.149 (-0.011)	0.062
d_5	0.203 (-0.027)	0.114	0.213 (-0.017)	0.085
Any cycle				
Markov Model 2.1				
d_1	0.038 (-0.004)	0.021	0.042 (-0.001)	0.016
d_2	0.105 (-0.011)	0.050	0.113 (-0.002)	0.038
d_3	0.234 (-0.021)	0.092	0.250 (-0.005)	0.068
d_4	0.405 (-0.033)	0.132	0.428 (-0.010)	0.095
d_5	0.598 (-0.046)	0.152	0.626 (-0.019)	0.109
Binary Model 2.8				
d_1	0.041 (-0.001)	0.022	0.042 (< 0.001)	0.014
d_2	0.111 (-0.004)	0.055	0.114 (-0.001)	0.035
d_3	0.241 (-0.014)	0.107	0.251 (-0.004)	0.070
d_4	0.407 (-0.032)	0.153	0.428 (-0.011)	0.102
d_5	0.590 (-0.055)	0.176	0.624 (-0.02)	0.115

[1] ESD is empirical SD of the 500 estimates

Table 2.8: Parameter estimates from 500 simulated datasets containing $N = 10$ and $N = 30$ patients receiving one of five dose groups $d_1 \dots d_5$ over six cycles with α, β, ρ based on Model 2.1, γ from Model 2.7, δ from Model 2.8 and having **unequal** patients on all the dose levels.

	Estimate(Bias)	MSD[1]	ESD[2]	Coverage
N=10 patients				
α	0.965 (0.035)	0.688	0.566	99.8
β	0.530 (-0.030)	0.489	0.340	100
ρ	0.818 (-0.018)	0.137	0.046	100
γ	0.924 (0.076)	0.704	0.590	99.8
δ	0.995 (0.005)	0.543	0.570	94.6
N=30 patients				
α	0.986 (0.014)	0.446	0.398	95.8
β	0.487 (0.013)	0.329	0.236	99.0
ρ	0.803 (-0.003)	0.131	0.058	100
γ	0.945 (0.055)	0.455	0.419	92.2
δ	0.980 (0.020)	0.311	0.299	95.6

[1] MSD is mean of the SD from 500 estimates

[2] ESD is empirical SD of the 500 estimates

Table 2.9: Probability of toxicity estimates based on 500 simulated datasets containing $N = 10$ and $N = 30$ patients receiving one of five dose groups $d_1 \dots d_5$ over six cycles comparing estimates based on Model 2.1 to estimates based on Models 2.7 and 2.8. Results presented for $\alpha = 1, \beta = 0.5$ and $\rho = 0.8$ and with **unequal** patients on the five dose levels.

	$N = 10$		$N = 30$	
	Estimate(Bias)	ESD[1]	Estimate(Bias)	ESD[1]
First cycle				
Markov Model 2.1				
d_1	0.019(-0.001)	0.011	0.02(< 0.001)	0.008
d_2	0.048(-0.002)	0.027	0.049(-0.001)	0.019
d_3	0.095(-0.005)	0.052	0.098(-0.002)	0.037
d_4	0.151(-0.009)	0.079	0.156(-0.004)	0.057
d_5	0.215(-0.015)	0.106	0.223(-0.007)	0.078
Binary Model 2.7				
d_1	0.018(-0.002)	0.012	0.019(-0.001)	0.008
d_2	0.046(-0.004)	0.028	0.047(-0.003)	0.02
d_3	0.091(-0.009)	0.054	0.094(-0.006)	0.039
d_4	0.144(-0.016)	0.082	0.15(-0.010)	0.06
d_5	0.206(-0.024)	0.11	0.214(-0.016)	0.083
Any cycle				
Markov Model 2.1				
d_1	0.04(-0.003)	0.022	0.041(-0.001)	0.016
d_2	0.108(-0.007)	0.053	0.112(-0.003)	0.036
d_3	0.24(-0.015)	0.097	0.246(-0.008)	0.064
d_4	0.413(-0.025)	0.137	0.423(-0.016)	0.088
d_5	0.605(-0.04)	0.156	0.621(-0.024)	0.101
Binary Model 2.8				
d_1	0.042(< 0.001)	0.023	0.042(-0.001)	0.012
d_2	0.113(-0.003)	0.058	0.112(-0.003)	0.032
d_3	0.244(-0.011)	0.112	0.247(-0.007)	0.064
d_4	0.41(-0.028)	0.158	0.424(-0.015)	0.094
d_5	0.592(-0.052)	0.179	0.621(-0.024)	0.107

[1] ESD is empirical SD of the 500 estimates

Table 2.10: Target probability bound choices for $\mathring{P}(A_2)$ assuming bounds on $\mathring{P}(A_1)$ and $\mathring{P}(C)$.

	$\mathring{P}(C)$	
$\mathring{P}(A_1)$	0.30	0.40
0.05	0.0592	0.0878
0.10	0.0490	0.0779
0.15	0.0381	0.0673
0.20	0.0264	0.0559

Table 2.11: Table listing the probability of toxicity on the first cycle $P(A_1)$, the probability of toxicity on any cycle $P(C)$, the expected dose on the entire regimen, flags set to one if the regimen qualifies when using only option 1, $P^r(C) = 0.30$ and three instances of using option 2, $P^r(C) = 0.30$ and $P^r(A_1) = (0.05, 0.10, 0.20)$ corresponding to the 19 favorable dose regimen listed in Table 2.3. $\alpha = 1, \beta = 0.2, \rho = 0.8$ are the true values of the parameters.

Reg ID	Regimen						$P(A_1)$	$P(C)$	Exp Dose	$P^r(C) \leq 0.30$	$P^r(A_1) \leq 0.05$	$P^r(A_1) \leq 0.10$	$P^r(A_1) \leq 0.20$
1	1	1	1	1	1	1	0.02	0.04	5.86	1	1	1	1
2	2	2	2	2	2	2	0.05	0.10	11.29	1	1	1	1
3	3	3	3	3	3	3	0.10	0.22	15.84	1	0	1	1
4	4	4	4	4	4	4	0.16	0.36	19.33	0	0	0	0
5	5	5	5	5	5	5	0.23	0.52	21.54	0	0	0	0
6	1	1	1	2	2	2	0.02	0.08	8.63	1	1	1	1
7	2	2	2	3	3	3	0.05	0.18	13.71	1	1	1	1
8	3	3	3	4	4	4	0.10	0.31	17.81	0	0	0	0
9	4	4	4	5	5	5	0.16	0.46	20.76	0	0	0	0
10	2	2	2	1	1	1	0.05	0.07	8.57	1	1	1	1
11	3	3	3	2	2	2	0.10	0.15	13.46	1	0	1	1
12	4	4	4	3	3	3	0.16	0.26	**17.43**	1	0	0	1
13	5	5	5	4	4	4	0.23	0.40	20.21	0	0	0	0
14	1	1	2	2	3	3	0.02	0.15	11.22	1	1	1	1
15	2	2	3	3	4	4	0.05	0.26	**15.89**	1	1	1	1
16	3	3	4	4	5	5	0.10	0.41	19.48	0	0	0	0
17	5	5	4	4	3	3	0.23	0.34	18.67	0	0	0	0
18	4	4	3	3	2	2	0.16	0.22	15.34	1	0	0	1
19	3	3	2	2	1	1	0.10	0.13	10.96	1	0	1	1

Table 2.12: Trial/Patient summary results for Plan 1, maximizing expected dose over 500 simulated clinical trials. Columns present 1) mean dose received per patient, 2) mean toxicities over the trials with trial stopping early in parenthesis, 3) average patients having dose $\geq$ the recommended expected dose, 4) average patients having dose $\geq$ the target expected dose, 5) average patients having regimen exactly equal to the recommended regimen 6) average patients having regimen exactly equal to the target regimen, 7) average patients having regimen ≤ 2 from the recommended regimen and 8) average patients having regimen ≤ 2 from the target regimen.

	Mean Dose	Mean Toxicities*	Patients $\geq$R.Edose[a]	Patients $\geq$T.Edose[a]	Patients =R.Dist[b]	Patients =T.Dist[b]	P.R.Dist $\leq$2	P.T.Dist ≤ 2
Using $\boldsymbol{P^r(A_1)\&P^r(C)=0.30}$								
$\boldsymbol{\breve{P}(B)=0.3, \breve{P}(C)=0.4}$								
$\breve{P}(A_1)=0.05, \breve{P}(A_2)=0.09$	15.48	0.31 (48)	20.27	15.02	0.62	0.83	5.66	5.35
$\breve{P}(A_1)=0.1, \breve{P}(A_2)=0.08$	15.91	0.30 (25)	18.09	18.09	1.64	0.64	11.75	4.89
$\breve{P}(A_1)=0.2, \breve{P}(A_2)=0.06$	16.21	0.27 (11)	16.73	17.22	3.96	1.83	10.86	8.57
$\boldsymbol{\breve{P}(B)=0.4, \breve{P}(C)=0.4}$								
$\breve{P}(A_1)=0.05, \breve{P}(A_2)=0.09$	15.53	0.31 (48)	20.28	14.96	0.60	0.83	5.40	5.24
$\breve{P}(A_1)=0.1, \breve{P}(A_2)=0.08$	15.99	0.31 (25)	17.98	17.68	1.59	0.69	11.3	4.67
$\breve{P}(A_1)=0.2, \breve{P}(A_2)=0.06$	16.33	0.28 (11)	16.82	17.38	4.46	2.16	11.26	9.31
Using $\boldsymbol{P^r(C)=0.30}$								
$\boldsymbol{\breve{P}(B)=0.3, \breve{P}(C)=0.4}$								
$\breve{P}(A_1)=0.05, \breve{P}(A_2)=0.09$	15.48	0.31 (48)	12.84	12.71	0.45	0.13	8.15	2.09
$\breve{P}(A_1)=0.1, \breve{P}(A_2)=0.08$	15.91	0.3 (25)	15.46	13.51	1.81	0.68	11.09	4.95
$\breve{P}(A_1)=0.2, \breve{P}(A_2)=0.06$	16.21	0.27 (11)	16.73	17.22	3.95	1.83	10.81	8.57
$\boldsymbol{\breve{P}(B)=0.4, \breve{P}(C)=0.4}$								
$\breve{P}(A_1)=0.05, \breve{P}(A_2)=0.09$	15.53	0.31 (48)	12.81	12.77	0.43	0.12	7.68	2.22
$\breve{P}(A_1)=0.1, \breve{P}(A_2)=0.08$	15.99	0.31 (25)	15.29	13.37	1.72	0.65	10.71	5.44
$\breve{P}(A_1)=0.2, \breve{P}(A_2)=0.06$	16.33	0.28 (11)	16.82	17.38	4.47	2.16	11.24	9.31

* Values in parenthesis indicates the number of trials out of 500 that stopped early

[a] T/R.Edose - Target/Recommended expected dose.

[b] T/R.Dist - Distance of patient regimen from Target/Recommended regimen.

Table 2.13: Trial/Patient summary results for Plan 2- matching a regimen with $R_l \leq 3$ over 500 simulated clinical trials.

	Mean Dose	Mean Toxicities*	Patients ≥R.Edose[a]	Patients ≥T.Edose[a]	Patients =R.Dist[b]	Patients =T.Dist[b]	P.R.Dist ≤2	P.T.Dist ≤ 2
Using $\boldsymbol{P^r(A_1)}$&$\boldsymbol{P^r(C) = 0.30}$								
$\boldsymbol{\mathring{P}(B) = 0.3, \mathring{P}(C) = 0.4}$								
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	15.57	0.31 (48)	20.18	15.02	0.61	0.82	5.52	5.18
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	16.07	0.31 (25)	17.98	17.8	1.44	0.64	10.85	4.7
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	17.33	0.31 (11)	19.75	18.98	4.57	3.33	12.83	11.81
$\boldsymbol{\mathring{P}(B) = 0.4, \mathring{P}(C) = 0.4}$								
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	15.55	0.31 (48)	20.28	15.03	0.6	0.82	5.36	5.21
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	16.10	0.31 (25)	17.89	17.88	1.39	0.63	10.89	4.71
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	17.34	0.31 (11)	19.76	19.00	4.64	3.36	12.94	11.85
Using $\boldsymbol{P^r(C) = 0.30}$								
$\boldsymbol{\mathring{P}(B) = 0.3, \mathring{P}(C) = 0.4}$								
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	15.57	0.31 (48)	12.85	12.92	0.42	0.11	7.47	2.36
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	16.07	0.31 (25)	15.60	13.71	1.87	1.08	10.74	6.13
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	17.33	0.31 (11)	19.75	18.98	4.7	3.33	12.99	11.81
$\boldsymbol{\mathring{P}(B) = 0.4, \mathring{P}(C) = 0.4}$								
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	15.55	0.31 (48)	13	12.88	0.42	0.11	7.49	2.36
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	16.10	0.31 (25)	15.62	13.79	1.85	1.09	10.76	6.10
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	17.34	0.31 (11)	19.76	19.00	4.73	3.36	13.04	11.85

* Values in parenthesis indicates the number of trials out of 500 that stopped early

[a] T/R.Edose - Target/Recommended expected dose.

[b] T/R.Dist - Distance of patient regimen from Target/Recommended regimen.

Table 2.14: Trial/Patient summary results for Plan 2- matching a regimen with $R_l \leq 2$ instead of 3 over 500 simulated clinical trials.

	Mean Dose	Mean Toxicities*	Patients ≥R.Edose[a]	Patients ≥T.Edose[a]	Patients =R.Dist[b]	Patients =T.Dist[b]	P.R.Dist ≤2	P.T.Dist ≤ 2
Using $\boldsymbol{P^r(A_1)}$&$\boldsymbol{P^r(C) = 0.30}$								
$\boldsymbol{\mathring{P}(B) = 0.3, \mathring{P}(C) = 0.4}$								
$\mathring{P}(A_1) = 0.05, \breve{P}(A_2) = 0.09$	15.54	0.31 (48)	20.4	14.99	0.58	0.82	5.31	5.21
$\mathring{P}(A_1) = 0.1, \breve{P}(A_2) = 0.08$	16.11	0.31 (25)	17.97	17.88	1.41	0.63	10.97	4.71
$\mathring{P}(A_1) = 0.2, \breve{P}(A_2) = 0.06$	17.34	0.31 (11)	19.74	18.99	4.65	3.35	12.86	11.82
$\boldsymbol{\mathring{P}(B) = 0.4, \mathring{P}(C) = 0.4}$								
$\mathring{P}(A_1) = 0.05, \breve{P}(A_2) = 0.09$	15.54	0.31 (48)	20.2	14.93	0.60	0.84	5.38	5.2
$\mathring{P}(A_1) = 0.1, \breve{P}(A_2) = 0.08$	16.08	0.31 (25)	17.91	17.84	1.40	0.63	10.82	4.71
$\mathring{P}(A_1) = 0.2, \breve{P}(A_2) = 0.06$	17.34	0.31 (11)	19.73	19.00	4.60	3.37	12.81	11.84
Using $\boldsymbol{P^r(C) = 0.30}$								
$\boldsymbol{\breve{P}(B) = 0.3, \breve{P}(C) = 0.4}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.09$	15.54	0.31 (48)	12.83	12.88	0.41	0.12	7.63	2.36
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.08$	16.11	0.31 (25)	15.7	13.79	1.81	1.09	10.64	6.07
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.06$	17.34	0.31 (11)	19.73	18.99	4.75	3.35	12.98	11.82
$\boldsymbol{\breve{P}(B) = 0.4, \breve{P}(C) = 0.4}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.09$	15.54	0.31 (48)	12.75	12.83	0.42	0.12	7.47	2.36
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.08$	16.08	0.31 (25)	15.58	13.74	1.81	1.08	10.68	6.11
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.06$	17.34	0.31 (11)	19.73	19.00	4.70	3.37	12.91	11.84

* Values in parenthesis indicates the number of trials out of 500 that stopped early

[a] T/R.Edose - Target/Recommended expected dose.

[b] T/R.Dist - Distance of patient regimen from Target/Recommended regimen.

Table 2.15: Trial/Patient summary results for Plan 2- matching a regimen with $\min(R_l)$ over 500 simulated clinical trials.

	Mean Dose	Mean Toxicities*	Patients ≥R.Edose[a]	Patients ≥T.Edose[a]	Patients =R.Dist[b]	Patients =T.Dist[b]	P.R.Dist ≤2	P.T.Dist ≤ 2
Using $\boldsymbol{P^r(A_1), P^r(C) = 0.30}$								
$\boldsymbol{\check{P}(B) = 0.3, \check{P}(C) = 0.4}$								
$\check{P}(A_1) = 0.05, \check{P}(A_2) = 0.09$	13.91	0.26 (48)	17.28	15.12	12.14	10.82	18.06	15.31
$\check{P}(A_1) = 0.1, \check{P}(A_2) = 0.08$	15.15	0.27 (25)	15.91	17.84	2.84	2.40	11.75	9.03
$\check{P}(A_1) = 0.2, \check{P}(A_2) = 0.06$	17	0.30 (11)	19.2	18.76	4.41	3.46	10.96	10.30
$\boldsymbol{\check{P}(B) = 0.4, \check{P}(C) = 0.4}$								
$\check{P}(A_1) = 0.05, \check{P}(A_2) = 0.09$	13.92	0.27 (50)	17.21	15.53	11.81	10.73	18.69	15.72
$\check{P}(A_1) = 0.1, \check{P}(A_2) = 0.08$	15.2	0.28 (25)	16.12	18.01	2.42	2.42	12.26	9.27
$\check{P}(A_1) = 0.2, \check{P}(A_2) = 0.06$	17.01	0.30 (11)	19.24	18.79	4.61	3.48	11.24	10.40
Using $\boldsymbol{P^r(C) = 0.30}$								
$\boldsymbol{\check{P}(B) = 0.3, \check{P}(C) = 0.4}$								
$\check{P}(A_1) = 0.05, \check{P}(A_2) = 0.09$	13.91	0.26 (48)	7.91	11.58	2.08	0.15	7.31	1.88
$\check{P}(A_1) = 0.1, \check{P}(A_2) = 0.08$	15.15	0.27 (25)	13.64	13.49	3.24	1.30	10.33	4.81
$\check{P}(A_1) = 0.2, \check{P}(A_2) = 0.06$	17	0.3 (11)	19.19	18.76	4.54	3.46	11.03	10.30
$\boldsymbol{\check{P}(B) = 0.4, \check{P}(C) = 0.4}$								
$\check{P}(A_1) = 0.05, \check{P}(A_2) = 0.09$	13.92	0.27 (50)	8.03	11.49	1.60	0.14	7.50	1.95
$\check{P}(A_1) = 0.1, \check{P}(A_2) = 0.08$	15.2	0.28 (25)	13.95	13.39	2.98	1.32	10.73	4.99
$\check{P}(A_1) = 0.2, \check{P}(A_2) = 0.06$	17.01	0.30 (11)	19.22	18.79	4.72	3.48	11.3	10.40

* Values in parenthesis indicates the number of trials out of 500 that stopped early

[a] T/R.Edose - Target/Recommended expected dose.

[b] T/R.Dist - Distance of patient regimen from Target/Recommended regimen.

Table 2.16: Regimen recommendation summary for Plan 1 -maximizing the expected dose over 500 simulated clinical trials. Columns correspond to 1) the mean distance between the target (T) and recommended (R) regimen, 2) the proportion of trials have an exact match between the T and R regimens, 3) the expected dose for the T regimen, 4) the expected dose for the R regimen, 5) the probability of toxicity on any cycle under the T regimen and 6) the probability of toxicity on any cycle under the R regimen.

	Mean(T-R)[a] Distance[*]	Prop of T-R=0[*]	Target Exp Dose[*]	Rec Exp Dose[*]	Target Toxicities	Rec Toxicities
Using $\boldsymbol{P^r(A_1)\&P^r(C)=0.30}$						
$\boldsymbol{\breve{P}(B)=0.3, \breve{P}(C)=0.4}$						
$\breve{P}(A_1)=0.05, \breve{P}(A_2)=0.09$	2.77 (452)	0.51 (452)	15.89 (500)	15 (452)	0.26	0.22
$\breve{P}(A_1)=0.1, \breve{P}(A_2)=0.08$	2.92 (475)	0.25 (475)	15.89 (500)	16.25 (475)	0.26	0.26
$\breve{P}(A_1)=0.2, \breve{P}(A_2)=0.06$	3.99 (489)	0.22 (489)	17.43 (500)	17.03 (489)	0.26	0.28
$\boldsymbol{\breve{P}(B)=0.4, \breve{P}(C)=0.4}$						
$\breve{P}(A_1)=0.05, \breve{P}(A_2)=0.09$	2.84 (452)	0.51 (452)	15.89 (500)	15 (452)	0.26	0.22
$\breve{P}(A_1)=0.1, \breve{P}(A_2)=0.08$	2.96 (475)	0.23 (475)	15.89 (500)	16.18 (475)	0.26	0.26
$\breve{P}(A_1)=0.2, \breve{P}(A_2)=0.06$	4.04 (489)	0.18 (489)	17.43 (500)	17.17 (489)	0.26	0.28
Using $\boldsymbol{P^r(C)=0.30}$						
$\boldsymbol{\breve{P}(B)=0.3, \breve{P}(C)=0.4}$						
$\breve{P}(A_1)=0.05, \breve{P}(A_2)=0.09$	4.70 (489)	0.16 (489)	17.43 (500)	17.11 (489)	0.26	0.28
$\breve{P}(A_1)=0.1, \breve{P}(A_2)=0.08$	4.32 (489)	0.21 (489)	17.43 (500)	16.82 (489)	0.26	0.28
$\breve{P}(A_1)=0.2, \breve{P}(A_2)=0.06$	3.96 (489)	0.22 (489)	17.43 (500)	17.04 (489)	0.26	0.28
$\boldsymbol{\breve{P}(B)=0.4, \breve{P}(C)=0.4}$						
$\breve{P}(A_1)=0.05, \breve{P}(A_2)=0.09$	4.61 (489)	0.18 (489)	17.43 (500)	17.14 (489)	0.26	0.28
$\breve{P}(A_1)=0.1, \breve{P}(A_2)=0.08$	4.22 (489)	0.19 (489)	17.43 (500)	16.79 (489)	0.26	0.28
$\breve{P}(A_1)=0.2, \breve{P}(A_2)=0.06$	3.99 (489)	0.18 (489)	17.43 (500)	17.19 (489)	0.26	0.28

[*] Values in parenthesis indicates the number of trials out of 500 that had a regimen selection
[a] T - Target/True regimen, R - Recommended regimen

Table 2.17: Regimen recommendation summary for Plan 2 - matching the regimen with $R_l \leq 3$ over 500 simulated clinical trials.

	Mean(T-R)[a] Distance[*]	Prop of T-R=0[*]	Target Exp Dose[*]	Rec Exp Dose[*]	Target Toxicities	Rec Toxicities
Using $\boldsymbol{P^r(A_1)}$&$\boldsymbol{P^r(C) = 0.30}$						
$\boldsymbol{\mathring{P}(B) = 0.3, \mathring{P}(C) = 0.4}$						
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	2.84 (452)	0.5 (452)	15.89 (500)	15.09 (452)	0.26	0.22
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	2.86 (475)	0.25 (475)	15.89 (500)	16.21 (475)	0.26	0.26
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	3.85 (489)	0.22 (489)	17.43 (500)	17.21 (489)	0.26	0.28
$\boldsymbol{\mathring{P}(B) = 0.4, \mathring{P}(C) = 0.4}$						
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	2.83 (452)	0.5 (452)	15.89 (500)	15.02 (452)	0.26	0.22
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	2.84 (475)	0.24 (475)	15.89 (500)	16.26 (475)	0.26	0.26
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	3.84 (489)	0.22 (489)	17.43 (500)	17.24 (489)	0.26	0.28
Regimen selected using $\boldsymbol{P^r(C)}$						
$\boldsymbol{\mathring{P}(B) = 0.3, \mathring{P}(C) = 0.4}$						
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	4.64 (489)	0.17 (489)	17.43 (500)	17.16 (489)	0.26	0.28
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	4.24 (489)	0.21 (489)	17.43 (500)	16.78 (489)	0.26	0.28
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	3.78 (489)	0.22 (489)	17.43 (500)	17.24 (489)	0.26	0.28
$\boldsymbol{\mathring{P}(B) = 0.4, \mathring{P}(C) = 0.4}$						
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	4.65 (489)	0.17 (489)	17.43 (500)	17.12 (489)	0.26	0.28
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	4.23 (489)	0.21 (489)	17.43 (500)	16.83 (489)	0.26	0.28
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	3.76 (489)	0.22 (489)	17.43 (500)	17.26 (489)	0.26	0.28

[*] Values in parenthesis indicates the number of trials out of 500 that had a regimen selection
[a] T - Target/True regimen, R - Recommended regimen

Table 2.18: Regimen recommendation summary for 500 simulated clinical trials executed using Plan 2, matching the regimen with $R_l \leq 2$.

	Mean(T-R)[a] Distance[*]	Prop of T-R=0[*]	Target Exp Dose[*]	Rec Exp Dose[*]	Target Toxicities	Rec Toxicities
Regimen selected using $\boldsymbol{P^r(A_1)}$&$\boldsymbol{P^r(C)}$						
$\boldsymbol{\mathring{P}(B) = 0.3, \mathring{P}(C) = 0.4}$						
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	2.87 (452)	0.5 (452)	15.89 (500)	14.95 (452)	0.26	0.22
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	2.83 (475)	0.25 (475)	15.89 (500)	16.26 (475)	0.26	0.26
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	3.87 (489)	0.22 (489)	17.43 (500)	17.26 (489)	0.26	0.28
$\boldsymbol{\check{P}(B) = 0.4, \check{P}(C) = 0.4}$						
$\check{P}(A_1) = 0.05, \check{P}(A_2) = 0.09$	2.93 (452)	0.49 (452)	15.89 (500)	15.03 (452)	0.26	0.22
$\check{P}(A_1) = 0.1, \check{P}(A_2) = 0.08$	2.81 (475)	0.25 (475)	15.89 (500)	16.25 (475)	0.26	0.26
$\check{P}(A_1) = 0.2, \check{P}(A_2) = 0.06$	3.89 (489)	0.22 (489)	17.43 (500)	17.24 (489)	0.26	0.28
Regimen selected using $\boldsymbol{P^r(C)}$						
$\boldsymbol{\mathring{P}(B) = 0.3, \mathring{P}(C) = 0.4}$						
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$	4.65 (489)	0.17 (489)	17.43 (500)	17.13 (489)	0.26	0.28
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.08$	4.24 (489)	0.20 (489)	17.43 (500)	16.82 (489)	0.26	0.28
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.06$	3.80 (489)	0.22 (489)	17.43 (500)	17.28 (489)	0.26	0.28
$\boldsymbol{\check{P}(B) = 0.4, \check{P}(C) = 0.4}$						
$\check{P}(A_1) = 0.05, \check{P}(A_2) = 0.09$	4.62 (489)	0.18 (489)	17.43 (500)	17.13 (489)	0.26	0.28
$\check{P}(A_1) = 0.1, \check{P}(A_2) = 0.08$	4.25 (489)	0.21 (489)	17.43 (500)	16.80 (489)	0.26	0.28
$\check{P}(A_1) = 0.2, \check{P}(A_2) = 0.06$	3.81 (489)	0.22 (489)	17.43 (500)	17.26 (489)	0.26	0.28

[*] Values in parenthesis indicates the number of trials out of 500 that had a regimen selection
[a] T - Target/True regimen, R - Recommended regimen

Table 2.19: Regimen recommendation summary for 500 simulated clinical trials executed using Plan 2, matching the regimen with regimen distance equal to $\min(R_l)$ and switching to Plan 1 not allowed.

	Mean(T-R)[a] Distance*	Prop of T-R=0*	Target Exp Dose*	Rec Exp Dose*	Target Toxicities	Rec Toxicities
Regimen selected using $\boldsymbol{P^r(A_1)}$&$\boldsymbol{P^r(C)}$						
$\boldsymbol{\mathring{P}(B)=0.3, \mathring{P}(C)=0.4}$						
$\mathring{P}(A_1)=0.05, \mathring{P}(A_2)=0.09$	2.71 (452)	0.51 (452)	15.89 (500)	14.87 (452)	0.26	0.22
$\mathring{P}(A_1)=0.1, \mathring{P}(A_2)=0.08$	2.89 (475)	0.27 (475)	15.89 (500)	16.13 (475)	0.26	0.26
$\mathring{P}(A_1)=0.2, \mathring{P}(A_2)=0.06$	3.89 (489)	0.21 (489)	17.43 (500)	17.17 (489)	0.26	0.28
$\boldsymbol{\mathring{P}(B)=0.4, \mathring{P}(C)=0.4}$						
$\mathring{P}(A_1)=0.05, \mathring{P}(A_2)=0.09$	2.64 (450)	0.52 (450)	15.89 (500)	14.91 (450)	0.26	0.22
$\mathring{P}(A_1)=0.1, \mathring{P}(A_2)=0.08$	2.82 (475)	0.27 (475)	15.89 (500)	16.16 (475)	0.26	0.26
$\mathring{P}(A_1)=0.2, \mathring{P}(A_2)=0.06$	3.77 (489)	0.22 (489)	17.43 (500)	17.21 (489)	0.26	0.28
Using $\boldsymbol{P^r(C)=0.30}$						
$\boldsymbol{\check{P}(B)=0.3, \check{P}(C)=0.4}$						
$\check{P}(A_1)=0.05, \check{P}(A_2)=0.09$	4.68 (489)	0.17 (489)	17.43 (500)	16.85 (489)	0.26	0.28
$\check{P}(A_1)=0.1, \check{P}(A_2)=0.08$	4.26 (489)	0.2 (489)	17.43 (500)	16.69 (489)	0.26	0.28
$\check{P}(A_1)=0.2, \check{P}(A_2)=0.06$	3.83 (489)	0.21 (489)	17.43 (500)	17.19 (489)	0.26	0.28
$\boldsymbol{\check{P}(B)=0.4, \check{P}(C)=0.4}$						
$\check{P}(A_1)=0.05, \check{P}(A_2)=0.09$	4.7 (489)	0.18 (489)	17.43 (500)	16.81 (489)	0.26	0.28
$\check{P}(A_1)=0.1, \check{P}(A_2)=0.08$	4.17 (489)	0.21 (489)	17.43 (500)	16.74 (489)	0.26	0.28
$\check{P}(A_1)=0.2, \check{P}(A_2)=0.06$	3.72 (489)	0.22 (489)	17.43 (500)	17.23 (489)	0.26	0.28

* Values in parenthesis indicates the number of trials out of 500 that had a regimen selection

[a] T - Target/True regimen, R - Recommended regimen

Table 2.20: Trial/Patient summary -using only $\mathring{P}(C)$ flag in carrying out 100 simulated trials.

	Mean Dose	Mean Toxicities*	Patients ≥R.Edose[a]	Patients ≥T.Edose[a]	Patients =R.Dist[b]	Patients =T.Dist[b]	P.R.Dist ≤2	P.T.Dist ≤ 2
Plan 1: $\boldsymbol{P^r(A_1)\&P^r(C)=0.30}$								
$\boldsymbol{\mathring{P}(C)=0.30}$ $\mathring{P}(A_1)=0.20, \mathring{P}(A_2)=0.03$	15.91	0.24 (2)	13.55	15.48	3.26	1.25	11.08	6.28
$\boldsymbol{\mathring{P}(C)=0.40}\mathring{P}(A_1)=0.20, \mathring{P}(A_2)=0.06$	16.40	0.27 (2)	16.67	17.68	4.87	1.97	11.69	8.68
Using $\boldsymbol{P^r(C)=0.30}$ in regimen selection								
$\boldsymbol{\mathring{P}(C)=0.30}$ $\mathring{P}(A_1)=0.20, \mathring{P}(A_2)=0.03$	15.91	0.24 (2)	13.5	15.48	2.72	1.25	10.57	6.28
$\boldsymbol{\mathring{P}(C)=0.40}\mathring{P}(A_1)=0.20, \mathring{P}(A_2)=0.06$	16.40	0.27 (2)	16.64	17.68	4.19	1.97	11.04	8.68
Plan 2:$\boldsymbol{P^r(A_1)\&P^r(C)=0.30}$								
$\boldsymbol{\mathring{P}(C)=0.30}\mathring{P}(A_1)=0.20, \mathring{P}(A_2)=0.03$	16.45	0.27 (2)	18.62	18.39	5.09	8.20	10.04	12.67
$\boldsymbol{\mathring{P}(C)=0.40}$ $\mathring{P}(A_1)=0.20, \mathring{P}(A_2)=0.06$	17.20	0.30 (2)	19.19	19.25	5.37	3.63	12.02	10.91
Using $\boldsymbol{P^r(C)=0.30}$ in regimen selection								
$\boldsymbol{\mathring{P}(C)=0.30}\mathring{P}(A_1)=0.20, \mathring{P}(A_2)=0.03$	16.45	0.27 (2)	18.59	18.39	4.68	8.20	9.60	12.67
$\boldsymbol{\mathring{P}(C)=0.40}$ $\mathring{P}(A_1)=0.20, \mathring{P}(A_2)=0.06$	17.20	0.30 (2)	18.93	19.25	5.10	3.63	11.66	10.91

* Values in parenthesis indicates the number of trials out of 100 that stopped early

[a] T/R.Edose - Target/Recommended expected dose.

[b] T/R.Dist - Distance of patient regimen from Target/Recommended regimen.

Table 2.21: Regimen recommendation summary using $\mathring{P}(C)$ instead of $\mathring{P}(B)$ on later cycles for Plan 1 and 2, with Plan 2 executed by setting R_l to the minimum distance possible in 100 simulated trials.

	Mean(T-R)[a] Distance[*]	Prop of T-R=0[*]	Target Exp Dose[*]	Rec Exp Dose[*]	Target Toxicities	Rec Toxicities
Plan 1: Regimen selected using $\boldsymbol{P^r(A_1)}$&$\boldsymbol{P^r(C)}$						
$\boldsymbol{\check{P}(C) = 0.30}$, $\check{P}(A_1) = 0.20, \check{P}(A_2) = 0.03$	4.26 (98)	0.14 (98)	17.43 (100)	18.06 (98)	0.26	0.27
$\boldsymbol{\check{P}(C) = 0.40}$,$\check{P}(A_1) = 0.20, \check{P}(A_2) = 0.06$	4.01 (98)	0.15 (98)	17.43 (100)	17.62 (98)	0.26	0.28
Regimen selected using $\boldsymbol{\tilde{P}(C)}$						
$\boldsymbol{\check{P}(C) = 0.30}$,$\check{P}(A_1) = 0.20, \check{P}(A_2) = 0.03$	4.13 (98)	0.17 (98)	17.43 (100)	18.12 (98)	0.26	0.28
$\boldsymbol{\check{P}(C) = 0.40}$,$\check{P}(A_1) = 0.20, \check{P}(A_2) = 0.06$	3.89 (98)	0.18 (98)	17.43 (100)	17.67 (98)	0.26	0.28
Plan 2: Regimen selected using $\boldsymbol{P^r(A_1)}$&$\boldsymbol{P^r(C)}$						
$\boldsymbol{\check{P}(C) = 0.30}$,$\check{P}(A_1) = 0.20, \check{P}(A_2) = 0.03$	4.21 (98)	0.20 (98)	17.43 (100)	17.62 (98)	0.26	0.27
$\boldsymbol{\check{P}(C) = 0.40}$,$\check{P}(A_1) = 0.20, \check{P}(A_2) = 0.06$	3.89 (98)	0.21 (98)	17.43 (100)	17.77 (98)	0.26	0.28
Regimen selected using $\boldsymbol{P^r(C)}$						
$\boldsymbol{\check{P}(C) = 0.30}$, $\check{P}(A_1) = 0.20, \check{P}(A_2) = 0.03$	4.09 (98)	0.23 (98)	17.43 (100)	17.67 (98)	0.26	0.28
$\boldsymbol{\check{P}(C) = 0.40}$, $\check{P}(A_1) = 0.20, \check{P}(A_2) = 0.06$	3.64 (98)	0.26 (98)	17.43 (100)	17.87 (98)	0.26	0.28

[*] Values in parenthesis indicates the number of trials out of 100 that had a regimen selection
[a] T - Target/True regimen, R - Recommended regimen

2.6 Appendix

2.6.1 Outline of the code written in JAGS

The code presented below corresponds to applying Model 2.1 in simulations for parameter estimation from the posterior samples.

```
#Defining the model.bug file
model {
#Define the likelihood for each of the N subjects
for (i in 1:N) {
prob[i] <- 1- exp( -alpha* ( dose[i] - rho*maxprevdose[i])*step(dose[i] -rho*maxprevdose[i])
- beta*dose[i]*cumdose[i] )
response[i] ~ dbern(prob[i])
}
#Setting up the priors
#prior on α - E(α) = 1 and Var(α) = 2
mu1 <- -0.8047190; tau1 <- 0.6213349
alpha ~ dlnorm(mu1,tau1)
#prior on ρ
a1 <- 5 ; b1 <- 1
rho ~ dbeta(a1,b1)
#Prior on β - E(β) = 0.5 and Var(β) = 1
mu2 <- -1.498; tau2 <- 0.621
beta ~ dlnorm(mu2,tau2)
}
#Initializing the parameters
inits< −list(list(alpha=1,beta=0.1,rho=0.2),
list(alpha=0.5,beta=0.8,rho=0.9) )
parameters < −c("alpha","beta", "rho")
#updating the simulations
data < −list("response"=response,"maxprevdose"=maxprevdose,
"cumdose"=cumdose,"dose"=dose,"N"=N)
jags <- jags.model(file="prior.bug", data = data, inits=inits, n.chains = 2, n.adapt
```

```
= 5000)
adapt(jags,n.iter=1000)
update(jags,10000) # burin samples
sim1< −coda.samples(jags, parameters, 100000,thin=20)
#check for convergence
plot(sim1)
gelman.plot(sim1)
gelman.diag(sim1)
geweke.plot(sim1)
geweke.diag(sim1)
autocorr(sim1)
autocorr.plot(sim1)
#report the mean and quantiles of the posterior distributions
y3< −summary(sim1)
ystat< −data.frame(y3statistics)
yquant = data.frame(y3quantiles)
```

2.6.2 Parameter Estimation with different burn-in period

In this section simulation results are presented explaining the rationale behind the decisions for setting the variance of the priors and the choice of burn-in period.

The effect of changing the variance of the prior distributions of α and β was studied with $N = (10, 30)$ patients in completed trials. The probability skeleton $q_g = (0.02, 0.05, 0.10, 0.15, 0.23)$ was used to obtain doses d_g. Patients were assumed to receive the same dose on all the $K = 6$ cycles and the N patients were divided equally among the five dose levels. The conditional probability of toxic response $p_{i,k}$ for patient i on cycle k was calculated using Model 2.1 for dose d_g and fixed parameters $\alpha = 1$, $\beta = 0.5$ and $\rho = 0.8$. A Bernoulli $(p_{i,k})$ random variable was used to assign the response for each patient at every cycle. Further doses were assigned

until a DLT response was observed or until completion of $K = 6$ cycles. Simulation results are presented over 500 such replicates/datasets.

The use of probability skeleton requires that the prior mean of α be set at one. The prior mean of β was set to 0.5, matching the true value used for data generation in simulations. To study the effect of the priors at the estimation stage, two different cases were considered. In the first case the prior means used were $E(\alpha) = 1$ and $E(\beta) = 0.5$ matching the true values used to generate the data. The SD was set to two times the mean, $SD(\alpha) = 2$ and $SD(\beta) = 1$.

In the second case the means of the prior were $E(\alpha) = 1$ and $E(\beta) = 0.5$ matching the true values used to generate the data while the standard deviation was set to five times the mean, $SD(\alpha) = 5$ and $SD(\beta) = 2.5$

In both the cases $Beta(5, 1)$ prior was used on ρ. The goal was to observe the effect of varying the informativeness of the priors on the estimation of the parameters. After an adaptive phase of 10K samples and a burn-in period of 10K samples, 100K MCMC samples were drawn from the posterior with a thinning of 20. The effect of using a longer burn-in period was also studied by drawing an additional 100K MCMC samples effectively increasing the burn-in period to 110K (10K from the initial burn-in and 100K from the first sample). The goal of studying estimates from two different burn-in periods was to study and determine the appropriate burn-in period for use in further simulations.

The simulation results of the parameter estimates from these simulations are

presented in Table 2.22 for $N = (10, 30, 90)$ patients. Within each of the sample size, the rows are grouped by the two cases followed by the two burn-in periods with each row corresponding to estimates of either α, β or ρ . Each of the four columns report (1) the estimated value and the mean bias from the true value, (2) the mean SD (MSD) of the estimates across 500 iterations, (3) the empirical SD (ESD) of the 500 estimates, (4) the credible interval coverage rate across the 500 replicates/datasets.

The corresponding probability estimates are presented in Table 2.23 for $N = 10$ patients and Table 2.24 for $N = 30$ patients. The columns indicate the mean estimate of probability of toxicity on the first, the second, the sixth cycle and the overall probability of toxicity on any of the cycles along with the bias from the true values in parenthesis. The rows are grouped by the two cases studying the effect of varying variances of the prior and within each case the results are further grouped by the two different burn-in periods with each of the rows corresponding to one of the dose levels. The following are the conclusions from these simulations results.

- Using a longer burn-in period does not provide improve the parameter estimates in terms of the bias. Hence using the shorter burn-in period of 10K is recommended for further simulation setups.

- The bias difference between the two cases not very different and showed that having the SD equal to two times the mean with a CV of 2 provided a fair degree of variability. There was not much gain from using a larger variance.

- The bias of the parameter estimates seems quite high in Table 2.22 but when compared to the conditional probability estimates in Tables 2.23 and 2.24 the bias was not high. So it seems like the model estimates the values of the conditional probabilities of toxicity correctly even though the parameters might have a slight bias from the true values. Subsequent tables will display either the probabilities or the parameter estimates to prevent replication of information.

- These results demonstrate the best the model can achieve in determining the parameters and the probabilities since the prior means are completely aligned with the true values used to generate the data. This provides a sense of how well the model does in terms of estimation and will provide a benchmark in assessing the model fit in later simulation settings.

- An improvement in bias and efficiency was seen in terms of lower SD of the estimates with an increase in the sample size. Also the empirical SD of the estimates and the mean SD from the 500 datasets was comparable as the sample size increased. This confirms that the having a larger sample size improves estimates and that the distribution of the estimates obtained is alike to the sampling distribution.

Table 2.22: Parameter Estimates Based on Model 2.1 from 500 simulated datasets containing $N = (10, 30)$ patients receiving one of five dose groups $d_1 \dots d_5$ over six cycles comparing different burin periods. Results presented for $\alpha = 1, \beta = 0.5$ and $\rho = 0.8$

	Estimate(Bias)	MSD [1]	ESD [2]	Coverage
N=10				
Burn in 10K	**Case 1**			
α	0.945 (-0.055)	0.683	0.584	99.6
β	0.515 (0.015)	0.463	0.308	100
ρ	0.820 (0.02)	0.137	0.040	100
Burn in 110K				
α	0.945 (-0.055)	0.683	0.584	99.6
β	0.515 (0.015)	0.463	0.308	100
ρ	0.820 (0.02)	0.137	0.040	100
Burn in 10K	**Case 2**			
α	0.932 (-0.068)	0.725	0.697	95.2
β	0.481 (-0.019)	0.498	0.387	94.6
ρ	0.813 (0.013)	0.140	0.039	100
Burn in 110K				
α	0.932 (-0.068)	0.724	0.697	95.2
β	0.481 (-0.019)	0.498	0.386	94.8
ρ	0.813 (0.013)	0.140	0.039	100
N=30				
Burn in 10K	**Case 1**			
α	0.957 (-0.043)	0.458	0.417	96.6
β	0.511 (0.011)	0.338	0.296	97.8
ρ	0.806 (0.006)	0.132	0.056	100
Burn in 110K				
α	0.957 (-0.043)	0.457	0.416	96.6
β	0.511 (0.011)	0.339	0.296	98.0
ρ	0.806 (0.006)	0.132	0.055	100
Burn in 10K	**Case 2**			
α	0.969 (-0.031)	0.477	0.455	94.6
β	0.464 (-0.036)	0.353	0.338	95.0
ρ	0.793 (-0.007)	0.136	0.057	100
Burn in 110K				
α	0.968 (-0.032)	0.476	0.455	94.4
β	0.464 (-0.036)	0.353	0.338	95.0
ρ	0.793 (-0.007)	0.136	0.057	100

[1] MSD is mean of the SD from 500 estimates

[2] ESD is empirical SD of the 500 estimates

Table 2.23: Table comparing the effect of using two different burn-in periods and the effect of using two different priors (Case 1 and Case 2) on the parameters. Estimates (bias) from the true value of conditional probability of toxicity on each of the six cycles estimated using Model 2.1 from 500 simulated datasets containing $N = 10$ patients each receiving one of five dose groups $d_1 \dots d_5$.Results presented for $\alpha = 1, \beta = 0.5$ and $\rho = 0.8$

	Cycle 1	Cycle 2	Cycle 6	Any Cycle
	Est(bias)	Est(bias)	Est(bias)	Est(bias)
Case 1 Burn in 10K				
d_1	0.019 (-0.001)	0.004 (< 0.001)	0.005 (< 0.001)	0.04 (-0.002)
d_2	0.047 (-0.003)	0.011 (-0.001)	0.016 (< 0.001)	0.109 (-0.006)
d_3	0.093 (-0.007)	0.025 (-0.001)	0.047 (-0.001)	0.242 (-0.013)
d_4	0.148 (-0.012)	0.047 (-0.002)	0.104 (-0.001)	0.415 (-0.023)
d_5	0.210 (-0.020)	0.081 (-0.002)	0.197 (-0.003)	0.609 (-0.036)
Burn in 110K				
d_1	0.019 (-0.001)	0.004 (< 0.001)	0.005 (< 0.001)	0.04 (-0.002)
d_2	0.047 (-0.003)	0.011 (-0.001)	0.016 (< 0.001)	0.109 (-0.006)
d_3	0.093 (-0.007)	0.025 (-0.001)	0.047 (< 0.001)	0.242 (-0.013)
d_4	0.148 (-0.012)	0.047 (-0.002)	0.104 (-0.001)	0.416 (-0.023)
d_5	0.210 (-0.020)	0.081 (-0.002)	0.197 (-0.003)	0.609 (-0.036)
Case 2 Burn in 10K				
d_1	0.019 (-0.001)	0.004 (< 0.001)	0.005 (< 0.001)	0.039 (-0.003)
d_2	0.046 (-0.004)	0.011 (-0.001)	0.016 (< 0.001)	0.106 (-0.009)
d_3	0.091 (-0.009)	0.025 (-0.002)	0.045 (-0.003)	0.233 (-0.021)
d_4	0.144 (-0.016)	0.046 (-0.003)	0.099 (-0.006)	0.399 (-0.040)
d_5	0.204 (-0.026)	0.078 (-0.005)	0.185 (-0.015)	0.582 (-0.063)
Burn in 110K				
d_1	0.019 (-0.001)	0.004 (< 0.001)	0.005 (< 0.001)	0.039 (-0.003)
d_2	0.046 (-0.004)	0.011 (-0.001)	0.016 (< 0.001)	0.106 (-0.009)
d_3	0.091 (-0.009)	0.025 (-0.002)	0.045 (-0.003)	0.233 (-0.022)
d_4	0.144 (-0.016)	0.046 (-0.003)	0.099 (-0.006)	0.398 (-0.040)
d_5	0.204 (-0.026)	0.078 (-0.005)	0.185 (-0.015)	0.582 (-0.063)

Table 2.24: Table comparing the effect of using two different burn-in periods and the effect of using four different priors on the parameters. Parameter Estimates with bias from the true value of conditional probability of no toxicity on each of the six cycles estimated using Model 2.1 from 500 simulated datasets containing $N = 30$ patients each receiving one of five dose groups $d_1 \dots d_5$.Results presented for $\alpha = 1, \beta = 0.5$ and $\rho = 0.8$

	Cycle 1	Cycle 2	Cycle 6	Any Cycle
	Est(bias)	Est(bias)	Est(bias)	Est(bias)
Case 1 Burn in 10K				
d_1	0.019 (-0.001)	0.004(< 0.001)	0.005(< 0.001)	0.041 (-0.002)
d_2	0.048 (-0.002)	0.011 (< 0.001)	0.016 (< 0.001)	0.111 (-0.004)
d_3	0.095 (-0.005)	0.025 (-0.001)	0.047 (< 0.001)	0.246 (-0.008)
d_4	0.152 (-0.008)	0.048 (-0.001)	0.104 (-0.001)	0.425 (-0.014)
d_5	0.217 (-0.013)	0.081 (-0.002)	0.197 (-0.002)	0.623 (-0.022)
Burn in 110K				
d_1	0.019 (-0.001)	0.004 (< 0.001)	0.005 (< 0.001)	0.041 (-0.002)
d_2	0.048 (-0.002)	0.011 (< 0.001)	0.016 (< 0.001)	0.111 (-0.004)
d_3	0.095 (-0.005)	0.025 (-0.001)	0.047(< 0.001)	0.246 (-0.008)
d_4	0.151 (-0.009)	0.048 (-0.001)	0.104 (-0.001)	0.425 (-0.014)
d_5	0.217 (-0.013)	0.081 (-0.002)	0.198 (-0.002)	0.623 (-0.022)
Case 2 Burn in 10K				
d_1	0.019 (-0.001)	0.004 (< 0.001)	0.005 (< 0.001)	0.041 (-0.002)
d_2	0.048 (-0.002)	0.011 (< 0.001)	0.016 (< 0.001)	0.110 (-0.005)
d_3	0.096 (-0.004)	0.025 (-0.001)	0.045 (-0.003)	0.242 (-0.013)
d_4	0.153 (-0.007)	0.047 (-0.002)	0.098 (-0.007)	0.413 (-0.025)
d_5	0.218 (-0.012)	0.079 (-0.004)	0.184 (-0.016)	0.604 (-0.040)
Burn in 110K				
d_1	0.019 (-0.001)	0.004 (< 0.001)	0.005 (< 0.001)	0.041 (-0.002)
d_2	0.048 (-0.002)	0.011 (< 0.001)	0.016 (< 0.001)	0.11 (-0.005)
d_3	0.096 (-0.004)	0.025 (-0.001)	0.045 (-0.003)	0.242 (-0.013)
d_4	0.153 (-0.007)	0.047 (-0.002)	0.098 (-0.007)	0.414 (-0.025)
d_5	0.218 (-0.012)	0.079 (-0.004)	0.184 (-0.016)	0.604 (-0.040)

2.6.3 Patient profiles for mixed dose assignment

The following table shows the potential dose course for the 30 patients to be used in the simulations to compare the efficiency gain in using dose variation within patients.

Table 2.25: Table showing the dose level assignment at each cycle for the $N = 30$ patients so that each of the dose occurs 36 times over all the patients.

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
1	1	1	1	1	1	1
2	1	1	1	1	1	1
3	1	1	1	1	1	1
4	2	2	2	2	2	2
5	3	3	3	3	3	3
6	4	4	4	4	4	4
7	5	5	5	5	5	5
8	5	5	5	5	5	5
9	5	5	5	5	5	5
10	1	1	1	2	2	2
11	1	1	1	2	2	2
12	2	2	2	3	3	3
13	2	2	2	3	3	3
14	3	3	3	4	4	4
15	3	3	3	4	4	4
16	4	4	4	5	5	5
17	4	4	4	5	5	5
18	2	2	2	1	1	1
19	3	3	3	2	2	2
20	4	4	4	3	3	3
21	5	5	5	4	4	4
22	1	1	2	2	3	3
23	1	1	2	2	3	3
24	2	2	3	3	4	4
25	2	2	3	3	4	4
26	3	3	4	4	5	5
27	3	3	4	4	5	5
28	5	5	4	4	3	3
29	4	4	3	3	2	2
30	3	3	2	2	1	1

2.6.4 Example of an adaptive trial in progress

Based on the algorithm presented in Section 2.4.4.1 for Plan 1, an example of a trial in progress is presented in this section for demonstrating the dose assignment

in practice. The target probability bounds used during the execution are $\mathring{P}(C) = 0.40, \mathring{P}(B) = 0.30, \mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.09$ and the true values of the parameters are $\alpha = 1, \beta = 0.2, \rho = 0.8$.

Table 2.26 presents the current patient profile in the trial. The rows correspond to the unique patients added sequentially in the trial. The columns correspond to the six cycles with the dose level assigned to the patient and the response in parenthesis. A zero signifies no DLT while a 1 denotes a DLT, a cross is placed in all cycles once a DLT response is observed for a patient. There are 12 patients in the trial and decisions need to be made for dose assignment to patients 8, 9, 10, and 11 and a new patient 13. Before patient 12 was added to the trial the parameter estimates were $\hat{\alpha} = 0.421(0.404), \hat{\beta} = 1.142(0.8377)$ and $\hat{\rho} = 0.826(0.147)$ with the posterior standard deviations in parenthesis. The updated current estimates of the parameters are $\hat{\alpha} = 0.7635(0.566)$ $\hat{\beta} = 0.917(0.711)$ and $\hat{\rho} = 0.829(0.147)$. Notice that the estimate of $\hat{\alpha}$ increases in response to the DLT observed by patient 12 on cycle 1.

The probability of toxicity on dose levels 1 through 4 for patient 8 are 0.010,0.025, 0.051,0.103 of which dose level 3 has probability of toxicity $\leq \mathring{P}(A_2) = 0.09$ and is also able to provide a regimen combination that satisfies the $\mathring{P}(C) = 0.40, \mathring{P}(B) = 0.30$ and hence is assigned to patient 8 on cycle 6. Using the true values of the parameters and the current dose assignment the true probability of toxicity is calculated and a Bernoulli response is generated. In a similar fashion the remaining patients are

assigned doses and responses and the updated patient profile is presented in Table 2.27. The updated parameter estimates are now $\hat{\alpha} = 0.684(0.519), \hat{\beta} = 1.323(0.829)$ and $\hat{\rho} = 0.838(0.146)$. Notice the decrease in estimate of $\hat{\alpha}$ since no fresh toxicities on the first cycle but there is an increase in $\hat{\beta}$ reflecting the toxicity on cycle 6 for patient 8.

At the end of the trial the completed patient profile is presented in Table 2.28. The parameter estimates at the conclusion of the trial are $\hat{\alpha} = 0.943(0.482)$, $\hat{\beta} = 0.738(0.4234)$ and $\hat{\rho} = 0.866(0.129)$. The probability of toxicity on the first cycle and on any cycle is calculated using the current estimates of the parameters for all the 19 regimens in Table 2.3 to select the recommended regimen and by using the true parameter values to select the target regimen. By setting $P^r(A_1) = \mathring{P}(A_1) = 0.05$ and $P^r(C) = 0.3$ and using the true parameter values the target regimen selected is 223344 while the recommended regimen is 222333 using the parameter estimates obtained at the conclusion of the trial. If only $P^r(C) = 0.3$ is used the target regimen is 444333 while the recommended regimen is 333333.

Table 2.26: Table showing the dose level assignment and patient responses in parenthesis for an adaptive trial in progress with accrual of 12 patients.

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
1	2 (0)	2 (0)	3 (0)	4 (0)	4 (1)	X
2	2 (0)	3 (0)	4 (0)	4 (0)	4 (0)	3 (0)
3	2 (0)	3 (0)	4 (0)	4 (0)	4 (0)	3 (0)
4	2 (0)	3 (0)	4 (0)	4 (0)	4 (0)	3 (0)
5	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
6	2 (0)	3 (0)	4 (1)	X	X	X
7	3 (0)	3 (0)	2 (0)	3 (0)	3 (0)	3 (0)
8	3 (0)	4 (0)	2 (0)	3 (0)	3 (0)	?
9	2 (0)	3 (0)	3 (0)	4 (0)	?	?
10	3 (0)	3 (0)	3 (0)	?	?	?
11	3 (0)	3 (0)	?	?	?	?
12	3 (1)	X	X	X	X	X

X- Terminated patient having a severe toxicity (1)

? Continuing patient

Table 2.27: Table showing the dose level assignment and patient responses in parenthesis for an adaptive trial in progress with accrual of 13 patients.

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
1	2 (0)	2 (0)	3 (0)	4 (0)	4 (1)	X
2	2 (0)	3 (0)	4 (0)	4 (0)	4 (0)	3 (0)
3	2 (0)	3 (0)	4 (0)	4 (0)	4 (0)	3 (0)
4	2 (0)	3 (0)	4 (0)	4 (0)	4 (0)	3 (0)
5	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
6	2 (0)	3 (0)	4 (1)	X	X	X
7	3 (0)	3 (0)	2 (0)	3 (0)	3 (0)	3 (0)
8	3 (0)	4 (0)	2 (0)	3 (0)	3 (0)	3(1)
9	2 (0)	3 (0)	3 (0)	4 (0)	4 (0)	?
10	3 (0)	3 (0)	3 (0)	3 (0)	?	?
11	3 (0)	3 (0)	3 (0)	?	?	?
12	3 (1)	X	X	X	X	X
13	2 (0)	?	?	?	?	?

X- Terminated patient having a severe toxicity (1)

? Continuing patient

Table 2.28: Table showing the dose level assignment and patient responses in parenthesis for a completed adaptive trial.

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
1	2 (0)	2 (0)	3 (0)	4 (0)	4 (1)	X
2	2 (0)	3 (0)	4 (0)	4 (0)	4 (0)	3 (0)
3	2 (0)	3 (0)	4 (0)	4 (0)	4 (0)	3 (0)
4	2 (0)	3 (0)	4 (0)	4 (0)	4 (0)	3 (0)
5	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
6	2 (0)	3 (0)	4 (1)	X	X	X
7	3 (0)	3 (0)	2 (0)	3 (0)	3 (0)	3 (0)
8	3 (0)	4 (0)	2 (0)	3 (0)	3 (0)	3(1)
9	2 (0)	3 (0)	3 (0)	4 (0)	4 (0)	3(0)
10	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
11	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
12	3 (1)	X	X	X	X	X
13	2 (0)	3 (1)	X	X	X	X
14	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
15	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
16	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
17	2 (0)	3 (0)	3 (0)	3 (0)	3 (1)	X
18	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
19	2 (0)	3 (0)	4 (0)	3 (0)	3 (0)	3 (0)
20	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
21	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
22	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
23	2 (0)	3 (0)	4 (0)	4 (0)	3 (0)	3 (0)
24	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
25	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
26	2 (0)	3 (1)	X	X	X	X
27	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
28	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
29	2 (0)	3 (0)	3 (0)	3 (0)	3 (0)	3 (0)
30	2 (1)	X	X	X	X	X

X- Terminated patient having a severe toxicity (1)

CHAPTER 3

Multivariate Markov Models for the Conditional Probability of Toxicity in Phase II Trials

3.1 Background and significance

The study of dose toxicity relationships in oncology occurs in phase I and phase II clinical trials and variations of these. We focus on the setting where patients are randomized to two, or more, dose groups administered over several cycles. An example of this study design can be found in [Worden et al., 2005, Chugh et al., 2007], which investigated two randomized doses of ifosamide plus doxorubicin and granulocyte colony-stimulating factor, hereafter called ifosamide for brevity. Depending on whether they had metastases, patients received either six or four cycles of either $6g/m^2$ or $12g/m^2$ of ifosamide given over a 4 day period at the beginning of a 21 day cycle. As is often the case, patients who experienced treatment toxicity were not continued in subsequent cycles. Adverse events are classified by Common Toxicity

Criteria as defined by National Cancer Institute [NCI, 2003] and grade 3 or higher toxic responses are considered to be a dose limiting toxicity (DLT). Proportions of continuing patients in each cycle who had a hemoglobin DLT are show in Figure 3.1. Two important features of the data are (1) the conditional nature of toxicity proportions observed at each cycle that are based on previously toxicity-free patients and (2) the trend towards higher conditional toxicity rates as dose accumulates over cycles in the high dose group. Possible analytical tools for this data are generalized estimating equations (GEE) or generalized linear mixed models (GLMM) that account for correlation within a patient treated over multiple cycles. These models allow covariates for cycle, dose, cumulative dose and other mitigating factors in modeling the probability of toxicity. An example of analyzing such data can also be found in [Legedza and Ibrahim, 2000] that applies to phase I clinical trials and as a special case of [Doussau et al., 2013] that uses ordinal outcomes also in a phase I setting. All these methods do not explicitly model the tendency to discontinue cycles for patients who have demonstrated previous DLT, although the resulting estimated toxicity rates may be conditional in nature. In fact, conditional probabilities of toxicity are precisely what are needed in order to advise patients during subsequent therapy. Developing models that explicitly capture these conditional probabilities are key.

A large number of parameters may be necessary to capture all features of the dose-toxicity relationship, and with the typically small sample size available for modeling

purposes, a Bayesian approach is an attractive alternative. Our model is an extension to transitional models with first-order Markov chains [Agresti, 2002] considering only the previous cycle. In this chapter, we use a Markov model to explicitly model conditional probabilities of toxicity in a cycle given that the patient did not have a DLT in the past. The proposed model allows for a cumulative effect of dose on toxicity after the first cycle and allows covariates to influence the toxicity profile over cycles. A parameter is included to reflect an individual's tendency to respond consistently with past dose experience. In Section 3.2, we formulate a model useful in this phase II clinical trial setting, describe the Bayesian estimation method and provide intuition on model behavior. We then study finite sample operating characteristics through simulation in Section 3.3. In Section 3.4 we apply the methods to the ifosfamide study described earlier and follow with a discussion in Section 3.5.

3.2 Methodology

In Section 3.2.1, we define the data structure and the proposed dose-toxicity model. Calculations for the expected total dose over K potential treatment cycles as well as the expected number of completed cycles is presented in Section 3.2.2. Technical details of dose modeling via skeleton probabilities of toxicity during the first cycle of treatment are covered in Section 3.2.3. Selection of priors and the formulation of the posterior distribution is presented in Section 3.2.4 followed by a section 3.2.5 reviewing model selection strategies.

3.2.1 Proposed Markov model

We assume patients $i = 1 \ldots N$ are randomized to one of two doses $S_g, g = 1, 2$. In Section 3.2.3, we will show that it is convenient to rescale doses S_g to d_g; a strategy useful for incorporating beliefs about toxicity on the first cycle, similar to what was proposed in [O'Quigley et al., 1990, Lee and Cheung, 2009]. For convenience, d_g terms will hereafter be referred to as the dose assigned per cycle to group g, even though numerically d_g must be transformed back to the S_g scale to reflect actual doses.

Each dose group is scheduled to undergo K cycles of treatment, however, individual patients complete K_i cycles, where K_i may be less than K if experiencing a DLT. On each cycle $k = 1, \ldots, K_i$, patient i randomized to group g receives dose $d_{i,k} = d_g$. Since dose is constant across cycles, we will typically use the notation $d_{i,1}$ for the dose given to patient i at each cycle. A patient's cumulative dose prior to cycle k is $D_{i,k} = (k-1) \times d_{i,1}$, with $D_{i,1} = 0$. We also use the convention that $d_{i,k-1} = 0$ for cycle $k = 1$.

A Bernoulli random variable, $Y_{i,k}$, denotes the occurrence of a DLT for patient i on cycle k. In general, $Y_{i,k} = 0$ for $k = 1, \ldots, K_i - 1$, indicating no DLT on these cycles. If patient i completes $K_i = K$ cycles, then $Y_{i,K}$ may be either zero or one, depending on manifestation of a DLT in the final cycle. Additional patient covariates ($\mathbf{Z}$) are available for modeling the dose-toxicity relationship. So the observed data becomes $(Y_{i,1}, \ldots, Y_{i,K_i}, Z_i, d_{i,1})$.

We define $p_{i,k} = P(Y_{i,k} = 1|Y_{i,k-1} = 0, \ldots, Y_{i,1} = 0, Z_i, d_{i,1})$ as the conditional probability of a DLT for patient i on cycle k given that patient i has experienced no previous DLTs. Model 3.1 for $p_{i,k}$ is

$$\log(1 - p_{i,k}) = -g_1(\boldsymbol{\alpha}, \boldsymbol{Z})(d_{i,1} - g_2(\boldsymbol{\rho}, \boldsymbol{Z})d_{i,k-1}) - g_3(\boldsymbol{\beta}, \boldsymbol{Z})D_{i,k}, \tag{3.1}$$

where $\boldsymbol{\alpha}, \boldsymbol{\rho}$, and $\boldsymbol{\beta}$ parameterize the relationship between dose, $\mathbf{Z}$ and $p_{i,k}$. In the simplest case, $g_1(\cdot), g_2(\cdot)$ and $g_3(\cdot)$ are identity functions not involving $\mathbf{Z}$, so that Model 3.1 reduces to Model 3.2 below:

$$\log(1 - p_{i,k}) = -\alpha(d_{i,1} - \rho d_{i,k-1}) - \beta D_{i,k} \tag{3.2}$$

or equivalently, $p_{i,k} = 1 - \exp[-\alpha(d_{i,1} - \rho d_{i,k-1}) - \beta D_{i,k}]$. (3.3)

Intuition behind the parameters is easiest to follow for the special case in Model 3.2, where α, ρ and β are 1-dimensional parameters. The parameter, α, accounts for DLT encountered on cycle 1. The parameter, $0 \leq \rho \leq 1$, allows for a reduced probability of toxicity related to dose $d_{i,1}$ in a subsequent cycle if patient i has previously tolerated this dose; this term captures dependency in short-term toxicity outcomes between cycles. The effect of cumulative dose from previous cycles is captured by β.

Figure 3.2 shows the dose-toxicity relationship across cycles and $p_{i,k}$ for two dose

groups for several parameter combinations with $K = 4$ cycles. For instance, in the top left panel where α is the only parameter driving the dose-toxicity relationship, a patient has an independent DLT response at each dose administration so that $p_{i,k}$ is constant across cycles. In the top right panel, the parameter ρ equals 1 indicating that a patient tolerating dose $d_{i,1}$ on cycle 1, will not experience future toxicity at this dose level. In the lower three panels, the influence of increasing β alters the dose-toxicity relationship based on the effect of cumulative dose.

3.2.2 Expected total dose and completed cycles

A higher dose might not be attractive if fewer cycles can be completed at that dose level based on DLTs. Investigators should gain a clear understanding of the expected number of completed cycles for a dose level as well as the expected total dose over the entire trial based on Model 3.1 or the special case without covariates, Model 3.2. An individual i's total expected dose is based on probabilities $P(Y_{i,1} = 0, \ldots, Y_{i,K_i-1} = 0, Y_{i,K_i} = 1)$

$$= \quad (1-p_{i,1})(1-p_{i,2})\cdots(1-p_{i,K_i-1})p_{i,K_i} = p_{i,K_i}\prod_{j=1}^{K_i-1}(1-p_{i,j}) \qquad (3.4)$$

$$\text{and similarly, } P(Y_{i,1} = 0, \ldots, Y_{i,K-1} = 0, Y_{i,K} = 0) \quad = \quad \prod_{j=1}^{K}(1-p_{i,j}). \qquad (3.5)$$

Hence for person i, the expected number of completed cycles:

$$= p_{i,1} + \sum_{k=2}^{K-1} k p_{i,k} \prod_{j=1}^{k-1}(1 - p_{i,j}) + K \prod_{j=1}^{K-1}(1 - p_{i,j}) \tag{3.6}$$

and the expected total dose:

$$= p_{i,1} d_{i,1} + \sum_{k=2}^{K-1} k d_{i,1} p_{i,k} \prod_{j=1}^{k-1}(1 - p_{i,j}) + K d_{i,1} \prod_{j=1}^{K-1}(1 - p_{i,j}). \tag{3.7}$$

3.2.3 Model calibration using skeleton probabilities

For ease of interpretation of the parameter α, it is convenient to rescale the doses. As explained below they are rescaled based on initial guesses (skeleton probabilities) of the toxicity rate for the first cycle. For simplicity, we first describe use of skeleton probabilities for the simple case with no covariates, as in Model 3.2, and later suggest modifications for the more complex settings. Our strategy of defining skeleton probabilities and corresponding (transformed) dose values consistent with Model 3.2 is similar to that described by [O'Quigley et al., 1990, Lee and Cheung, 2009] in the context of the continual reassessment method in phase I studies as well as other contexts [Lee et al., 2011, Cheung and Elkind, 2010].

In our phase II setting, patients are given $S_1 = 6g/m^2$ of ifosamide or $S_2 = 12g/m^2$ of ifosamide over $K = 4$ cycles of treatment. Suppose q_g is an initial guess (skeleton probability) of a DLT at cycle one for dose S_g. Instead of using dose S_g in

Model 3.2, we use dose values $d_g, g = 1, 2$, satisfying $ln(1 - q_g) = -d_g$. For example, if we choose $q_1 = 0.10$, the resulting d_1 that stands in for $6g/m^2$ of ifosamide in Model 3.2 is $d_1 = -ln(1 - 0.10) \approx 0.11$. In defining d_g, we've conveniently normalized α to 1.0 if the skeleton probability for toxicity at cycle 1 is correct, making it easier to see deviations from the skeleton in the posterior distribution of α. Skeleton probabilities may be elicited from clinicians, previous animal studies or earlier phase clinical trials.

In Model 3.1, $g_1(\boldsymbol{\alpha}, \boldsymbol{Z})$, will be a known function of covariates, $\boldsymbol{Z}$, and skeleton probabilities for toxicity will need to be defined for reference values of $\boldsymbol{Z}$ in this relationship. For example, in Section 3.4, we consider the effect of gender (M, F) on $\boldsymbol{\alpha}$ via $g_1(\boldsymbol{\alpha}, \boldsymbol{Z}) = \alpha_1 + \alpha_2 I(\text{F})$. In this case, we use the male gender as the reference group, normalizing α_1 to one, as before, in obtaining d_1 and d_2 for this group. So the posterior distribution for α_1 deviating from one gives a sense of how on target the skeleton probabilities for the men were on the two doses. The posterior distribution of α_2 gives a sense of the effect of female gender on the toxicity rates for cycle 1.

3.2.4 Prior selection and posterior distribution

Based on the study design, patients contribute to the likelihood until they experience a DLT or the final Kth cycle is completed. That is, a person with toxicity on cycle K_i gives data $(Y_{i,1} = 0, Y_{i,K_i-1} = 0, \ldots, Y_{i,K_i} = 1, Z_i, d_{i,1})$ and a contribution to the likelihood as in equation 3.4. And a person completing K cycles without toxicity gives data $(Y_{i,1} = 0, \ldots, Y_{i,K-1} = 0, Y_{i,K} = 0, Z_i, d_{i,1})$ with likeli-

hood contribution as in equation 3.5. In general, subject i on cycle k contributes $L_{i,k}(Y_{i,k}|\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\rho}) = (p_{i,k})^{Y_{i,k}}(1-p_{i,k})^{1-Y_{i,k}}$ to the likelihood, with $p_{i,k}$ parameterized as in equation 3.1 and interpreted as the the probability of toxicity on cycle k conditional on having no prior DLTs in previous cycles. The resulting likelihood is,

$$L(Y|\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\rho}) = \prod_{i=1}^{N}\prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\rho}).$$

Our goal lies in estimating the posterior distribution of α, β and ρ and hence of $p_{i,k}$. Prior distributions on these parameters should reflect any auxiliary knowledge of the study design, with a large prior variance when this knowledge is limited. We first consider priors in the case with no covariates, so that $g_1(\boldsymbol{\alpha}, \boldsymbol{Z})=\alpha$, $g_2(\boldsymbol{\rho}, \boldsymbol{Z})=\rho$ and $g_3(\boldsymbol{\beta}, \boldsymbol{Z}) = \beta$. Modifications of this approach for more complex settings will be discussed after this more simple case is described. The priors are programmed using just another Gibbs sampler `(JAGS)` *rjags* [Plummer, 2011] package through [`R` Development Core Team, 2011].

3.2.4.1 Special case with no covariates

Recall that with no covariates, α relates to the probability of toxicity on cycle one, with $\alpha = 0$ implying no toxicity and increasing values of α giving larger toxicity probabilities. The skeleton described in Section 3.2.3 is calibrated to give α =1 when correctly specified. For many phase II trials the low dose is associated with very little

or no toxicity on the first cycle, so that a prior with a pre-specified point mass at $\alpha = 0$ and ranging across non-negative values is desirable. In Appendix 3.6.1.1, we describe our recommended prior on α, with cumulative distribution function (cdf) $F_\alpha(\alpha)$, mean $\mu_\alpha = 1$ and variance σ^2_α, as a mixture distribution of a lognormal density and a $q_\alpha \times 100\%$ point mass at zero. We denote the probability density function (pdf) of the lognormal component of the mixture distribution as $g_\alpha(x)$, with mean μ_g and variance σ^2_g. The cdf of α then becomes $F_\alpha(x) = q_\alpha + (1 - q_\alpha) \int_0^x g_\alpha(u)du$. The μ_g parameter of the lognormal distribution shifts as a function of the point mass percentage to maintain a mean one prior; that is, $\mu_g = (1 - q_\alpha)^{-1}$ yields $\mu_\alpha = 1$. The lognormal variance parameter, σ^2_g depends on the desired variance for the prior mixture, σ^2_α, as well as the point mass probability, q_α, that is, $\sigma^2_g = (1-q_\alpha)^{-1}\sigma^2_\alpha - \mu^2_g q_\alpha$. For convenience, an example of `JAGS` code for generating this prior is included on lines 10 through 17 of Appendix 3.6.1.3.

The ρ parameter ($0 \leq \rho \leq 1$) captures dependency in toxicity outcomes within a patient, with values near zero indicating that the current toxicity outcome is virtually unaffected by previous tolerance of dose and values near one indicating an almost certain chance of tolerating previously administered doses. In Appendix 3.6.1.2, we develop this prior as a mixture distribution of pre-specified $q_\rho \times 100\%$ point masses at 0 and 1 and a trapezoidal density with height b at zero and slope $m = 2(1 - b)$ comprising the remainder of the distribution over $(0, 1)$. The cdf of ρ is $F_\rho(x) = q_\rho + (1 - 2q_\rho) \int_0^x [b + 2(1 - b)u]\, du + q_\rho I(x = 1)$. The b parameter governs whether

the trapezoidal shape favors low or high values of ρ, with $b = 1$ reducing to a uniform shape and $b = 0$ or $b = 2$ reducing to triangular shapes with positive and negative slopes. An example of `JAGS` code for this prior is located on lines 18 through 30 of Appendix 3.6.1.3.

The remaining prior that we define is for the β parameter, which captures the effect of cumulative dose on the conditional probability of toxicity. When $\beta = 0$, the probability of toxicity does not change based on cumulative dose. When $\beta > 0$, toxicities are more likely to occur as dose accumulates. The model also allows the possibility of developing an increased tolerance for dose with repeated exposure, that is, β may be negative subject to the constraint that toxicity probabilities remain in the [0, 1] range. A lower bound for β is obtained by noting that $d_{i,k} = d_{i,1}$ for $k \geq 1$, in equation 3.3 giving $0 < \alpha(d_{i,1} - \rho d_{i,1}) + \beta(k-1)d_{i,1} < \infty, \forall k$. It is convenient to define the prior in terms of a shared boundary at all cycles, $\beta > -\alpha(1-\rho)/(K-1)$. In particular, for the ifosamide study we investigate later, $K = 4$ cycles and the lower boundary on the prior for β becomes $-\alpha(1-\rho)/3$, depending on α and ρ.

In constructing a prior for β conditional on α and ρ, with cdf $F_\beta(\beta|\alpha, \rho)$, we use a Normal$(\mu, \sigma^2)$ distribution truncated on the left by $-\alpha(1-\rho)/3$. To program this truncated Normal prior with mean μ_β and standard deviation σ_β in JAGS, we need to input mean μ and standard deviation σ of an untruncated Normal distribution giving mean μ_β and variance σ_β upon left truncation at $-\alpha(1-\rho)/3$. This can be achieved in the following way using JAGS. Expressions for deriving μ_β and σ_β based on known

parameters (μ, σ^2) of the untruncated Normal distribution exist [Johnson and Kotz, 1970, Greene, 2003]. The variance, σ, of the untruncated distribution is involved in calculating both μ_β and σ_β for the truncated distribution so that arbitrary specification of these prior parameters does not always give a viable untruncated distribution to work with. The `R` function, `findbetaroots`, in Appendix 3.6.1.4 solves for parameters (μ, σ^2) required by the `JAGS` program based on desired prior parameters for β, $(\mu_\beta, \sigma_\beta)$ or indicates that no possible solution exists for that combination of values. For convenience, Tables 3.8 and 3.9 lists helpful examples of prior parameters for β and the corresponding parameters of the untruncated distribution that are used in `JAGS`. An example of `JAGS` code for this prior on lines 31 through 37 of Appendix 3.6.1.3.

A few more definitions using Stieltjes notation help in characterizing the posterior distribution. Let $h_\alpha(\alpha) = q_\alpha I(\alpha = 0) + \{1 - q_\alpha\} I(\alpha > 0) g(\alpha)$ and $h_\rho(\rho) = q_\rho I(\rho = 0) + \{1 - 2q_\rho\} I(0 < \rho < 1)\{b + 2(1-b)\rho\} + q_\rho I(\rho = 1)$ capture either a probability mass or a density function as appropriate. In addition define $dF_\alpha(\alpha) = h_\alpha(\alpha) d\alpha$ and $dF_\rho(\rho) = q_\rho I(\rho = 0) + \{1 - 2q_\rho\} I(0 < \rho < 1)\{b + 2(1-b)\rho\} d\rho + q_\rho I(\rho = 1)$. Also let $f_\beta(\beta|\alpha, \rho)$ be the prior density function of β. The posterior distribution for α, β and ρ given the observed data Y is then $f(\alpha, \beta, \rho|Y) =$

$$\frac{\prod_{i=1}^{N} \prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha, \beta, \rho) f_\beta(\beta|\alpha, \rho) h_\alpha(\alpha) h_\rho(\rho)}{\int_0^1 \int_0^\infty \left[\int_{\frac{-\alpha(1-\rho)}{K-1}}^{\infty} \prod_{i=1}^{N} \prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha, \beta, \rho) f_\beta(\beta|\alpha, \rho) d\beta \right] dF_\alpha(\alpha) dF_\rho(\rho)}. \quad (3.8)$$

The posterior distribution of α, β, ρ from Model 3.2 can be estimated via Markov Chain Monte Carlo (MCMC) methods [Robert and Casella, 1999] using `JAGS` [Plummer, 2011] called in [`R` Development Core Team, 2011]. `JAGS` includes several algorithms for sampling from the posterior distributions produced from the MCMC iterations, for instance the standard Gibbs sampler is available for this purpose. Parallel chains starting from different initial values for each parameter (α, β, ρ) are followed through to convergence after an appropriate burn-in period. After convergence the posterior distributions of the parameters are available as well as functions of these parameters, such as the desired conditional toxicity profiles. The mean of the posterior distributions are used as estimates of the quantities of interest.

3.2.4.2 Covariate specific priors

Priors from section 3.2.4.1 can be easily extended to allow dependence of $\boldsymbol{\alpha}, \boldsymbol{\rho}, \boldsymbol{\beta}$ on $\mathbf{Z}$ in Model 3.1. As an instructive example, we again consider the case where $\mathbf{Z}$ includes gender (M, F) and dose group (d_1, d_2) so that $g_1(\boldsymbol{\alpha}, \boldsymbol{Z}) = \alpha_1 I(\mathrm{M}, d_1) + \alpha_2 I(\mathrm{M}, d_2) + \alpha_3 I(\mathrm{F}, d_1) + \alpha_4 I(\mathrm{F}, d_2)$ for a total of four required priors. Each of these four priors can be built just as in Section 3.2.4.1 if there is no prior information suggesting deviations from Model 3.2, i.e., the model that parameterizes a single α to account for dose-toxicity on cycle 1. We include an additional subscript to distinguish between priors for α. That is, with a q_{α_j} point mass at zero, the prior for $\alpha_j, j = 1, \ldots, 4$ has CDF $F_{\alpha_j}(x) = q_{\alpha_j} + \{1 - q_{\alpha_j}\} \int_0^x g_{\alpha_j}(u)du$, where density

function $g_{\alpha_j}(u)$ is a lognormal $\{(1-q_{\alpha_j})^{-1}, \sigma^2_{\alpha_j} = (1-q_{\alpha_j})^{-1}\sigma^2_{\alpha_j} - \mu^2_{g_j} q_{\alpha_j}\}$. When prior information on gender related toxicity is available, prior means of α_3 and α_4 may be chosen to reflect this additional knowledge. The priors for α_1 and α_2 are generally left with a mean of one since the skeleton discussed in Section 3.2.3 calibrated these values to one based on initial assumptions about dose-toxicity on cycle 1. In the case where we desire a prior with mean, μ_{α_j}, and variance, $\sigma^2_{\alpha_j}$, we would define density function $g_{\alpha_j}(x)$ as lognormal $\{\mu_{\alpha_j}(1-q_{\alpha_j})^{-1}, (1-q_{\alpha_j})(\mu^2_{g_j} q_{\alpha_j} + \sigma^2_{g_j})\}$.

Prior parameterization of $\boldsymbol{\rho}$ is technically straightforward using trapezoidal shapes described in section 3.2.4.1. However for limited sample sizes it makes sense to model a common prior for ρ. Priors for $\boldsymbol{\alpha}$ and ρ affect prior definition of $\boldsymbol{\beta}$. Recall that when parameterizing the prior for β in Section 3.2.4.1, the range of the prior was $[-\alpha(1-\rho)/(K-1), \infty)$, with negative β indicating an increased dose tolerance upon repeated exposure and positive β indicating increased toxicity with accumulating dose. Continuing our instructive example, suppose $g_3(\boldsymbol{\beta}, \boldsymbol{Z}) = \beta_1 I(\mathrm{M}, d_1) + \beta_2 I(\mathrm{M}, d_2) + \beta_3 I(\mathrm{F}, d_1) + \beta_4 I(\mathrm{F}, d_2)$. Priors for each of the four $\beta' s$ can be constructed as in Section 3.2.4.1, provided that the lower bound of each β_j is maintained to be consistent with values of α_j and ρ for those with the same covariates, $\mathbf{Z}$. In particular, in the case with gender and dose influencing all parameters, the range of β_j is restricted on the left by $-\alpha_j(1-\rho)/3, j = 1, \ldots, 4$.

Depending upon the number of priors set up on $\boldsymbol{\alpha}, \boldsymbol{\rho}, \boldsymbol{\beta}$ the likelihood and the posterior distribution will change accordingly in Equation 3.8.

For instance, assuming four subgroups for α_j and $\beta_j, j = 1, \ldots, 4$ that correspond to levels of gender and dose and assuming a common ρ across all $\mathbf{Z}$, the posterior distribution, $f(\alpha_1, \ldots, \alpha_4, \beta_1, \ldots, \beta_4, \rho|Y)$, becomes

$$\frac{\prod_{j=1}^{4}\prod_{i=1}^{N_j}\prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha_j, \beta_j, \rho) f_{\beta_j}(\beta_j|\alpha_j, \rho) h_{\alpha_j}(\alpha_j) h_\rho(\rho)}{\int_0^1 \left[\prod_{j=1}^{4}\int_0^\infty \left\{\int_{\frac{-\alpha_j(1-\rho)}{3}}^{\infty}\prod_{i=1}^{N_j}\prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha_j, \beta_j, \rho) f_{\beta_j}(\beta_j|\alpha_j, \rho) d\beta_j\right\} dF_{\alpha_j}(\alpha_j)\right] dF_\rho(\rho)}$$

where $\sum_{j=1}^{4} N_j = N$ is the sum of patients in the four categories of gender and dose.

3.2.5 Model selection

There is no restriction requiring the same covariates be included in parameterizations of $g_1(\boldsymbol{\alpha}, \boldsymbol{Z})$, $g_2(\boldsymbol{\rho}, \boldsymbol{Z})$ and $g_3(\boldsymbol{\beta}, \boldsymbol{Z})$. We recommend two common model selection criteria: (1) a plot of observed and predicted values of toxicity, along with 95% credible bands for the true probabilities of toxicity over the cycles. (2) The deviance information criteria (DIC) [Spiegelhalter et al., 2002] is calculated by adding the effective number of parameters (p_D) to the expected deviance, where the expected deviance is the deviance at the posterior mean parameter values and the effective number of parameters are estimated using the approach suggested by [Plummer, 2002, 2008]. Smaller DIC values indicate the preferred model.

3.3 Operating characteristics

To study operating characteristics of Model 3.2 we consider a trial with two dose groups (S_1, S_2) receiving a maximum of $K = 4$ cycles with 50 patients per group for a total of $N = 100$. Following Section 3.2.3, skeleton probabilities are set at $q_g = (0.05, 0.10)$, that is, 5% and 10% of patients are expected to have a DLT on the first cycle in the low and high dose groups, respectively. Then the transformed doses, $d_g, g = 1, 2$, used to stand in for $S_g, g = 1, 2$, in Model 3.2 become $d_1 = -\ln(1-0.05)$ and $d_2 = -\ln(1-0.10)$. Model 3.2 defines conditional probabilities of toxicity during the trial with $\alpha = 1$ and varying values of $\beta = \{0, 0.2\}$ and $\rho = \{0.25, 0.75\}$, i.e., 4 different simulated cases. Simulated toxicity outcomes across cycles are based on Bernoulli$(p_{i,k})$ random variables until a DLT is observed or the 4^{th} cycle is completed. Five hundred independent datasets (simulation replications) were created to assess coverage rates, bias and standard deviations.

Priors are set on α, β, ρ as in Section 3.2.4.1 for Model 3.2. In particular, the prior on α is a mixture with point mass $q_\alpha = 4\%$ on zero and a lognormal density component with parameters $\mu_g = 1.04$ and $\sigma_g^2 = 9.33$ giving prior mean and variance for α, $\mu_\alpha = 1$ and $\sigma_\alpha^2 = 9$. The prior for ρ is a mixture with point masses $q_\rho = 2\%$ at zero and one and an intercept $b = 0.20$ trapezoidal density, resulting in prior mean and variance $\mu_\rho = 0.6280$ and $\sigma_\rho^2 = 0.0736$. This prior indicates a moderate to high correlation in toxicity responses within-patient. The prior set on β is a Normal distribution truncated at $-1/3$ having mean, $\mu_\beta = 2$ and variance, $\sigma_\beta^2 = 4$,

providing a coefficient of variance of two. Since truncation of this distribution of β is conditional on prior values of α and ρ, current values from the MCMC simulation are used for the truncation point at each sampling. Based on Model 3.2, toxicity probabilities are also sampled and monitored for convergence; an adaptive phase of 1,000 samples is used to choose the best sampling algorithm and an additional 10K samples are discarded as part of the burn-in period.

The conditional probability of toxicity for each of the four simulated cases are presented in Table 3.1 and displayed in Figure 3.3. Each of the four columns of Table 3.1 indicate treatment cycles during the study. Rows are separated according to different parameter selections (cases) and dose group within case, d_1 or d_2. Reported values are (1) the estimated conditional probability of toxicity for an arbitrary patient i at cycle k, $\hat{p}_{i,k}$, (2) the mean bias, $\hat{p}_{i,k} - p_{i,k}$, across 500 iterations (3) the mean standard deviation (SD) of $\hat{p}_{i,k}$ across 500 iterations, (4) the empirical SD of the 500 $\hat{p}_{i,k}$ estimates, (5) the credible interval coverage rate of $p_{i,k}$ across the 500 iterations and (6) the mean number of patients who enter the following cycle toxicity-free. Within a particular case (1 - 4), the same posterior values of α, β and ρ are used to calculate finite sample characteristics of the eight cells of dose and cycle combinations.

Table 3.1 indicates that estimates of the conditional probability of toxicity have very low bias and that mean SD and empirical SD are comparable. This low bias is evident in Figure 3.3 where solid shapes (true conditional probabilities) and hollow shapes (estimated conditional probabilities) are very close to one another. In cases

1 and 3, where ρ takes on the lower value of 0.25, there is a higher decrease in the average number of patients making it through successive cycles as additional patients exhibit toxicity patterns that remove them from the study. Cases 2 and 4, with $\rho = 0.75$, have a higher tendency to avoid toxicity once they have tolerated their first cycle. In cases 3 and 4, where $\beta = 0.2$, there is a tendency for slightly more patients to discontinue due to accumulated toxicity, particularly in cycles 3 and 4. Case 4 shows the most impact of patients dropping out due to accumulating toxicity in cycles 3 and 4, since in this case the high value of $\rho = 0.75$ usually causes those with single dose susceptibility to be eliminated in cycle one rather than later cycles.

The major focus of estimation in these small studies is typically on estimated probabilities of toxicity, which seem to have little bias regardless of the model's ability to clearly identify individual parameter estimates. However, we summarize results from simulation studies of parameter estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\rho}$ in Table 3.2. Each row corresponds to one of the four simulated cases and presents (1) the mean of the estimated values, (2) the mean bias of the estimates from the true value of the parameter, (3) the mean SD of the parameter estimates, (4) the empirical SD from the 500 simulated datasets and (5) the coverage rate for the true parameter value in the credible intervals.

Higher values of ρ make it easier to isolate information on β since after tolerating cycle 1, toxicity observed after the first cycle is more likely to be caused by accumulating toxicity captured by β. This is reflected in Table 3.2 cases 2 and 4 where

$\rho = 0.75$ and bias for β is at its lowest. Lower values of ρ allow the current dose on every cycle to play a higher role in the manifestation of the toxicity, so that information on α is better identified. Table 3.2 cases 1 and 3, where $\rho = 0.25$, provide the lowest bias in estimation of α by correctly attributing toxicity to the effect of the current dose. The ability to estimate ρ, when the prior is not compatible with the true model, is a challenge in these small studies. In additional simulations, not shown, larger sample sizes do improve estimation of ρ, as well as the other parameters. As in all early phase studies, model assumptions are relied upon in making inferences. In the following section we perform additional sensitivity analyses for misspecification of the skeleton probabilities on estimation of conditional probability of toxicity.

3.3.1 Sensitivity to choice of skeleton probabilities

Given fixed values of $\alpha = 1$, $\beta = 0.2$ and $\rho = 0.75$, as in case 4 above, data was simulated as before. Recall that the probabilities of toxicity during the first cycle on doses 1 and 2 are 0.05 and 0.10 and that these probabilities were assumed in creating the skeleton used for analysis in the previous section. As opposed to the previous section, this section mis-specifies the probability skeleton when performing the analysis. In scenario 1, a skeleton that is 1.5 times higher than that used to generate the data is used in the analysis. That is, the probabilities of toxicity on cycle one are assumed to be 0.075 and 0.150 for doses 1 and 2, respectively. In

scenario 2, the probabilities of toxicity on cycle 1 were assumed to be 0.06 and 0.15 for the two dose levels, i.e., 20% and 50% overestimates of toxicity on cycle 1 for dose levels 1 and 2, respectively. Otherwise, the assumed priors and estimation procedure remained unchanged from the previous section.

Table 3.3 provides results on estimation of the conditional probabilities of toxicity across 500 iterations of the simulation study. The columns correspond to the four cycles. Within different skeleton misspecifications (scenarios 1 and 2) the rows are grouped by the two dose levels, d_1 or d_2. Reported values are (1) the estimated conditional probability of toxicity for an arbitrary patient i at cycle k, $\hat{p}_{i,k}$, (2) the mean bias, $\hat{p}_{i,k} - p_{i,k}$, across 500 iterations (3) the mean SD of $\hat{p}_{i,k}$ across 500 iterations, (4) the empirical SD of the 500 $\hat{p}_{i,k}$ estimates, (5) the credible interval coverage rate of $p_{i,k}$ across the 500 iterations and (6) the mean number of patients who enter the following cycle toxicity-free. Within a particular scenario (1 or 2), the same posterior values of α, β and ρ are used to calculate finite sample characteristics of the eight cells of dose and cycle combinations.

Results are comparable to those from case 4 in Table 3.1, where the correct skeleton was used in the analysis. The only exception is seen in scenario 2, cycle 1, dose 1, where bias is higher and coverage is lower than desired. Since the empirical and mean SD are very close in this case, the coverage is likely being effected by the bias for this term.

3.4 Application to the ifosamide study

The original ifosamide study was a phase II randomized clinical trial comparing the toxicity and efficacy of doxorubicin with high-dose ifosamide or standard-dose ifosamide in patients with soft-tissue sarcoma [Worden et al., 2005, Chugh et al., 2007]. The treatment was given for 4 consecutive days at the beginning of each 21 day cycle. The original study considered six cycles for metastatic disease and four cycles for localized disease but for convenience we consider just the first four cycles for each group. We evaluate 77 patients with data on toxicity, where 39 of these were randomized to the standard $6g/m^2$ ifosamide dose group and 38 of these were randomized to receive $12g/m^2$ of ifosamide. We use a patient's minimum hemoglobin (HGB) value during a cycle to define a DLT in this example, so that a patient is removed from the study if their HGB value drops below 8 mg. This criteria is defined as a grade 3/4 toxicity by NCI. Table 3.4 and Figure 3.1 present the empirical data from the study with conditional probabilities at each of the four cycles for the two dose groups.

Based on DIC criteria, Model 3.2 was improved by allowing differential cumulative effects of dose by dose group resulting in Model 3.9 of the form:

$$\ln(1 - p_{i,k}) = -\alpha(d_{i,1} - \rho d_{i,k-1}) - \beta_1 D_{i,k} I(d_1) - \beta_2 D_{i,k} I(d_2). \tag{3.9}$$

Resulting empirical and model-based conditional probabilities of toxicity by dose

group on each of four cycles are shown in Figure 3.4, along with 95% credible intervals. Priors for α and ρ used to perform the analysis are identical to those used in Section 3.3. Priors for β_1 and β_2 are identical to the prior used for β in Section 3.3.

Parameter estimates model are shown in Table 3.5 indicate indicate a particularly high cumulative effect of the dose in patients on the high group.

Following Section 3.2.2, the total expected completed cycles for the low and high dose groups, respectively, are 3.69 (3.68 observed on average) and 3.19 (3.23 observed on average) with corresponding expected total doses of 22.108 g/m^2 (22.06 g/m^2 observed on average) and 38.259 g/m^2 (38.70 g/m^2 observed on average). A patient in the high-dose group tends to receive more total ifosamide than a patient in the low-dose group before a DLT, but is less likely to successfully complete all four cycles without a DLT 29.3% and 71.3% respectively.

Upon further study, DLTs in the high-dose group are especially high in women, but not necessarily men, as seen in Figure 3.5 and Table 3.6. DIC criteria suggested further improvement in the model with inclusion of gender terms for the parameter α and gender by dose group terms for the cumulative toxicity. To account for the steep quadratic trend in toxicities for the females on the high-dose group a quadratic (squared) term was also included in the model and results in Model 3.10 as follows,

$$\begin{aligned}\log(1-p_{i,k}) = &-\{\alpha_1 + \alpha_2 I(F)\}(d_{i,1} - \rho d_{i,k-1}) - \beta_1 D_{i,k} I(M, d_1) \\ &- \beta_2 D_{i,k} I(F, d_1) - \beta_3 D_{i,k} I(M, d_2) - \beta_4 D_{i,k}^2 I(F, d_2).\end{aligned} \tag{3.10}$$

Figures 3.6 and 3.7 display the empirical and model-based conditional probabilities of toxicity, along with 95% credible intervals, by gender on the low and high dose groups, respectively, using Model 3.10. Priors were identical to those used in Section 3.3. That is, priors on α_1 and α_2 were identical to the prior on α in Section 3.3, the prior on ρ was left unchanged, and priors on β_1, β_2, β_3 and β_4 were identical to the prior used for β in Section 3.3.

Parameter estimates of the model are shown in Table 3.7 indicate that the toxicity due to cumulative effect of the dose is high in females in comparison to males on both the dose groups.

The total expected completed cycles for the low-dose female group is 3.55 (3.62 observed) with a total expected dose of 21.27 g/m^2 (21.71 g/m^2 observed). Women on the high-dose are expected to complete 2.95 cycles (2.97 observed) with a total expected dose of 35.44 g/m^2 (35.59 g/m^2 observed). Men are expected to complete roughly the same number of cycles in the low and high dose groups, 3.78 (3.58) and 3.46 (3.36) cycles, respectively, for total expected doses of 22.70 g/m^2 (21.49 g/m^2) and 41.51 g/m^2 (40.28 g/m^2) in the corresponding dose groups.

3.5 Discussion

We have presented a novel conditional probability model for the dose toxicity relationship in data arising from a Phase II study setting having patients with multiple cycles of the same dose over their treatment course. The conditional nature of

the model takes into account that patients having a DLT on a particular cycle will not continue further cycles. The use of ρ allows dependence of toxicity in a current cycle to depend on tolerance on previous cycles. The α and β parameters capture the effect of current and cumulative dose effects at each cycle. This three parameter model may be all that is estimable in small studies, but the model offers flexibility to include additional parameters to account for covariate dependent effects on toxicity.

Priors and skeletons described in this work offer a wide variety of prior beliefs to be included in the analyses. Our investigation of mis-specified skeleton probabilities showed very little effect on model performance.

One limitation of the model is that there must be sufficient cycles in the study to allow plausible estimation of ρ and β parameters that are based on data beyond cycle one; two cycles would not be sufficient to disentangle the effects of cumulative dose from tolerance to previous dose. As most studies of this type have between three and six cycles of therapy, this limitation should not impede use of the model in practice.

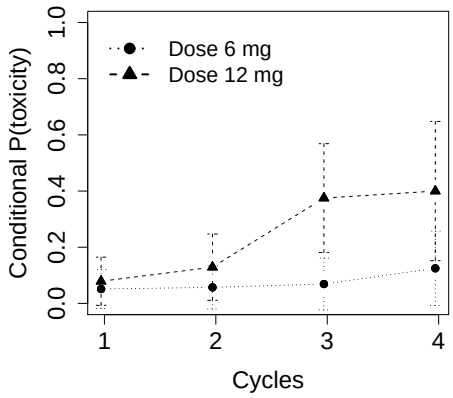

Figure 3.1: Observed proportion of low HGB in continuing subjects with one standard error intervals in the ifosamide trial.

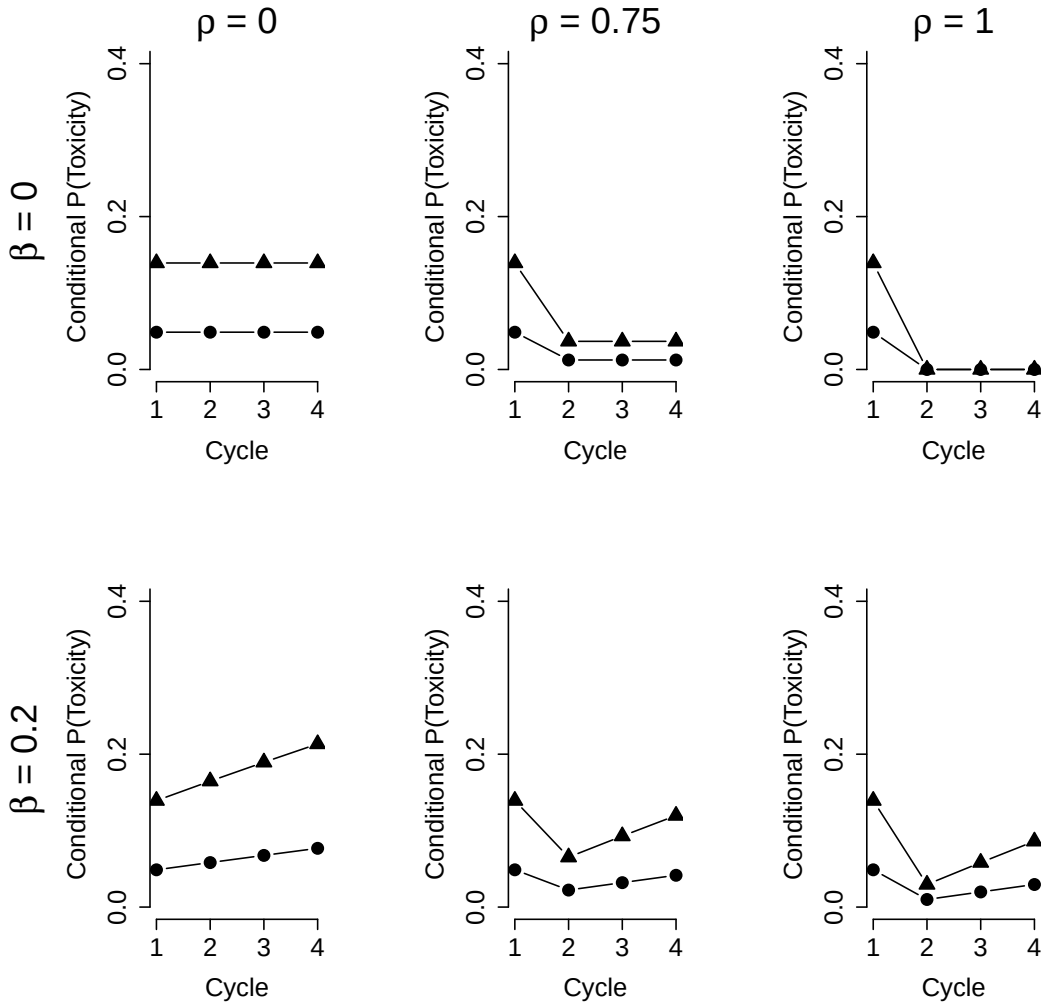

Figure 3.2: Conditional P(Toxicity) on cycle k, $p_{i,k}$ based on Model 3.2, for two dose levels with $\alpha = 1$, $\beta = (0, 0.2)$ and $\rho = (0, 0.75, 1)$. Solid triangles and circles correspond to the high and low dose groups respectively.

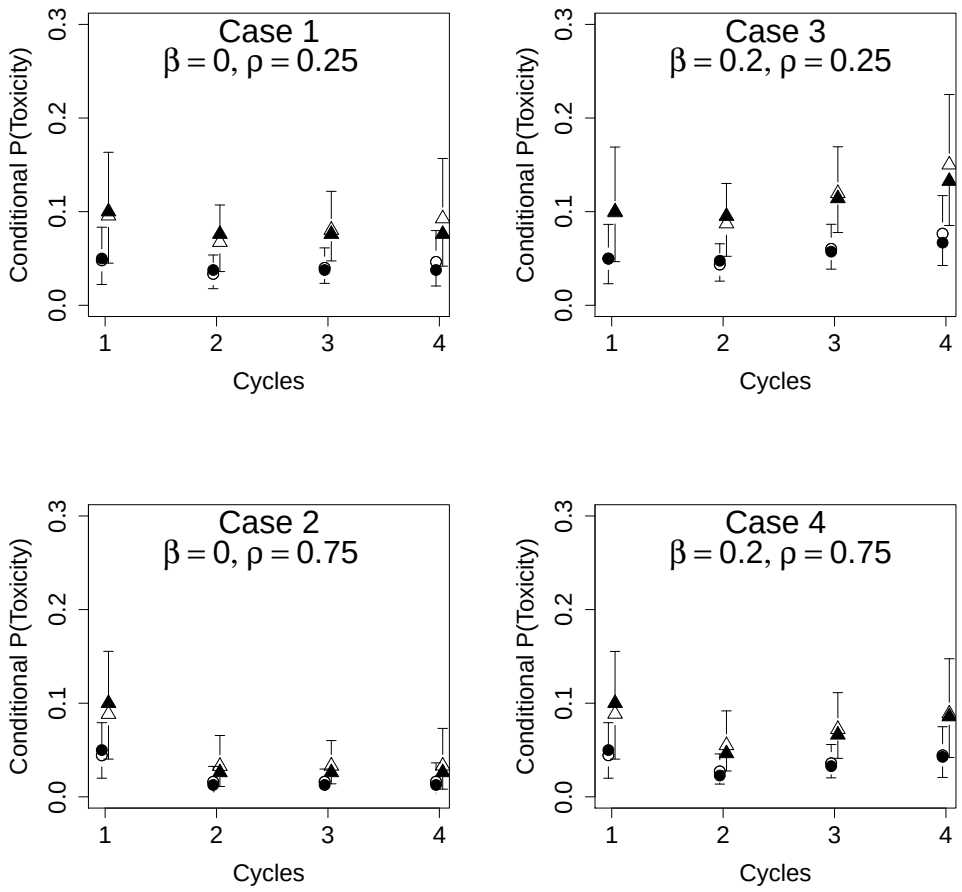

Figure 3.3: Simulation study results. Plot of conditional P(Toxicity), based on equation 3.2 for two dose levels with $\alpha = 1$ $\beta = (0, 0.2)$ and $\rho = (0.25, 0.75)$. Cases 1 through 4 in the Table 3.1. The panels in the first and second column correspond to $\beta = 0$ and $\beta = 0.2$ respectively while the top row corresponds to $\rho = 0.25$ and the bottom row to $\rho = 0.75$. The solid circles and triangles are the true values of the conditional P(toxicity) at each of the cycles on the lower and higher dose group respectively. The corresponding hollow circles and triangles are the average of the estimates of the values and the average of the limits of the 95% credible bands.

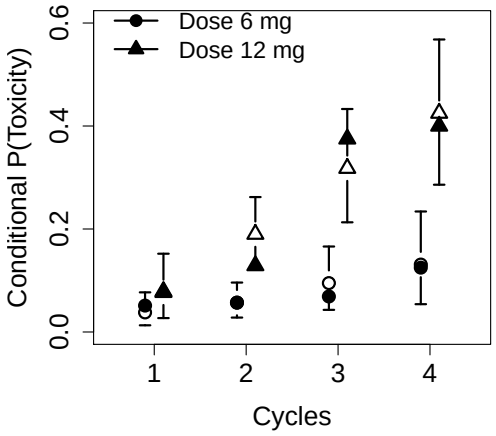

Figure 3.4: Estimates of the conditional probabilities of toxicity from Model 3.9 for patients on the low and high dose group in hollow circles and triangles respectively with observed values in solid circles and triangles

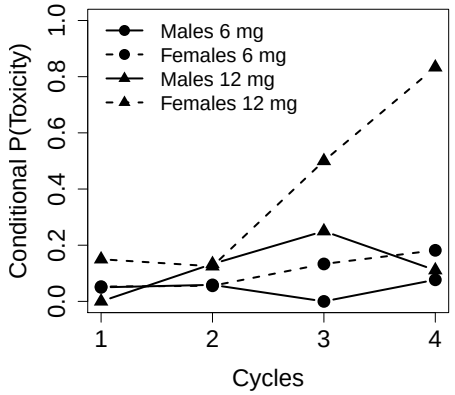

Figure 3.5: Empirical conditional probabilities of toxicity for patients on the low (in solid circles) and high (in solid triangles) dose groups by gender, males in solid lines and females in dotted lines.

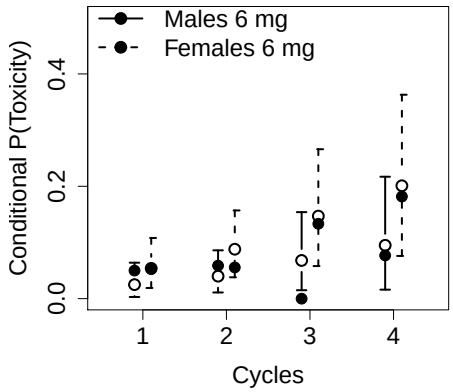

Figure 3.6: Estimates of the conditional probabilities of toxicity from Model 3.10 for patients on the **low dose group** in hollow circles with empirical values in solid circles, using two different β terms to estimate the cumulative effect for males (solid lines) and females (dotted lines).

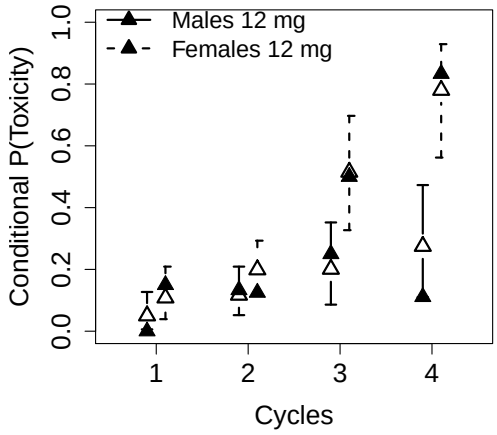

Figure 3.7: Estimates of the conditional probabilities of toxicity from Model 3.10 for patients on the **high dose group** in hollow triangles with empirical values in solid triangles, using two different β terms to estimate the cumulative effect for males (solid lines) and females (dotted lines).

Table 3.1: Estimates of conditional probability of toxicity based on Model 3.2 from 500 simulated datasets containing $N = 100$ patients receiving one of two dose groups d_1 and d_2 over four cycles. Results presented for $\alpha = 1$ and various combinations of $\beta = (0, 0.2)$ and $\rho = (0.25, 0.75)$

Cycle	1	2	3	4
	Case 1: $\alpha = 1$ $\beta = 0$ $\rho = 0.25$			
$\mathbf{d_1}$ Mean Estimated $\hat{p}_{ik}$	0.0480	0.0335	0.0400	0.0464
Mean Bias $\hat{p}_{ik}$	-0.0020	-0.0043	0.0023	0.0087
Mean SD $\hat{p}_{ik}$	0.0158	0.0093	0.0097	0.0153
Empirical SD $\hat{p}_{ik}$	0.0157	0.0077	0.0092	0.0145
Coverage rate of p_{ik}	93.0	94.6	95.2	92.8
Average patients	47.5	45.8	44.0	42.4
$\mathbf{d_2}$ Mean Estimated $\hat{p}_{ik}$	0.0956	0.0674	0.0803	0.0926
Mean Bias $\hat{p}_{ik}$	-0.0044	-0.0086	0.0043	0.0166
Mean SD $\hat{p}_{ik}$	0.0305	0.0183	0.0191	0.0297
Empirical SD $\hat{p}_{ik}$	0.0305	0.0152	0.0181	0.0282
Coverage rate of p_{ik}	93.0	94.6	95.2	92.8
Average patients	45.1	41.7	38.5	35.6
	Case 2: $\alpha = 1$ $\beta = 0$ $\rho = 0.75$			
$\mathbf{d_1}$ Mean Estimated $\hat{p}_{ik}$	0.0443	0.0161	0.0162	0.0163
Mean Bias $\hat{p}_{ik}$	-0.0057	0.0034	0.0035	0.0036
Mean SD $\hat{p}_{ik}$	0.0153	0.0071	0.0059	0.0085
Empirical SD $\hat{p}_{ik}$	0.0157	0.0051	0.0051	0.0073
Coverage rate of p_{ik}	89.8	99.2	94.6	96.2
Average patients	47.5	46.9	46.8	45.8
$\mathbf{d_2}$ Mean Estimated $\hat{p}_{ik}$	0.0883	0.0327	0.0330	0.0331
Mean Bias $\hat{p}_{ik}$	-0.0117	0.0067	0.0070	0.0071
Mean SD $\hat{p}_{ik}$	0.0297	0.0142	0.0119	0.0170
Empirical SD $\hat{p}_{ik}$	0.0305	0.0102	0.0103	0.0147
Coverage rate of p_{ik}	89.8	99.2	94.6	96.2
Average patients	44.9	43.9	42.7	41.6
	Case 3: $\alpha = 1$ $\beta = 0.2$ $\rho = 0.25$			
$\mathbf{d_1}$ Mean Estimated $\hat{p}_{ik}$	0.0498	0.0435	0.0603	0.0765
Mean Bias $\hat{p}_{ik}$	-2e-04	-0.0040	0.0030	0.0096
Mean SD $\hat{p}_{ik}$	0.0163	0.0103	0.0123	0.0191
Empirical SD $\hat{p}_{ik}$	0.0155	0.0086	0.0122	0.0183
Coverage rate of p_{ik}	95.2	95.0	95.0	92.6
Average patients	47.4	45.1	42.4	39.5
$\mathbf{d_2}$ Mean Estimated $\hat{p}_{ik}$	0.0991	0.0872	0.1195	0.1501
Mean Bias $\hat{p}_{ik}$	-9e-04	-0.0081	0.0054	0.0176
Mean SD $\hat{p}_{ik}$	0.0315	0.02	0.0235	0.0359
Empirical SD $\hat{p}_{ik}$	0.0301	0.0169	0.0234	0.0344
Coverage rate of p_{ik}	95.2	95.0	95.0	92.6
Average patients	44.9	40.7	36.1	31.3
	Case 4: $\alpha = 1$ $\beta = 0.2$ $\rho = 0.75$			
$\mathbf{d_1}$ Mean Estimated $\hat{p}_{ik}$	0.0444	0.0272	0.0359	0.0444
Mean Bias $\hat{p}_{ik}$	-0.0056	0.0044	0.0031	0.0018
Mean SD $\hat{p}_{ik}$	0.0153	0.0084	0.0091	0.0139
Empirical SD $\hat{p}_{ik}$	0.0162	0.0064	0.0088	0.0134
Coverage rate of p_{ik}	89.8	98.0	95.2	96.4
Average patients	47.5	46.5	44.9	43.1
$\mathbf{d_2}$ Mean Estimated $\hat{p}_{ik}$	0.0885	0.0549	0.0721	0.0887
Mean Bias $\hat{p}_{ik}$	-0.0115	0.0086	0.0060	0.0031
Mean SD $\hat{p}_{ik}$	0.0297	0.0166	0.0180	0.0271
Empirical SD $\hat{p}_{ik}$	0.0315	0.0128	0.0173	0.0261
Coverage rate of p_{ik}	89.8	98.0	95.2	96.4
Average patients	45.1	42.9	40.1	36.8

[1] SD refers to the standard deviation of the estimates.

Table 3.2: Estimates of parameters based on Model 3.2 from 500 simulated datasets containing $N = 100$ patients receiving one of two dose groups d_1 and d_2 over four cycles. Results presented for $\alpha = 1$ and various combinations of $\beta = (0, 0.2)$ and $\rho = (0.25, 0.75)$.

	α	β	ρ
Case 1 : $\alpha = 1, \beta = 0, \rho = 0.25$			
Estimated Value	0.9651	0.1332	0.4370
Mean Bias	-0.0349	0.1332	0.1870
Mean SD	0.3260	0.1672	0.2652
Empirical SD	0.3232	0.1446	0.1357
Coverage rate	93.0	89.2	99.6
Case 2 : $\alpha = 1, \beta = 0, \rho = 0.75$			
Estimated Value	0.8886	0.0022	0.6025
Mean Bias	-0.1114	0.0022	-0.1475
Mean SD	0.3153	0.1027	0.2560
Empirical SD	0.3227	0.0729	0.1375
Coverage rate	89.8	99.8	99.8
Case 3 : $\alpha = 1, \beta = 0.2, \rho = 0.25$			
Estimated Value	1.0018	0.3455	0.4649
Mean Bias	0.0018	0.1455	0.2149
Mean SD	0.3382	0.1982	0.2761
Empirical SD	0.3203	0.1608	0.1287
Coverage rate	95.2	93.2	99.8
Case 4 : $\alpha = 1, \beta = 0.2, \rho = 0.75$			
Estimated Value	0.8909	0.1758	0.5626
Mean Bias	-0.1091	-0.0242	-0.1874
Mean SD	0.3153	0.1438	0.2724
Empirical SD	0.3331	0.1193	0.1329
Coverage rate	89.8	98.0	99.0

Table 3.3: Estimates of conditional probability of no toxicity based on Model 3.2 from 500 simulated datasets containing $N = 100$ patients receiving one of two dose groups d_1 and d_2 over four cycles. Results presented for $\alpha = 1, \beta = 0.2$ and $\rho = 0.75$ for different mis-specified probability skeletons.

Cycle	1	2	3	4
	Scenario 1: $d_1 = 0.0780, d_2 = 0.1625$			
d_1				
Mean Estimated $\hat{p}_{ik}$	0.0456	0.0271	0.0357	0.0441
Mean Bias $\hat{p}_{ik}$	-0.0044	0.0043	0.0029	0.0014
Mean SD $\hat{p}_{ik}$	0.0155	0.0084	0.0091	0.0139
Empirical SD $\hat{p}_{ik}$	0.0157	0.0064	0.0088	0.0132
Coverage rate of true p_{ik}	93.2	97.6	94.8	96.2
Average patients	47.5	46.432	44.8	42.9
d_2				
Mean Estimated $\hat{p}_{ik}$	0.0922	0.0556	0.0728	0.0893
Mean Bias $\hat{p}_{ik}$	-0.0078	0.0093	0.0066	0.0036
Mean SD $\hat{p}_{ik}$	0.0304	0.0170	0.0181	0.0273
Empirical SD $\hat{p}_{ik}$	0.0308	0.013	0.0176	0.0261
Coverage rate of true p_{ik}	94.4	97.6	94.6	96.6
Average patients	45.0	42.9	40.1	36.7
	Scenario 2: $d_1 = 0.0620, d_2 = 0.1625$			
d_1				
Mean Estimated $\hat{p}_{ik}$	0.0389	0.0232	0.0306	0.0379
Mean Bias $\hat{p}_{ik}$	-0.0111	4e-04	-0.0022	-0.0048
Mean SD $\hat{p}_{ik}$	0.0133	0.0072	0.0078	0.0119
Empirical SD $\hat{p}_{ik}$	0.0135	0.0055	0.0076	0.0114
Coverage rate of true p_{ik}	84.2	97.6	93.4	92.8
Average patients	47.5	46.4	44.8	42.9
d_2				
Mean Estimated $\hat{p}_{ik}$	0.0982	0.0596	0.0781	0.0959
Mean Bias $\hat{p}_{ik}$	-0.0018	0.0133	0.012	0.0102
Mean SD $\hat{p}_{ik}$	0.0323	0.0181	0.0194	0.0292
Empirical SD $\hat{p}_{ik}$	0.0329	0.0139	0.0190	0.0280
Coverage rate of true p_{ik}	93.0	96.0	92.2	95.0
Average patients	45.0	42.9	40.1	36.7

Table 3.4: Grade 3/4 dose dose limiting toxicities (DLTs) observed when the hemoglobin levels dropped below 8mg on the two dose groups in patients completing the previous cycle without a DLT over the four cycles of treatment.

	Cycle 1	Cycle 2	Cycle 3	Cycle 4
dose 6	2/39	2/35	2/29	3/24
dose 12	3/38	4/31	9/24	6/15

Table 3.5: Parameter estimates of the parameters with SD obtained using Model 3.9

	Estimate	SD
α	0.76	0.33
β_1	0.81	0.37
β_2	1.66	0.44
ρ	0.54	0.30
DIC	167.90	-

Table 3.6: Dose limiting toxicities (DLTs), grouped by gender, inpatients completing the previous cycle without a DLT when the hemoglobin levels dropped below 8mg over four cycles of the treatment

	Cycle 1	Cycle 2	Cycle 3	Cycle 4
	Males			
dose 6	1/20	1/17	0/14	1/13
dose 12	0/18	2/15	3/12	1/9
	Females			
dose 6	1/19	1/18	2/15	2/11
dose 12	3/20	2/16	6/12	5/6

Table 3.7: Parameter estimates of the parameters with SD obtained using Model 3.10

	Estimate	SD
α_1	0.49	0.33
α_2	0.60	0.53
β_1	0.58	0.42
β_2	1.33	0.65
β_3	0.97	0.43
β_4	15.62	4.70
ρ	0.55	0.30
DIC	160.49	-

3.6 Appendix

3.6.1 Further details on prior distributions discussed in Section 3.2.4.1

3.6.1.1 Prior distribution for α

Recall that α is a mixture distribution of a lognormal random variable and a point mass at zero of size q_α. In the case with no covariates, the prior mean of α should be 1.0 to be consistent with skeleton calibration beliefs described in Section 3.2.3. When α depends on covariates, we may desire more flexibility in defining prior means. In this section, we derive parameters for the lognormal component of the mixture distribution that give desired means and variances of the overall mixture distribution. Let X be the lognormal random variable in the mixture distribution of α with density $g_\alpha(x)$, mean μ_g and variance σ_g^2. If B is Bernoulli(q_α) then the mixture random variable we desire can be written as $\alpha = (1 - B) * X$. The mean of

α is

$$\mu_\alpha = E\{E(\alpha|B)\} = E\{(1-B)E(X)\} = \mu_g * (1-q_\alpha).$$

When $q_\alpha = 1$, $E(\alpha)$ becomes 0 as we would expect if the prior was a point mass at zero, and $q_\alpha = 0$ gives $E(\alpha) = \mu_g$, the mean of the lognormal, X. For $0 < q_\alpha < 1$, we can obtain a mean μ_α prior for α by defining the mean of X to be $\mu_g = \mu_\alpha(1-q_\alpha)^{-1}$. The variance of α is

$$\begin{aligned}
\sigma_\alpha^2 &= V\{E(\alpha|B)\} + E\{V(\alpha|B)\} \\
&= V\{\mu_g(1-B)\} + E\{(1-B)^2\sigma_g^2\} \\
&= \mu_g^2 q_\alpha(1-q_\alpha) + \sigma_g^2(1-q_\alpha) \\
&= (1-q_\alpha)(\mu_g^2 q_\alpha + \sigma_g^2)
\end{aligned}$$

Again, $q_\alpha = 1$ gives $\sigma_\alpha^2 = 0$ (pointmass) and $q_\alpha = 0$ gives $\sigma_\alpha^2 = \sigma_g^2$, the variance of the lognormal. For $0 < q_\alpha < 1$, we can obtain prior variance σ_α^2 for α by defining $\sigma_g^2 = (1-q_\alpha)^{-1}\sigma_\alpha^2 - \mu_g^2 q_\alpha$. Appendix 3.6.1.3, includes code for this prior on lines 10 through 17.

3.6.1.2 Prior distribution for ρ

Recall that $\rho \in [0,1]$ follows a mixture distribution with $q_\rho \times 100\%$ point masses at 0 and 1 and a trapezoidal density shape for $0 < \rho < 1$. The family of trapezoidal density shapes included in the mixture distribution for ρ is a special case of that

seen in [van Dorp and Kotz, 2003] and [van Dorp et al., 2007]. In this section we define the mixture distribution for ρ in more detail and give means and variances associated with this prior distribution.

We define the random variable, X, to have density $g_\rho(x) = I(0 < x < 1)(mx + b)$, with b and m standing in for the intercept and slope of the line that determines the trapezoidal density shape over $(0, 1)$. To be a proper density function that integrates to 1.0 the constraint $m = 2(1 - b)$ must be satisfied. Hence, although m is convenient for defining the trapezoidal shape, $g_\rho(x)$ may be defined in terms of b only, $g_\rho(x) = b + 2(1 - b)x$. The expected value and variance of X are found to be $(4 - b)/6$ and $(2 + 2b - b^2)/36$ respectively. Special cases include (i) b=0, giving $g_\rho(x) = 2x$, with a positive slope and higher mass on values favoring one (mean=2/3); (ii) b=1, giving $g_\rho(x) = 1$, a Uniform distribution, (mean=1/2); and (iii) $b = 2$, giving $g_\rho(x) = 2 - 2x$ with a negative slope and higher mass on values closer to zero (mean=1/3).

To incorporate X into a mixture distribution with point masses at 0 and 1, we define a multinomial random variable $M = (M_1, M_2, M_3)$, where M_1, M_2 and M_3 are dependent Bernoulli random variables with $M_1 + M_2 + M_3 = 1$ and probabilities $(q_\rho, 1 - 2q_\rho, q_\rho)$.

The mixture distribution of ρ can be written as $M_2X + M_3$, that is,

$$\rho = \begin{cases} 0 & \text{when } M_1 = 1, \text{ with probability } q_\rho \\ \text{X} & \text{when } M_2 = 1, \text{ with probability } 1 - 2q_\rho \\ 1 & \text{when } M_3 = 1, \text{ with probability } q_\rho \end{cases}$$

The CDF of ρ becomes,

$$\begin{aligned} F_\rho(x) &= P(M_1 = 1) + P(\rho \leq x | M_2 = 1)P(M_2 = 1) + I(x = 1)P(M_3 = 1) \\ &= q_\rho + (1 - 2q_\rho) \int_0^x [b + 2(1 - b)u]\, du + q_\rho I(x = 1). \end{aligned}$$

Also,

$$\begin{aligned} E(\rho) &= E(M_2X + M_3) \\ &= E(M_2)E(X) + E(M_3) \\ &= (1 - 2q_\rho)E(X) + q_\rho \\ &= (1 - 2q_\rho)(\frac{4 - b}{6}) + q_\rho. \end{aligned}$$

and

$$\begin{aligned}
V(\rho) &= V\{E(\rho|M)\} + E\{V(\rho|M)\} \\
V\{E(\rho|M)\} &= V\{E(M_2X + M_3)|M\} \\
&= V\{M_2E(X) + M_3)\} \\
&= E(X)^2V(M_2) + V(M_3) + 2E(X)Cov(M_2, M_3) \\
&= E(X)^2(1-2q_\rho)2q_\rho + q_\rho(1-q_\rho) - 2E(X)(1-2q_\rho)q_\rho \\
&= 2E(X)(1-2q_\rho)q_\rho\{E(X)-1\} + q_\rho(1-q_\rho)
\end{aligned}$$

$$\begin{aligned}
E\{V(\rho|M)\} &= E\{V(M_2X + M_3)|M\} \\
&= E\{V(M_2X)|M\} \\
&= E\{M_2^2V(X)\} \\
&= V(X)E(M_2^2) \\
&= V(X)\{1-2q_\rho\}
\end{aligned}$$

$$
\begin{aligned}
V(\rho) &= 2E(X)(1-2q_\rho)q_\rho\{E(X)-1\} + q_\rho(1-q_\rho) + V(X)\{1-2q_\rho\} \\
&= 2\left(\frac{4-b}{6}\right)(1-2q_\rho)q_\rho\left\{\frac{4-b}{6}-1\right\} + q_\rho(1-q_\rho) + \{1-2q_\rho\}\left\{\frac{2+2b-b^2}{36}\right\} \\
&= q_\rho(1-q_\rho) - 2\left(\frac{(4-b)(b+2)}{36}\right)(1-2q_\rho)q_\rho + \{1-2q_\rho\}\left\{\frac{2+2b-b^2}{36}\right\} \\
&= \frac{3-(1-b)^2}{36} + \left(\frac{1}{2} + \frac{-(8+2b-b^2)-(2+2b-b^2)}{36}\right)2q_\rho + \left(\frac{4(8+2b-b^2)}{36}-1\right)q_\rho^2 \\
&= \frac{3-(1-b)^2}{36} + \left(\frac{1}{2} + \frac{-10-4b+2b^2)}{36}\right)2q_\rho + \left(\frac{(8+2b-b^2)-9}{9}\right)q_\rho^2 \\
&= \frac{3-(1-b)^2}{36} + \left(\frac{18-10-4b+2b^2)}{36}\right)2q_\rho + \left(\frac{(-1+2b-b^2}{9}\right)q_\rho^2 \\
&= \frac{3-(1-b)^2}{36} + \frac{b^2-2b+4}{9}q_\rho - \frac{(1-b)^2}{9}q_\rho^2
\end{aligned}
$$

When $q_\rho = 0$, the mean and variance of ρ are those of the trapezoidal distribution, X, defined earlier. Table 3.10 summarizes values of $E(\rho)$, $V(\rho)$ and m by choices for b and q_ρ. Appendix 3.6.1.3, includes code for this prior on lines 18 through 30.

3.6.1.3 Outline of the code written in JAGS

The code presented below corresponds to the applying the model to the data as presented in Section 3.4 for the simple case with no covariates.

```
1  #Defining the model.bug file
2  model {
3  c <- 1000 #a constant used in defining the mixture prior
4  #Define the likelihood for each of the N subjects
5  for (i in 1:N) {
```

```
prob[i] <- 1- exp( -alpha* ( dose[i] - rho*maxprevdose[i])*step(dose[i] -rho*maxprevdose[i])
- beta*cumdose[i])
response[i] ~ dbern(prob[i]) }
#Setting up the priors
#prior on α - E(α) = 1 and Var(α) = 9 and q_α = 0.04, the point mass at zero
Using these values μ_g = (1 − 0.04)^{-1} = 1.04 and σ_g^2 = (1 − q_α)^{-1}σ_α^2 − μ_g^2 q_α = 9.33
the mean and variance for the random variable X described in 3.6.1.1.
mu1 <- -1.093
tau1 <- 0.4416
alpha1 ~ dlnorm(mu1,tau1)
alphazero ~ dbern(0.96) #point mass of 0.04 at zero
alpha <- (1-alphazero)*0 + alphazero*alpha1
#prior on ρ based on 3.6.1.2
intercept <- 0.20 ; slope <- 2 - 2*intercept
zero1 ~ dpois(phi1)
phi1<- -log(intercept + slope*(rho1)) +c
rho1 ~ dunif(a1,b1)
a1 <- 0   b1 <- 1
#Mimicking the Multinomial by using a Uniform distribution for the two point
masses
u ~ dunif(0,1)
z1 <−(u <= 0.04)
z2 <−(u > 0.04&u < 0.96)
z3 <−(u >= 0.96)
rho <− z1 ∗ 0 + z2 ∗ rho1 + z3 ∗ 1
#Truncated Normal prior on β - μ_beta = −5.44 and σ_β^2 = 1/0.047 = 21.37
# Using the zeroes trick to simulate a non-standard prior
zero ~ dpois(phi)
phi <- 0.047*0.5*pow(beta+5.44,2) +c
beta ~ dunif(a2,b2)
a2 <- -alpha*(1-rho)/4
b2 <- 1000
#calculate the probabilities of toxicity
#define the two dose groups
d1 <−-log(1-0.05)
d2 <−-log(1-0.10)
#probability of no toxicity on each of the four cycles for the lower dose group
prob11 <- exp(-alpha*d1)
```

```
prob12 <- exp(-alpha*(d1-rho*d1)*step(d1-rho*d1) - beta*d1)
prob13 <- exp(-alpha*(d1-rho*d1)*step(d1-rho*d1) - beta*2*d1)
prob14 <- exp(-alpha*(d1-rho*d1)*step(d1-rho*d1) - beta*3*d1)
#probability of no toxicity on each of the four cycles for the higher dose group
prob21 <- exp(-alpha*d2)
prob22 <- exp(-alpha*(d2-rho*d2)*step(d2-rho*d2) - beta*d2)
prob23 <- exp(-alpha*(d2-rho*d2)*step(d2-rho*d2) - beta*2*d2)
prob24 <- exp(-alpha*(d2-rho*d2)*step(d2-rho*d2) - beta*3*d2)
}
#Initializing the parameters
inits<-list(list(zero=0,zero1=0,alpha1=1,beta=0.1,rho1=0.2,u=0.2,alphazero=0 ),
list(zero=0,zero1=0,alpha1=.5,beta=0,rho1=0.8,u=0.2,alphazero=0 ) )
parameters <-c("alpha","beta", "rho","prob11","prob12","prob13","prob14",
"prob21","prob22","prob23","prob24")
#updating the simulations
data <-list("response"=response,"maxprevdose"=maxprevdose,
"cumdose"=cumdose,"dose"=dose,"N"=N)
jags <- jags.model(file="prior.bug", data = data, inits=inits, n.chains = 2, n.adapt
= 5000)
adapt(jags,n.iter=1000)
update(jags,20500) # burin samples
sim1<-coda.samples(jags, parameters, 100000,thin=20)
#check for convergence
plot(sim1)
gelman.plot(sim1)
gelman.diag(sim1)
geweke.plot(sim1)
geweke.diag(sim1)
autocorr(sim1)
autocorr.plot(sim1)
#Monitor DIC
sim1.dic<-dic.samples(jags, n.iter=40000,n.thin=35, type="pD")
#report the mean and quantiles of the posterior distributions
y3<-summary(sim1)
ystat<-data.frame(y3statistics)
yquant = data.frame(y3quantiles)
```

3.6.1.4 Function code in R to find the parameters of the truncated Normal for β prior

```
# x[1] is the mean μ_β and x[2] is the standard deviation or σ_β
# y[1] gives the mean μ and y[2] the corresponding σ.
findbetaroots = function(x) {
y = numeric(2)
a.low=0 #the truncation point when ρ = 0
a.low=-1/3 # the truncation point when ρ ≠ 0
pdf.fun = dnorm( (a.low-x[1]/x[2]),0,1)
cdf.fun= 1-pnorm((a.low-x[1])/x[2],0,1)
y[1]=x[1] + x[2]*pdf.fun/cdf.fun - 3
alp= a.low - x[1]/ x[2]
delta.alp=pdf.fun/cdf.fun * (pdf.fun/cdf.fun - alp)
y[2]=x[2]^2 * (1-delta.alp) - 3
}
xstart = c(-2,5)
out=nleqslv(xstart, findbetaroots, control=list(btol=.01))
```

Table 3.8: Parameters for obtaining truncated Normal prior, $f_\beta(\beta|\alpha,\rho) \sim N(\mu,\sigma^2)$ on β from a Normal$(\mu_\beta,\sigma^2_\beta)$ distribution truncated at zero.

μ/σ^2	0.5	1	2	3	4	5
1	0.40, 1.10	-	-	-	-	-
2	2.00, 0.51	1.89, 1.21	0.80, 4.39	-5.16, 17.32	-	-
3	2.99, 0.50	2.99, 1.01	2.89, 2.30	2.57, 4.28	1.83, 7.50	0.31, 13.08
4	4.00, 0.50	3.99, 1.00	3.98, 2.04	3.93, 3.28	3.79, 4.82	3.54, 6.81
5	5.00, 0.50	4.99, 1.00	4.99, 2.00	4.99, 3.06	4.95, 4.22	4.89, 5.56

Table 3.9: Parameters for obtaining truncated Normal prior, $f_\beta(\beta|\alpha,\rho) \sim N(\mu,\sigma^2)$ on β from a Normal$(\mu_\beta,\sigma^2_\beta)$ distribution truncated at $-1/3$, $\alpha = 1$, $\rho = 0$, $k = 3$.

μ/σ^2	0.5	1	2	3	4	5
1	0.90, 0.64	-0.26, 2.68	-	-	-	-
2	2.00, 0.50	1.96, 1.09	1.54, 3.08	-0.028, 7.73	-5.44, 21.37	-42.64, 109.16
3	2.99, 0.50	2.99, 1.00	2.95, 2.16	2.77, 3.75	2.83, 6.06	1.64, 9.52
4	4.00, 0.50	3.99, 1.00	3.99, 2.02	3.96, 3.17	3.88, 4.53	3.72, 6.21
5	5.00, 0.50	5.00, 1.00	4.99, 2.00	4.99, 3.03	4.97, 4.15	4.93, 5.38

Table 3.10: Values of slope $m, E(\rho)$ and $V(\rho)$ for different values of the intercept b and point mass q_ρ based on the expressions derived in Section 3.6.1.2

	q_ρ=0	q_ρ=0.02	q_ρ=0.04
$b = 0, m = 2$	0.6670, 0.0556	0.6600, 0.0644	0.6533, 0.0732
$b = 0.2, m = 1.6$	0.6330, 0.0656	0.6280, 0.0736	0.6227, 0.0816
$b = 0.4, m = 1.2$	0.6000, 0.0733	0.5960, 0.0808	0.5920, 0.0882
$b = 0.6, m = 0.8$	0.5667, 0.0789	0.5640, 0.0859	0.5613, 0.0929
$b = 0.8, m = 0.4$	0.5333, 0.0822	0.5320, 0.0890	0.5307, 0.0957
$b = 1, m = 0$	0.5000, 0.0833	0.5000,0.0900	0.5000, 0.0967
$b = 1.2, m = -0.4$	0.4667, 0.0822	0.4680, 0.0890	0.4693, 0.0957
$b = 1.4, m = -0.8$	0.4333, 0.0789	0.4360, 0.0859	0.4387, 0.0929
$b = 1.6, m = -1.2$	0.4000, 0.0733	0.4040, 0.0808	0.4080, 0.0882
$b = 1.8, m = -1.6$	0.3667, 0.0656	0.3720, 0.0736	0.3773, 0.0816
$b = 2.0, m = -2.0$	0.3333, 0.0556	0.3400, 0.0644	0.3467, 0.0732

CHAPTER 4

Adaptive Phase I Clinical Trial Design Using Markov Models in Oncology for Patients with Ordinal Outcomes with Repeated Measures

4.1 Introduction/Background

Most dose finding clinical trial designs including the '3+3' [Storer, 1989] and the continual reassessment method (CRM) [O'Quigley et al., 1990] are based on dichotomizing the response into either a dose limiting toxicity (DLT) or no DLT, typically using the Common Toxicity Criteria defined by National Cancer Institute [NCI, 2003] that classifies grade 3 or higher toxic response as DLT. The literature that considers expanded levels of toxicity [Iasonos et al., 2011, Lee et al., 2012, Ivanova and Kim, 2009] confirms loss of information by grouping grade 1 or 2 toxicities into the 'no DLT' category. There are many examples of the extension of the CRM to an ordinal response [Lee et al., 2011, Van Meter et al., 2012, 2011, Bekele and Thall,

2004, Yuan et al., 2007, Wang et al., 2000] showing benefits in using ordinal response outcomes but are primarily for trials involving single dose administrations of the drug to the patient.

Instead of allowing patients to receive only a single dose, [Simon et al., 1997] provided the rationale for the accelerated titration design where patients could have intra-patient dose escalation. A random effects models with a continuous response was considered in this design with cut-offs for classifying the toxicity into four categories. This model was used to simulate data for the evaluation of the accelerated titration method, but the model was not used for data analysis. Motivated by considerations of pharmacokinetics [Legedza and Ibrahim, 2000] developed a model for repeated toxicity measures for each patient but dichotomized the response at each cycle for the patients. Both these methods assumed that the dose would remain constant for the patient on every cycle in the study.

There are ethical benefits in allowing a patient to receive a higher dose on subsequent cycles if no DLTs are manifested in the previous cycles. Alternatively the dose could be reduced if a higher than expected number of toxicities are observed in other patients in the study. Such a study conduct with repeated doses per patient in a binary outcome setting has been described in Chapter 2. Potential efficiency gains in estimating the dose-toxicity relationship through use of dose escalation/de-escalation have also been demonstrated in Chapter 2 especially in the setting of small sample sizes.

The use of ordinal outcomes is especially important in the setting when a patient is allowed to have multiple doses with possible dose escalation schemes so that the dose could be de-escalated or kept the same on the next cycle if a mild toxicity is observed on the current cycle. More recently [Doussau et al., 2013] provided a mixed effects proportional odds model to incorporate ordinal outcomes in a phase I setting to identify the probability of a severe toxicity and trend in risk of toxicity with time at the end of the trial. This method does not explicitly model the tendency to discontinue cycles for patients who have demonstrated previous DLT, although the resulting estimated toxicity rates may be conditional in nature. In addition the cumulative effect of the dose is not captured and patients are not allowed to escalate or de-escalate doses. In this setting where there are multiple cycles, doses may change and there are more than two levels of toxicity a large number of parameters may be necessary to capture all features of the dose-toxicity relationship. Furthermore, with the typically small sample size available for modeling purposes, a Bayesian approach is an attractive alternative. The ordinal Markov model presented in this chapter is an extension to the dichotomous Markov Model 2.1 from Chapter 2 which is an example of transitional models with Markov chains [Agresti, 2002] considering the toxicity on previous cycles.

In this chapter we will discuss extensions to the Markov Model in Chapter 2 for ordinal outcomes in phase I clinical trials using the concepts of the cumulative logit models. The binary outcomes are extended to three ordinal outcomes - none, mild

and severe denoted by 1, 2 and 3 respectively. A patient experiencing a severe toxicity would be considered a DLT and would not receive any further doses, while patients having a mild or no toxicity on the previous cycle would be considered eligible for further dose assignments. If a patient experiences a severe toxicity on any cycle they are typically taken off the study and they would not provide further data for the assessment of toxicity. The data for each patient would either consist of a series of ones (for example 11111) for a patient with no toxicity or a series of ones with some mild toxicities, (for example 1122) and some terminating with a severe toxicity (for example 11223). Typically in practice if a mild toxicity is observed the dose is lowered or kept the same on the next cycle. As an alternative to the 2 state Markov model we develop a three state Markov model with the states being 1, 2 and 3. Because 3 is a terminating state, we only need to consider the transition probabilities out of state 1 and 2. We explicitly model conditional probabilities of toxicity in a cycle given that the patient is either toxicity-free or has mild toxicities to date. The proposed model allows for a cumulative effect of dose on toxicity after the first cycle and allows covariates to influence the toxicity profile over cycles. A parameter is included to reflect an individual's tendency to respond consistently with past dose experience. Another parameter is included to account for the ordinal responses. In Section 4.2, we formulate the model, describe the Bayesian estimation method and provide intuition on model behavior. We then study finite sample operating characteristics and adaptive design through simulations in Sections 4.3 and 4.4 respectively and end

with a discussion in Section 4.5.

4.2 Methodology

4.2.1 Notation and data structure

We assume that there are five increasing dose levels of an experimental study drug represented by $d_g, g = 1, \ldots 5$, that will be studied in $i = 1 \ldots N$ patients. Each patient i completes $K_i \leq 6$ cycles, where K_i may be less than six if a patient experiences a DLT. On each cycle $k = 1, \ldots, K_i$, patient i receives a dose $d_{i,k}$ equal to one of the five values of d_g. The toxicity response for patient i on cycle k is $Y_{i,k}$ with $Y_{i,k} = 1$ indicating none toxicity and $Y_{i,k} = 2$ or $Y_{i,k} = 3$ indicating a mild or severe toxicity respectively. Patients stop receiving further administrations of the the drug if they experience a severe toxicity which is usually a DLT. Thus the possible patterns of $Y_{i,k}$ values for a patient are a sequence of ones or a sequence of ones and twos followed by a three. The observed data is $\{(Y_{i,1}, \ldots, Y_{i,K_i}, d_{i,1}, \ldots, d_{i,K_i}), i = 1, \ldots, N\}$. A patient's cumulative dose prior to cycle $k = 1, \ldots, K_i$ is $D_{i,k} = \sum_{j=1}^{k} d_{i,j-1}$, so that $D_{i,1} = 0$.

Borrowing on the idea from Chapter 2, the history of tolerating the previous doses is captured through the dose term $d_{i,k}^{\ddagger}$ which is designed to represent the maximum dose level a patient has tolerated with no toxicity on the past cycles. If a patient i had a mild toxicity on any dose level then $d_{i,k}^{\ddagger} = \phi \times \min\ (d_{i,j} \mid Y_{i,j} = 2, j = 1 \ldots k-1)$

and if the patient had no toxicities in the past then $d^{\ddagger}_{i,k} = \max(d_{i,j} \times I^{[Y_{i,j}=1]}, j = 1 \ldots k-1)$, where $I^{[Y_{i,j}=1]}$ denotes the indicator function that takes value 1 if $Y_{i,j} = 1$ and 0 otherwise. The intuition behind this is that the maximum dose level tolerated is set to be slightly lower by a factor ϕ, than the minimum dose at which the mild toxicity occurred. Note that $d^{\ddagger}_{i,k} = 0$ for cycle $k = 1$.

4.2.2 Proposed Markov model

For a patient i from a set of N patients, receiving $k(1 \ldots K)$ repeated cycles until a severe toxicity (grade 3 or higher) occurs, we define $p_{i,k(3)}$ to be the probability of severe outcome, $Y_{i,k} = 3$, at dose $d_{i,k}$ on cycle k and $p_{i,k(2+)}$ be the cumulative probability of severe and mild outcomes, $Y_{i,k} = 2$ or 3. Using the concepts of the cumulative logit models for modeling ordinal outcomes we define the following probabilities,

$$P(Y_{i,k} \geq 3) = P(Y_{i,k} = 3) = p_{i,k(3)}$$

$$P(Y_{i,k} \geq 2) = P(Y_{i,k} = 2, 3) = p_{i,k(2+)}$$

$$P(Y_{i,k} = 1) = p_{i,k(1)} = 1 - p_{i,k(2+)}$$

$$P(Y_{i,k} = 2) = p_{i,k(2)} = p_{i,k(2+)} - p_{i,k(3)}$$

Where $p_{i,k(1)}$, $p_{i,k(2)}$ and $p_{i,k(3)}$ are the probabilities of observing one of the three outcomes which follow a multinomial distribution. The ordinal Markov Model 4.1

for the ordinal outcomes is as follows,

$$1 - p_{i,k(3)} = \exp\left\{-\alpha\left(d_{i,k} - \rho d_{i,k}^{\ddagger}\right)^{+} - \beta D_{i,k} d_{i,k}\right\} \tag{4.1a}$$

$$1 - p_{i,k(2+)} = \exp\left\{(\theta+1)\left(-\alpha\left(d_{i,k} - \rho d_{i,k}^{\ddagger}\right)^{+} - \beta D_{i,k} d_{i,k}\right)\right\} \tag{4.1b}$$

where α, β, ρ and θ are required to be non-negative. Model intuition can be obtained by beginning with the first cycle, $k = 1$, when the probability of the severe toxicity is $p_{i,1(3)} = 1 - \exp(-\alpha d_{i,1})$ and $p_{i,1(2+)} = 1 - \exp(-(\theta+1)\alpha d_{i,1})$ is the probability of having a mild or severe toxicity. The only parameters that are relevant are α and θ on the first cycle. The probability of severe toxicity on the first cycle is captured by $\alpha \geq 0$. Because $\theta \geq 0$, the probability of severe toxicity is less than or equal to the probability of severe or mild toxicity. The parameter θ tries to capture the probability of observing a mild toxicity and can be thought of as a scaling parameter in the probability estimation of severe toxicity, which will be determined by the prevalence of mild toxicities.

On subsequent cycles, $k > 1$, the probability of toxicity is modified by the effect of the current dose and the cumulative dose captured by the two linear terms. In the first term $(d_{i,k} - \rho d_{i,k}^{\ddagger})$ consider first the role of ϕ which is set to a constant factor - ranging from $0 - 1$ and comes into play through $d_{i,k}^{\ddagger}$. A value of $\phi = 0$ implies that the patient cannot tolerate the dose level at which the mild toxicity occurred on a previous cycle while $\phi = 1$ implies that the mild toxicity can be ignored and assumed

(or override) that the patient could handle the full dose. The parameter ϕ in some sense captures the dependency of the patient responses on past doses to tolerability of the dose level in the presence of mild toxicity.

The term $(d_{i,k} - \rho d^{\ddagger}_{i,k})^+$ is not allowed to be negative to ensure that $p_{i,k} \in [0, 1]$. The parameter ρ can be thought of as reflecting the amount of memory about whether a dose was tolerable, with $\rho = 1$ reflecting perfect memory and $\rho = 0$ reflecting no memory. Thus this term tries to capture the within-patient correlation between dose cycles. If $\rho = 1$, $(d_{i,k} - \rho d^{\ddagger}_{i,k})^+$ reduces to $(d_{i,k} - d^{\ddagger}_{i,k})^+$ as the difference between the current assigned dose and the dose tolerated from previous cycles. If the current dose is less or equal to dose tolerated from previous cycles, the difference will be zero and will not contribute towards the probability estimate i.e., there is a strong memory that a higher or equal dose to the current one was tolerable hence the current dose is more likely to be tolerable. When $\rho = 0$, the term $(d_{i,k} - \rho d^{\ddagger}_{i,k})^+$ reduces to $d_{i,k}$ and thus there is no memory of the previous doses that had been tolerated. Intermediate values of ρ between zero and one have intermediate amount of memory.

The second linear term $\beta d_{i,k} D_{i,k}$, $\beta \geq 0$ is designed to capture the idea that there may be "damage" accumulated from prior doses and the amount of this "damage" plays a role in determining the probability of toxicity when a new dose is administered. The impact of the accumulated damage will be larger if $d_{i,k}$ is larger and will not be relevant if $d_{i,k} = 0$. In the case of a dose de-escalation and if the $(d_{i,k} - \rho d^{\ddagger}_{i,k})^+$ term is equal to zero then the cumulative effect is the only term that would account

for the probability of toxicity and will not be the driving force when the contribution from the current dose is too little or none at all.

Figure 4.1 plots the conditional probability of mild and severe toxicity in left and right columns respectively for different values of θ along the rows for fixed values of $\alpha = 1, \beta = 0, \phi = 0.8$ and $\rho = 0.8$ and aids in understanding the working properties of the ordinal Markov Model 4.1. The solid circles with dashed lines in the plots on the left panels represent the probability of mild toxicity on the first cycle at each of the five possible doses and is provided as a reference for comparison. While in the right panels the solid triangles represent the equivalent probability of severe toxicity on the first cycle at each of the five possible doses. The probabilities of toxicity on the second cycle are calculated at each of the five dose levels assuming that the patient received the third dose level on the first cycle. In both the mild/severe panels the open circles and crosses correspond to the probability of mild/severe toxicity on second cycle assuming mild and none toxicity respectively on the first cycle. In the first row when $\theta = 0.1$ there is not much difference between the probability of a mild toxicity on both cycle 1 and cycle 2. In contrast the probability of severe toxicity on the second cycle is lower than that on the first cycle. The probability of severe toxicity on second cycle given a mild toxicity on the first cycle is higher than the probability of severe toxicity given no toxicity on the first cycle as indicated by the open circles and crosses. Varying the values of θ does not have any effect on the probability of severe toxicity as evidenced by the identical left panel plots. For

increasing values of θ the probability of mild toxicity increases with the probability of observing a mild toxicity on the second cycle given that a mild toxicity occurred on the first cycle being higher than that when a none toxicity was observed on the first cycle. The non-zero value of $\rho = 0.8$ confers patients in cycle 2 to be less likely to experience a toxicity as a function of dose and hence the probability of toxicity on the second cycle is always lower than that on the first cycle.

The effect of using different values of ϕ is explored through Figure 4.2 with fixed values of $\alpha = 1, \beta = 0, \rho = 0.8$ and $\theta = 1$ with the layout of the panels and symbols in the plots holding the same definitions as Figure 4.1 and differing only by the varying values of ϕ along the rows. When $\phi = 0$ the probability of having a mild toxicity on the second cycle given that a mild toxicity was observed on the first cycle overlaps with the probability of observing a mild toxicity on the first cycle. Since $\phi = 0$ the patient is not given any credit on the second cycle for surviving the first cycle with a mild toxicity but surviving the first cycle with none toxicity confers a lower chance of mild toxicity on the second cycle as seen by the crosses. For increasing values of ϕ the probability of mild toxicity on the second cycle begins to differ from the first cycle and eventually for $\phi = 1$ the probability of a mild toxicity is the same irrespective of whether a mild or none toxicity was observed on the first cycle. The trends are similar in the case of the panels on the left for the probability of severe toxicity. There will be a positive contribution to the probability of both mild and severe toxicity on the second cycle when a non zero value of β is used in both these

plots. It is also easy to see that probability of toxicity on the first cycle will be higher/lower for increasing/decreasing values of α.

We can now compare the ordinal Markov model 4.1 to the model presented by Doussau [Doussau et al., 2013] using the mixed effects proportional odds model for three response outcomes as,

$$logit(P(Y_{i,j} \leq k|d_l)) \quad = \quad \alpha_k - \beta_1 d_l - u_i, u_i \sim N(0, \sigma_0^2)$$

Doussau's model has four parameters, $\alpha_1, \alpha_2, \beta_1$ and σ_0, which are equal to the number of parameters estimated using the Markov model. The correlation between patient responses is captured by the random effects term u_i. The authors also impose the condition that the dose levels remain constant within the patient. In contrast the ordinal Markov Model 4.1 allows patients the possibility of dose escalation or de-escalation especially when a mild toxicity is observed during the course of their treatment. Doussau's model formulation allows for a non-zero probability of toxicity in the absence of any dose, in general when dose (d_l) is zero the probability of toxicity should be zero, which is ensured in the ordinal Markov Model 4.1. Doussau's model does not capture the effect of the cumulative dose through an implicit parameter as is done in the Markov model through β. The probability of toxicity is assumed to be only due to the current dose although there exists a possibility for cumulative effect of toxicity. The random effect term u_i might be inadequate in capturing the additional

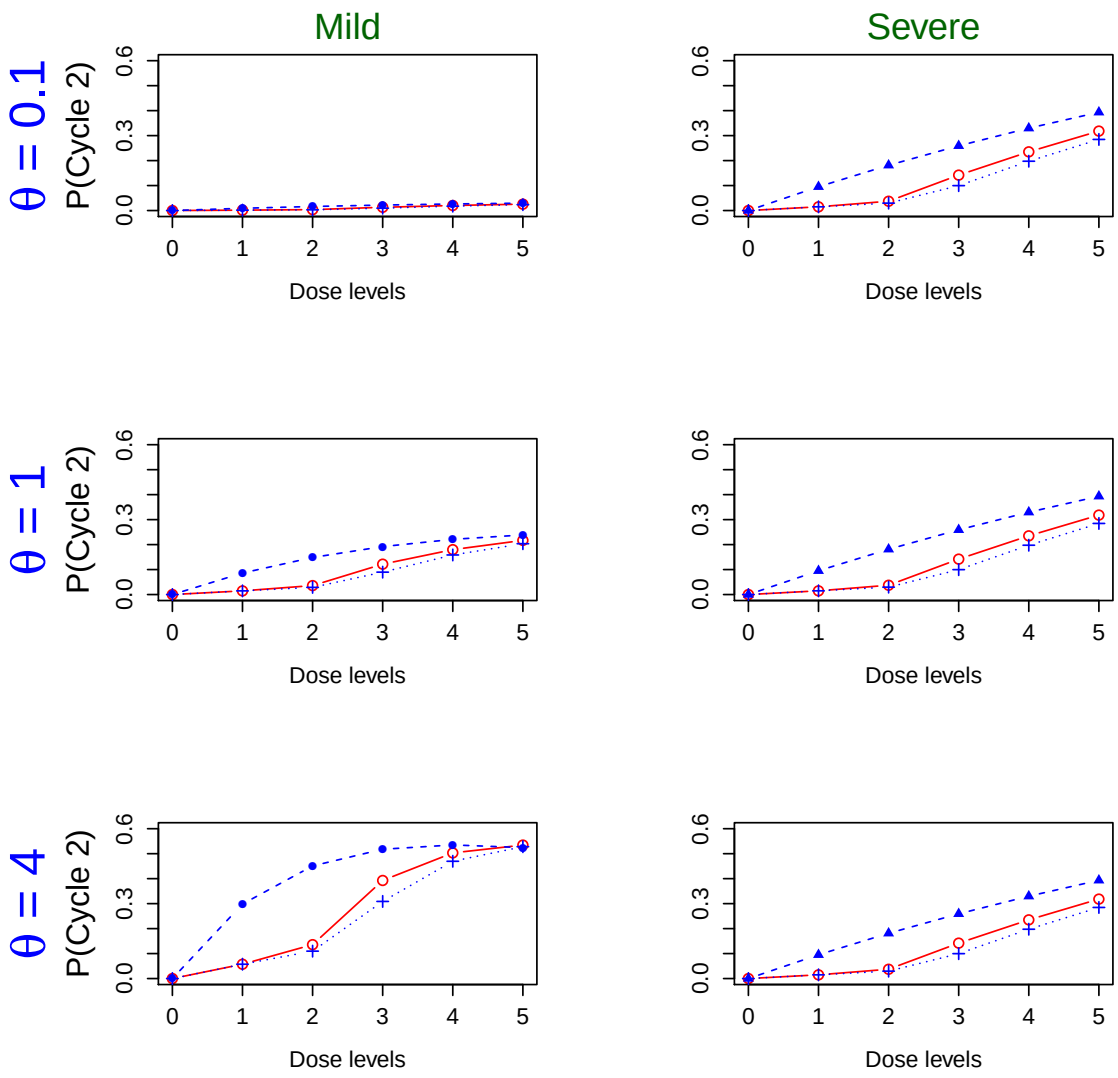

Figure 4.1: Conditional probability of mild and severe on the second cycle in the left and right columns respectively for different dose levels on the x-axis and assuming dose level three was given on cycle 1. The probability of mild and severe on the first cycle is in solid circles and triangles respectively for reference. The open circles correspond to assuming a mild toxicity while the crosses to none toxicity on cycle 1. Probabilities on cycle 2 are arranged by increasing values of θ from top to bottom with fixed values of $\alpha = 1, \beta = 0, \phi = 0.8$ and $\rho = 0.8$

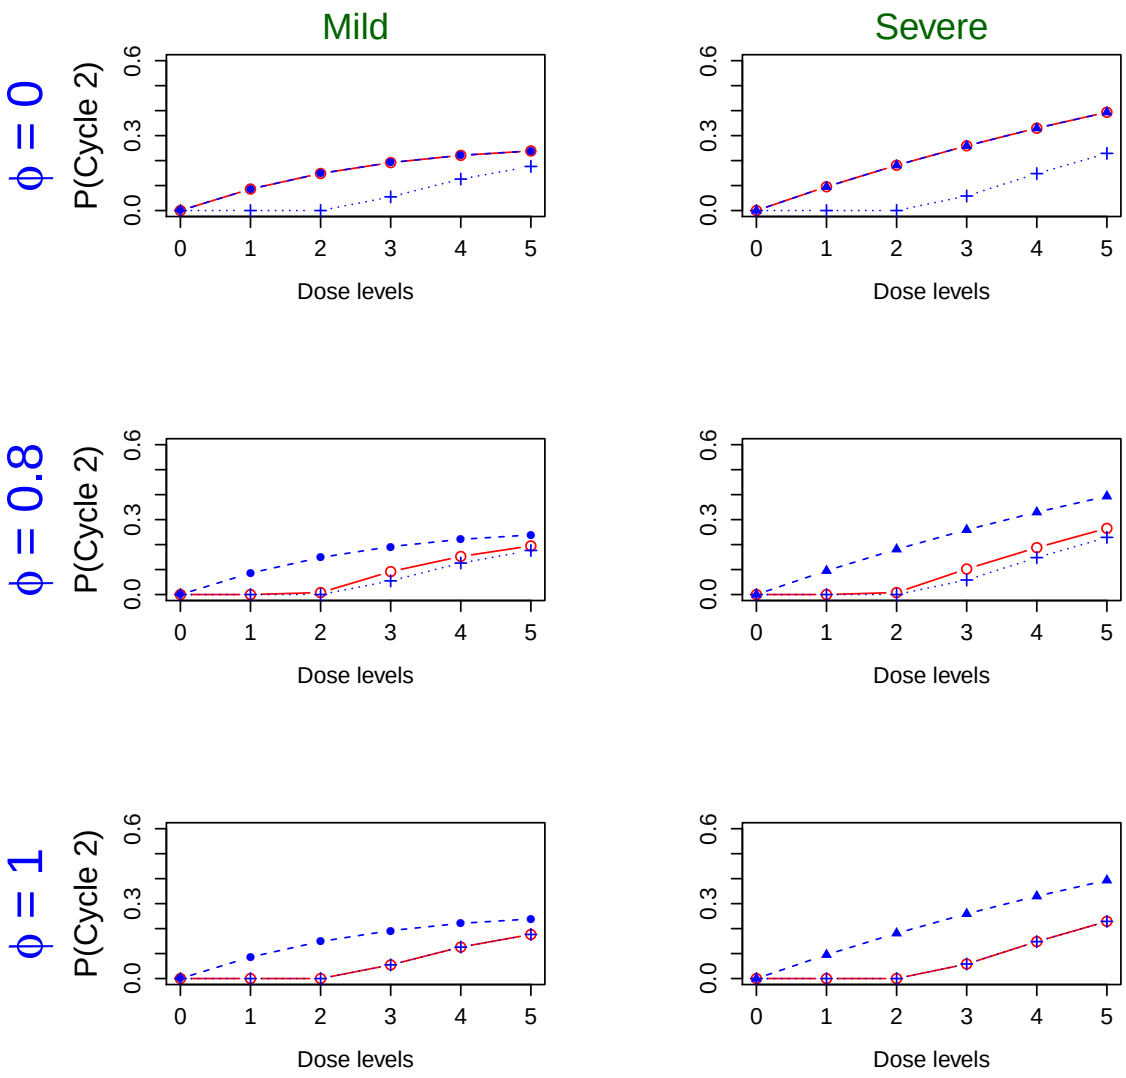

Figure 4.2: Conditional probability of mild and severe on the second cycle in the left and right columns respectively for different dose levels on the x-axis and assuming dose level three was given on cycle 1. The probability of mild and severe on the first cycle is in solid circles and triangles respectively. The open circles correspond to assuming a mild toxicity while the crosses to no toxicity on cycle 1. Probabilities on cycle 2 are arranged by increasing values of ϕ from top to bottom with fixed values of $\alpha = 1, \beta = 0, \theta = 1$ and $\rho = 0.8$

effect of the cumulative dose. Maximum likelihood estimates were obtained using Laplace approximations and adaptive Gauss-Hermite quadrature through a package in R. In contrast we set priors on the parameters and use Bayesian MCMC methods for parameter estimation through R.

4.2.3 Probability Skeleton

The dose levels to be studied are transformed to d_g via pre-specified skeleton probabilities denoted by $q_g, g = 1 \ldots 5$. The skeleton probabilities incorporate prior knowledge of the dose-toxicity relationship and correspond to the probability of observing a severe toxicity on the first cycle for each of the dose levels. On the first cycle the probability of severe toxicity is parameterized only by α since, $1 - p_{i,1(3)} = exp(-\alpha d_{i,1})$. Using the probability skeleton values q_g, the corresponding values for d_g are calculated assuming $\alpha = 1$ and solving $d_g = -log(1 - q_g)$. These transformed values of d_g are used as doses in the model formulation. Such use of probability skeleton is seen in the works of other authors like [O'Quigley et al., 1990, Lee and Cheung, 2009] in the context of the CRM and [Lee et al., 2011, Cheung and Elkind, 2010] in other contexts.

4.2.4 Prior and posterior distribution

Based on the study design, patients contribute to the likelihood until they experience a DLT or the final Kth cycle is completed. In general, subject i contributes,

$$L_{i,k}(Y_{i,k}|\alpha,\beta,\rho,\theta) = \prod_{k=1}^{K_i}(p_{i,k(1)})^{I^{[Y_{i,k}=1]}}(p_{i,k(2)})^{I^{[Y_{i,k}=2]}}(p_{i,k(3)})^{I^{[Y_{i,k}=3]}}$$

to the likelihood, where $I^{[Y_{i,k}=j]}$ denotes the indicator function that takes value 1 if $Y_{i,k} = j$ and 0 otherwise. The resulting likelihood for the entire study population is given by,

$$L(Y|\alpha,\beta,\rho,\theta) = \prod_{i=1}^{N}\prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha,\beta,\rho,\theta).$$

Our goal lies in estimating the posterior distributions of $p_{i,k(1)}, p_{i,k(2)}, p_{i,k(3)}$, $k = 1,\ldots,K$ in terms of the posterior distributions of parameters α,β,ρ and θ. Prior distributions on these parameters should reflect any auxiliary knowledge of the toxicity profile for the drug/agents being used in the trial, with a large prior variance when this knowledge is limited. In setting the prior on $\alpha \geq 0$ a lognormal (μ,σ^2) is used as a suitable prior having the form:

$$\pi(\alpha|\mu,\sigma) = \frac{1}{\sqrt{2\pi\sigma^2}}\frac{\exp(-(\log\alpha-\mu)^2/2\sigma^2)}{\alpha}$$

The prior mean for α is set at 1 to align with the use of the probability skele-

ton information. Setting the variance of the prior to 4, provides a coefficient of variance (CV) of 2. Parameters μ and σ are estimated using the expressions for the mean and variance of the lognormal density, $E(\alpha|\mu,\sigma) = exp(\mu + \sigma^2/2)$ and $Var(\alpha|\mu,\sigma) = exp\{2(\mu+\sigma^2)\} - exp(2\mu+\sigma^2)$.

As mentioned earlier in Section 4.2.2, $\rho \in [0,1]$, captures the correlation within patients receiving multiple doses, with values near zero indicating that the toxicity outcome is not influenced by previously administered doses and a value near one indicating a lower chance of toxicity from a previously administered dose. Since there is limited information to estimate ρ, we use a prior distribution with a small variance to allow some uncertainty in ρ, rather then choosing a fixed value. A $Beta(a,b)$ prior is used on ρ having density of the form:

$$\pi(\rho|a,b) = (\rho)^{a-1}(1-\rho)^{1-b}$$

The hyperparameters are set to $a = 5$ and $b = 1$ and using the expressions for the mean $a/(a+b)$ and variance $ab/\{(a+b)^2(a+b+1)\}$ the prior on ρ has a mean of 0.833 and variance of 0.02

The probability of toxicity is assumed to increase with an increase in the cumu-

lative dose and hence the lognormal density is used as the prior on $\beta \geq 0$. The prior mean is set to 0.5 and variance is set 1 obtaining a CV of 2.

In the case of θ a lognormal prior with mean 4 and variance 16 is used to span the positive real axis conforming to its bounds of $0 \leq \theta \leq \infty$. Simulation studies proved that it was difficult to estimate all the five parameters and hence ϕ has been set to a constant. Simulation results explore the effects of using ϕ=0.8 or 0.9.

The posterior distribution for α, β, ρ and θ given the observed data Y is then $f(\alpha, \beta, \rho, \theta|Y) =$

$$\frac{\prod_{i=1}^{N}\prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha,\beta,\rho,\theta)\pi_\alpha(\alpha)\pi_\beta(\beta)\pi_\rho(\rho)\pi_\theta(\theta)}{\int_0^\infty\int_0^1\int_0^\infty\int_0^\infty\prod_{i=1}^{N}\prod_{k=1}^{K_i} L_{i,k}(Y_{i,k}|\alpha,\beta,\rho,\theta)\pi_\alpha(\alpha)\pi_\beta(\beta)\pi_\rho(\rho)\pi_\theta(\theta)d\alpha d\beta d\rho d\theta}.$$

The posterior distribution of α, β, ρ and θ can be estimated via Markov Chain Monte Carlo (MCMC) methods [Robert and Casella, 1999] using just another Gibbs sampler `(JAGS)` *rjags* [Plummer, 2011] package through [`R` Development Core Team, 2011]. `JAGS` includes several algorithms for sampling from the posterior distributions produced from the MCMC iterations, for instance the standard Gibbs sampler is available for this purpose. Details of setting up the MCMC simulations are given in Section 4.2.5.

4.2.5 Implementation in JAGS

The data likelihood and the density definitions of the priors are specified in a model file saved under a `.bug` extension. The model file and the data are passed into the `JAGS` for compilation along with the list of parameters, α, β, ρ and θ, that have to be monitored. The number of parallel chains to be run by `JAGS` are also defined at the compilation stage, where each parallel chain produces independent samples from the posterior distribution. The compiled model then needs to be initialized for all the parameters that need to be monitored in each of the chains. The `JAGS` code is provided in Appendix 4.6.1. Samplers are automatically assigned by `JAGS` at the initialization stage after a pre-specified adaptive phase for each of the parameters based on the likelihood definition of the model. A relatively large burn-in period of 1000K samples with posterior samples of 500K (thinned by 5) are used in simulations presented in later sections. A slightly longer burn-in period is used when there are fewer patients in the sequential trial design. Before using the samples from the two chains for reporting they are monitored and assessed through diagnostic tests. The correlation between samples generated at each iteration of the MCMC chain for each of the parameters needs to be sufficiently low. The posterior means, $\hat{\alpha}, \hat{\beta}, \hat{\rho}$ and $\hat{\theta}$, of the four parameters, are used to calculate the various probabilities of interest.

4.3 Operating Characteristics/Results

This section presents simulation results studying the working properties of the Ordinal Markov model in two different settings, 1) a static setting, demonstrating the parameter estimation and 2) an adaptive setting with patients recruited sequentially.

4.3.1 Parameter Estimation

A 100 datasets were generated each having $N = 30$ patients. Patients were distributed equally over the five dose levels receiving the same dose over six cycles. The probability skeleton used for the five dose levels was (0.02,0.05,0.10,0.16,0.23), implying that the lowest and highest doses expected a 2% and 23% chance of severe toxicity on the first cycle respectively. The priors used on the parameters are as discussed earlier in Section 4.2.4, α, β and θ with lognormal prior distributions while ρ with a Beta prior distribution. Ordinal responses were generated for the patients under two different scenarios. The true values of the parameters were $\alpha = 1, \beta = 0.5, \phi = 0.9, \rho = 0.8, \theta = 4$ in scenario 1 and $\alpha = 1, \beta = 0.8, \phi = 0.9, \rho = 0.8, \theta = 3$ in scenario 2.

Table 4.1 shows the results of the simulations from 100 datasets. The parameters in the rows are grouped by the two scenarios. The columns provide 1) true value 2) mean of the prior with SD in parenthesis 3) mean of the estimates from 100 datasets with the bias from the true value in parenthesis 4) mean SD (MSD) from the 100 datasets 5) empirical SD (ESD) of the parameter from 100 datasets and 6) coverage

rate of the true value in the 95% credible intervals. Based on the results in Table 4.1 there seems to be an acceptable level of bias in the parameter estimates. The estimates of θ are close to the true values used in generating the data. The values of the MSD and ESD are comparable (slightly higher) indicating sufficient variability, except for ρ where we used a tight prior. In comparison to the prior SD of the first three parameters the data provided information and hence the MSD was reduced considerably compared to the prior SD. But for ρ there was a small decrease from the prior because the data was minimally informative about this parameter. The coverage rates are between 89% and 100%. These results are from simulations where the patients receive the same dose on all the cycles. Similar simulation results could be obtained when patients are allowed to escalate or de-escalate. We conclude that the model is able to provide accurate estimates of the parameters.

4.4 Adaptive Trial Design and Simulation

This section describes application of the ordinal Markov Model 4.1 in designing a sequential clinical trial to be used in practice assuming that patients would be assigned to one of the five dose levels. On completion of the first cycle without any severe toxicity the patient would be eligible to either stay at the same dose level or escalate to a higher dose level or de-escalate to a lower dose level based on the recommendations of the algorithm using the ordinal Markov Model 4.1 and the specified safety criterion. It will be assumed that a new patient is ready to

be assigned a dose once other continuing patients in the trial have completed their dosing cycle. That is, all active patients in the study and the new patient to be enrolled all receive their dose assignment simultaneously. Maximizing the expected total dose is the optimizing strategy used in choosing the best dose to be assigned to a patient on the next cycle when multiple choices are presented. In the next section we begin by defining the safety criteria for dose assignment and the notation used for the probabilities in defining the criterion, followed by defining the particulars of the maximizing strategy to be used in dose assignment. Finally a simulation example of a trial and results evaluating trial conduct properties comparing various criterion through simulations will be presented in this section.

4.4.1 Safety Criteria

We begin by defining the safety criteria rules for dose assignment in carrying out a trial with dose escalation and/or de-escalation. Define $r_{g,k} = g, g = 1 \dots 5$ as one of the five dose levels on cycle k corresponding to $d_g, g = 1 \dots 5$ the transformed doses using the probability skeleton. The following commonly used dose escalation rules will be followed in defining the safety criteria to be used while carrying out an adaptive clinical trial based on ordinal Markov Model 4.1.

- The first and the second patient on the trial will be assigned the second lowest dose level, $r_{g,1} = 2$, on cycle 1. Given that there is no severe toxic response, the same dose level will be assigned on the second cycle for the first patient.

For subsequent patients and cycles the following rules will be used.

- A patient will only be allowed single dose jumps in dose escalations, i.e., a patient i completing cycle k on dose level $r_{g,k}$ could have a maximum dose level $min(r_{g,k}+1, 5)$ on cycle $k+1$.

- A patient can have a maximum of three dose levels in a dosing regimen, unless de-escalation to a lower dose is required i.e.,a patient i receiving dose level $r_{g,1}$ on cycle 1 can possibly receive $r_{g,1}+2$, as its highest dose level in the dosing regimen. In combination with the previous rule a patient i completing cycle k on dose level $r_{g,k}$, can have $r^{max}_{g,k+1} = min(r_{g,1}+2, r_{g,k}+1, 5)$ as its highest possible dose level on any cycle $k+1$ in the study.

- For a new patient, $\hat{i}$, on the first cycle, to ensure a considerable degree of safety especially during the early stages of the trial, the maximum dose level choice would be limited to $r^{max}_{g,1} = max(r^{\ddagger}_{g,1}+1, r^{\ddagger}_{g,k})$, where $r^{\ddagger}_{g,1}$ is the maximum of all the past dose levels assigned to the patients on cycle $k=1$ and $r^{\ddagger}_{g,k}$ is the maximum of all the past dose levels assigned to the patients in the study on cycles $k>1$. This ensures that new patient $\hat{i}$ on the first cycle will not jump a dose level that has not been assigned previously to any patient in the study (there could be a possibility of a dose jump on the first cycle).

- The study will conclude when none of the dose levels are included in the tolerable range as determined by the safety criteria defined below.

4.4.1.1 Notation for probabilities used in safety criteria

Monitoring the safety of the patients is ensured by assigning doses that satisfy a set of safety criteria and the following notation is useful to understand these rules. For a patient i who has completed cycle k on dose level $r_{g,k}$ without a DLT, the possible dose levels for this patient are $j = 1 \dots r^{max}_{g,k+1}$, allowing for dose de-escalation.

Define $A_{i,k+1,j}$ as the event of severe toxicity on cycle $k+1$ for patient i at dose level j given that there were no DLTs in the past. Hence $\mathrm{P}(\mathrm{A}_{i,k+1,j}) = p_{i,k+1(3)}$, where $p_{i,k+1(3)}$ is calculated using the Markov Model 4.1 for dose level j.

Given $r^{max}_{g,k+1}$ for patient i the potential regimens considered are denoted by the set $R^{regimen}_{i,k+1}$. Members, $r^{regimen}_{i,k+1,m}$, of this regimen set have length $K = 6$, corresponding to the number of cycles, where the first k elements are the doses received and future doses constrained by $min(r_{g,1} + 2, r_{g,k} + 1, 5)$ for cycles $(k+1) \dots K$.

Define $B_{i,k+1,m}$ as the event of having a severe toxicity on any cycle $k+1$ until K for a member m of the regimen set $R^{regimen}_{i,k+1}$, where $\mathrm{P}(\mathrm{B}_{i,k+1,m}) = 1 - \prod_{l=k+1}^{K}(1 - p_{i,m(3)})$. For regimen m from patient i's regimen set define $C_{i,k+1,m}$ as the event of severe toxicity for that regimen on any cycle from 1 through K i.e., $C_{i,k+1,m} = B_{i,1,m}$. $P(B_{i,k+1,m})$ captures the probability of severe toxicity for a patient on the remaining cycles while $\mathrm{P}(\mathrm{C}_{i,k+1,m})$ captures the overall probability of severe toxicity on the entire regimen m.

4.4.1.2 Target probability bounds

In conducting a clinical trial with multiple doses we define the following probability bounds $\mathring{P}(A)$, $\mathring{P}(B)$ and $\mathring{P}(C)$ based on clinical inputs. $\mathring{P}(A)$ is a vector of $K = 6$ probabilities corresponding to the acceptable conditional probability of severe toxicity on the next cycle. $\mathring{P}(A_1)$ is the acceptable probability of severe toxicity limit on first cycle, while $\mathring{P}(A_2) \ldots \mathring{P}(A_k) \ldots \mathring{P}(A_6)$ are the corresponding limits on subsequent cycles. Given the properties of the current Markov model used to specify the dose toxicity relationship wherein the conditional probabilities post cycle 1 are much lower and affected mainly by the cumulative effect, the bounds on the conditional probabilities are assumed to be equal to each other i.e., $\mathring{P}(A_2) = \mathring{P}(A_3) = \ldots = \mathring{P}(A_6)$, and will be referred to as $\mathring{P}(A_2)$. For a continuing patient, $\mathring{P}(B)$ is the acceptable probability of severe toxicity limit on all the remaining cycles while $\mathring{P}(C)$ is the acceptable probability of severe toxicity limit on the entire regimen. The relationship between these target probabilities is explored as follows.

Typically in single dose trials the acceptable level of toxicities is set at 30% which in the case of multiple dose trial would correspond to $\mathring{P}(C)$, the probability of severe toxicity bound on all the cycles. For the bounds to be consistent with each other, $1 - \mathring{P}(C) \leq \prod_{k=1}^{6}\{1 - \mathring{P}(A_k)\}$ which further reduces to $\{1 - \mathring{P}(C)\} \leq \{1 - \mathring{P}(A_1)\} \times \{1 - \mathring{P}(A_2)\}^5$. Assuming this relationship, provides an easy way to specify the bounds on $\mathring{P}(A_2)$, given the acceptable bounds on $\mathring{P}(A_1)$ and $\mathring{P}(C)$. In

the simulations presented slightly higher bounds are used to demonstrate differences in results. $\mathring{P}(B)$ controls the bounds on the probability of toxicity on the remainder of the cycles and its value is chosen in relation to $\mathring{P}(C)$ and would typically be $\leq \mathring{P}(C)$. We have set the value of $\mathring{P}(B) = \mathring{P}(C)$ in the simulations presented in this chapter.

4.4.2 Maximizing the Expected dose

With the safety criteria rules in place an optimizing strategy needs to be defined upfront in the event that multiple dose level qualify on the next cycle. As mentioned earlier the goal of the study is to maximize the total dose assigned to every patient and we proceed by deriving the expression for the expected dose followed by the algorithm for dose maximization.

4.4.2.1 Expected dose

In the presence graded toxicity outcomes the expected dose is calculated conditional on the past responses. A patient having a severe toxicity is considered to be a DLT and a terminating state hence the past responses of mild and none toxicities need to be accounted for in calculation of the expected dose. The following equations indicate the pattern in estimating the probabilities needed for the calculation of the

expected dose.

P(severe on cycle 1)

$$P(Y_{i,1} = 3) = p_{i,1(3)}$$

P(severe on cycle 2)

$$\begin{aligned} P(Y_{i,2} = 3) &= P(Y_{i,2} = 3, Y_{i,1} = 1) + P(Y_{i,2} = 3, Y_{i,1} = 2) \\ &= P(Y_{i,2} = 3|Y_{i,1} = 1)P(Y_{i,1} = 1) + \\ &\quad P(Y_{i,2} = 3|Y_{i,1} = 2)P(Y_{i,1} = 2) \\ &= p_{i,1(1)}p_{i,2(3|1)} + p_{i,1(2)}p_{i,2(3|2)} \end{aligned}$$

P(severe on cycle 3)

$$\begin{aligned} P(Y_{i,3} = 3) &= p_{i,1(1)}p_{i,2(1|1)}p_{i,3(3|11)} + p_{i,1(1)}p_{i,2(2|1)}p_{i,3(3|12)} + \\ &\quad p_{i,1(2)}p_{i,2(1|2)}p_{i,3(3|21)} + p_{i,1(2)}p_{i,2(2|2)}p_{i,3(3|22)} \end{aligned}$$

Similarly the P(mild on cycle 3) is given by,

$$\begin{aligned} P(Y_{i,3} = 2) &= p_{i,1(1)}p_{i,2(1|1)}p_{i,3(2|11)} + p_{i,1(1)}p_{i,2(2|1)}p_{i,3(2|12)} + \\ &\quad p_{i,1(2)}p_{i,2(1|2)}p_{i,3(2|21)} + p_{i,1(2)}p_{i,2(2|2)}p_{i,3(2|22)} \end{aligned}$$

And the P(none on cycle 3) is given by,

$$\begin{aligned} P(Y_{i,3}=1) &= p_{i,1(1)}p_{i,2(1|1)}p_{i,3(1|11)} + p_{i,1(1)}p_{i,2(2|1)}p_{i,3(1|12)} + \\ & \quad p_{i,1(2)}p_{i,2(1|2)}p_{i,3(1|21)} + p_{i,1(2)}p_{i,2(2|2)}p_{i,3(1|22)} \end{aligned}$$

The probability of observing a severe toxicity on cycle 3 for a given dose combination is a functional sum of the four combinations of past responses; none toxicities on the first two cycles, none toxicity on the first cycle and a mild toxicity on the second cycle, a mild toxicity on the first cycle and none toxicity on the second cycle and lastly both mild toxicities on the first two cycles. The exact probability of severe toxicity on every cycle k is calculated by forming a binary tree branch that tracks the past responses with $2^{(k-1)}$ combinations of mild and none toxicity. Using the expressions for exact probabilities of severe toxicity defined above and extending them mild and none toxicity, the expected dose for a patient i having a particular planned dose sequence $d_{i,1} \ldots d_{i,k}$ is

$$= \sum_{m=1}^{K} \left\{ \left(\sum_{j=1}^{m} d_{i,j} \right) P(Y_{i,m}=3) \right\} + \left(\sum_{j=1}^{K} d_{i,j} \right) P(Y_{i,K}=2) + \left(\sum_{j=1}^{K} d_{i,j} \right) P(Y_{i,K}=1).$$

4.4.2.2 Recommending a regimen

At the conclusion of the study, the estimates $\hat{\alpha}, \hat{\beta}, \hat{\rho}$ and $\hat{\theta}$ are used to estimate the overall probability of toxicity $\mathrm{P}(\mathrm{C}_j)$ for all of the $j = 1 \ldots 19$ regimens listed in

Table 2.3. During the conduct of the trial the target probability bounds used were $\mathring{P}(A)$, $\mathring{P}(B)$ and $\mathring{P}(C)$. These bounds especially, $\mathring{P}(C)$, are usually set at higher than acceptable values in practice and when selecting the final regimen we would use $P^r(A_1)$ and $P^r(C)$ which might be lower than or equal to $\mathring{P}(A_1)$ and $\mathring{P}(C)$ respectively. For example in running the trial $\mathring{P}(C) = 0.40$ could be used which implies an overall toxicity of 40% but in practice 30% toxicities are what we would want to see in the trials. Results will be presented when the final selection of the regimen is based on $P^r(A_1)$ and $P^r(C)$. The recommended regimen satisfies $P(A_1) \leq P^r(A_1)$ on the first cycle and $\mathrm{P(C}_j) \leq P^r(C)$ and has the highest possible maximum expected dose. The target regimen (T) is identified using the true parameter values of α, β, ρ and θ instead of the estimates obtained at the trial conclusion and is used as a reference for gauging the properties of the trial.

4.4.2.3 Algorithm maximizing the expected dose

The algorithmic plan for maximizing the expected dose on the study for each patient i is outlined below.

- For new patient $\hat{i}$ on cycle 1

    1. Estimate $\hat{P}(A_{\hat{i},1,j})$ using current estimates $\hat{\alpha}, \hat{\beta}, \hat{\rho}$ and $\hat{\theta}$ at each of the $j = 1 \ldots r_{g,1}^{max}$ dose levels. Where $r_{g,1}^{max} = max(r_{g,1}^{\ddagger} + 1, r_{g,k}^{\ddagger})$, $r_{g,1}^{\ddagger}$ is the maximum of all the past dose levels assigned to the patients on cycle $k = 1$

and $r^{\ddagger}_{g,k}$ is the maximum of all the past dose levels assigned to the patients in the study on cycles $k > 1$.

2. Subset the dose levels that satisfy $\hat{P}(A_{\hat{i},1,j}) \leq \mathring{P}(A_1)$ over all dose levels.
3. For the dose levels satisfying $\hat{P}(A_{\hat{i},1,j}) \leq \mathring{P}(A_1)$ subset the list of possible regimens from Table 2.3 and calculate the overall probability of severe toxicity $\hat{P}(C_{\hat{i},1,j})$ tracking all combinations of mild and none toxicities.
4. Select the dose level that has an overall probability of toxicity $\hat{P}(C_{\hat{i},1,j}) \leq \mathring{P}(C)$ and maximum expected dose.
5. If none of the doses satisfy $\hat{P}(A_{\hat{i},1,j}) \leq \mathring{P}(A_1)$ and if there are continuing patients in the study then wait until updated estimates $\hat{\alpha}, \hat{\beta}, \hat{\rho}$ and $\hat{\theta}$ allow doses to be assigned else the study is terminated.

- For continuing patient i on cycle $k > 1$,

1. List doses $\hat{P}(A_{i,k+1,j}) \leq \mathring{P}(A_2)$ from $r^{max}_{g,k+1} = min(r_{g,1} + 2, r_{g,k} + 1, 5)$ possible choices.
2. If there is more than one satisfying dose level then list the possible dose regimen set $R^{regimen}_{i,k+1}$.
3. Calculate the probability of severe toxicity $\hat{P}(B_{\hat{i},k+1,m})$ assuming all combinations of mild and none toxicities on the remainder of the cycles for each of the regimens m in $R^{regimen}_{i,k+1}$ and the corresponding expected dose using the current estimates $\hat{\alpha}, \hat{\beta}, \hat{\rho}$ and $\hat{\theta}$.

4. Select the dose level that has probability of toxicity $\hat{P}(B_{\hat{i},k+1,m}) \leq \mathring{P}(B)$ and maximizes the expected dose.

5. If none of the doses satisfy $\hat{P}(A_{\hat{i},k+1,j}) \leq \mathring{P}(A_2)$ and if there are continuing patients in the study then wait until updated estimates $\hat{\alpha}, \hat{\beta}, \hat{\rho}$ and $\hat{\theta}$ allow doses to be assigned else the study is terminated.

For the purposes of evaluating the properties of the simulation of clinical trials over multiple replications and comparing the properties of the target probabilities various test statistics will be calculated that can be grouped into 1) trial conduct or patient characteristics and 2) recommended regimen characteristics as explained below.

Patient characteristics

1. Mean dose per patient over all the replicates. In each of the replicates the total dose given to all the patients will be tracked and then averaged across the number of the patients in that trial. High values of mean dose are considered favorable indicating that the patients received higher quantities of the drug in the study.

2. Mean number of severe toxicities per study across all the replicates. The number of patients having a severe toxicity are averaged across the number of patients in the trial. Low values of severe toxicities are considered favorable.

3. The trials that stop early without recruiting all $N = 30$ patients. Lower number

of trials stopping are considered favorable because in the situations considered there is a regimen that is not too toxic.

4. Mean number of patients having a regimen that matches the recommended regimen.The distance from the recommended regimen is calculated for each of the patients in the study, and the proportion of patients having distance less than or equal to two are summarized and the average proportion across all the replications is presented. High values of patients matching or very similar to the recommended regimen are considered favorable implying that the regimen recommended at study conclusion was actually observed in patients in the trial.

Recommended regimen characteristics

1. Mean of expected dose using the true values of α, β, ρ and θ given the recommended regimen. High values of mean expected dose are considered favorable implying that the recommended regimen if completed on all the six cycles would provide the highest and safest amount of the study drug to the patients.

2. Mean of probability of severe toxicity on any cycle using the true values of α, β, ρ and θ given the recommended regimen. Low values of mean toxicities are considered favorable since the recommended regimen should provide low levels of severe toxicities.

3. The fraction of recommended regimens that have toxicity less than say 40% calculated using the true values is also presented. Higher proportions are fa-

vorable and imply that the recommended regimen does not have a high toxicity level.

4. A green region is defined as the list of five target regimens with highest expected dose that satisfy the safety constraints. The proportion of times the recommended regimen belongs to this green region is also presented. High values of the proportion values are considered favorable and indicating a higher degree of concordance between the recommended and target regimens.

4.4.3 Simulation results

Simulation results demonstrating the algorithm in sequential clinical trials are presented in this section. Additional details of the set-up are that each trial enrolls a maximum of $N = 30$ patients with each patient having a maximum of six cycles. A total of 100 replicates are conducted with the true values of $\alpha = 1, \beta = 0.5, \rho = 0.8$ and $\theta = 4.0$ used in generating the patient responses. Dose assignment is based on the algorithm outlined in Section 4.4.2.3 with a focus on maximizing the dose received by each patient on the trial. The skeleton probability for cycle 1 used for the doses was $0.02, 0.05, 0.10, 0.16, 0.23$. The priors used on the parameters are as outlined in the Section 4.2.4, specifically α, β and θ with lognormal prior distributions while ρ with a Beta prior distribution. Data from a 100 replicates/trials were simulated to study the design properties under different settings.

1. Effect of using two different values of $\phi = 0.8, 0.9$, implying that a patient could

handle about (80%, 90%) of the drug effect if a mild toxicity occurred at that dose level.

2. Effect of using different values of $\mathring{P}(A)$, $\mathring{P}(B)$ and $\mathring{P}(C)$ during trial conduct.

Table 4.2 presents simulation result summaries of trials carried out under different settings. Each row corresponds to summaries from 100 different replicates and are grouped firstly by $\phi = 0.8$ or $\phi = 0.9$ followed by $\mathring{P}(B) = 0.3$ and $\mathring{P}(C) = 0.4$ or $\mathring{P}(B) = \mathring{P}(C) = 0.4$ with three different combinations of $\mathring{P}(A_1)$ and $\mathring{P}(A_2)$. The columns correspond to the four patient and three regimen characteristics as mentioned in Section 4.4.2.3. While Table 4.2 has $\mathring{P}(A_1) = \mathring{P}(A_2)$, Table 4.3 presents results with $\mathring{P}(A_1) \neq \mathring{P}(A_2)$ and all other settings remaining the same.

The results from Table 4.2 indicate a very slight increase in the mean dose in trials with $\phi = 0.9$ as compared to $\phi = 0.8$. This aligns with the model intuition that lower values of ϕ imply the patient is able to handle a lower proportion of the dose when experiencing a mild toxicity. The mean dose also increases with a higher threshold of $\mathring{P}(B) = 0.4$ as compared to $\mathring{P}(B) = 0.30$ as expected since patients can be assigned a higher dose level with a lenient threshold of $\mathring{P}(B) = 0.4$. The mean dose increases with higher thresholds of $\mathring{P}(A_1)$ and $\mathring{P}(A_2)$ since higher dose levels qualify for dose assignment. Within a category of $\mathring{P}(B)$ and $\mathring{P}(C)$ the mean toxicity is the highest when $\mathring{P}(A_1)$ and $\mathring{P}(A_2)$ are at the highest threshold of 0.2 conveying the trade off between being lenient in dose assignment and patient safety. The number of trials

stopping early due to an inability to assign doses is highest when $\mathring{P}(A_1)$ and $\mathring{P}(A_2)$ are the most stringent at 0.05 essentially allowing only 5% severe toxicities at each cycle. The mean on the recommended dose is calculated using the true values given the recommended regimen. When $\mathring{P}(A_1) = 0.05$ the mean dose is lower than mean recommended dose implying that the patients received a lower dose during the study conduct. There is a close match between the mean dose and the recommended mean dose when $\mathring{P}(A_1) = 0.10$ and when $\mathring{P}(A_1) = 0.20$ the mean dose is higher than that mean of the recommended dose. The mean toxicity of the recommended regimen given the true values is always below 30% which is a very good feature. The fraction of recommended regimens that have toxicity less than say 40% is always $\geq$ 95%. Also the recommended regimens are mostly above 72% in the green region.

The results from Table 4.3 are generated with $\mathring{P}(A_1) \neq \mathring{P}(A_2)$ and better understanding of the results can be obtained by comparing the results to the corresponding row in Table 4.2. For example consider the first row corresponding to $\phi = 0.8$ and $\mathring{P}(C) = 0.4, \mathring{P}(B) = 0.30, \mathring{P}(A_1) = 0.05$ and the difference is only due to $\mathring{P}(A_2) = 0.05$ or $\mathring{P}(A_2) = 0.10$. As expected the higher value of $\mathring{P}(A_2)$ increases the mean dose assigned from 10.96 to 13.91 along with an increase in the mean toxicity from 0.17 to 0.29. The mean dose is now slightly higher than the mean of the recommended expected dose and the proportion of patients having the recommended regimen during the course of the trial also increases. In contrast notice the case when $\mathring{P}(A_1) = 0.20$ and $\mathring{P}(A_2) = 0.20$ or $\mathring{P}(A_2) = 0.05$ the mean dose decreases from 17.45

to 13.20 and the mean toxicities from 0.40 to 0.20 with none of the trials stopping early in both cases. In general Table 4.3 results indicate that having $\mathring{P}(A_2)$ lower than $\mathring{P}(A_1)$ provides lower values of mean toxicity and closer agreement between the mean dose and the recommended mean dose. Also the recommended regimens are mostly above 73% in the green region, the instance when the value drops to 69% could be ascribed to $\phi = 0.90$ and the stringent bounds during trial conduct and regimen selection at the end of the trial.

4.4.3.1 Comparison the dichotomous Markov model

Simulation results will now be presented showing the benefit in using an ordinal outcome Markov model in comparison to the dichotomous Markov Model 2.1. For the dichotomous outcome the algorithm set up for running the sequential trial is similar to that used earlier except that the ordinal outcome model is considered to be the true model when assigning the responses to the patients. When assigning dose levels to the patients, the mild response is assumed to be a none toxicity. All other details of the model set up are similar to that used in Section 4.4.3.

Tables 4.4 and 4.5 have a layout similar to Tables 4.2 and 4.3 respectively and allow an easy comparison. First consider the results from Table 4.4. There does not seem to be any effect of ϕ on the mean toxicities while there is a slight increase in the mean dose when $\phi = 0.9$ as compared to $\phi = 0.8$. The mean dose differs from the mean recommended dose in most cases and is comparable when $\mathring{P}(A_1) = 0.10$.

The proportion of recommended doses in the green are mostly around 50%. There are trials stopping early in every scenario the lowest value is when $\mathring{P}(A_1) = 0.20$. The results in Table 4.5 when $\mathring{P}(A_1) \neq \mathring{P}(A_2)$ can be explained by comparing the results from Table 4.4. For example the mean dose increases from 11.62 to 14.13 when $\mathring{P}(A_2) = 0.10$ as compared to $\mathring{P}(A_2) = 0.05$ with $\phi = 0.8$ and $\mathring{P}(A_1) = 0.05$. Differences in other results can be explained similarly.

In comparing the results across two modes of trial conduct with the truth as ordinal model and using either ordinal model or the dichotomous Markov model for the trial conduct there are some noticeable differences. Comparing the results from Table 4.4 to Table 4.2 and Table 4.5 to Table 4.3 we notice that the proportion of trials stopping early is higher for the dichotomous outcome. The proportion of patients having a distance ≤ 2 from the recommended regimen is lower and the proportion of recommended regimen being in the green are also very low. The mean dose and mean toxicity reverse equality based on the values of $\mathring{P}(C), \mathring{P}(B), \mathring{P}(A_1)$ and $\mathring{P}(A_2)$. The mean toxicities are usually higher but low only when $\mathring{P}(A_1) = \mathring{P}(A_2) = 0.20$. There seems to be a considerable improvement in the performance of the trial when the information provided by the mild toxicities is taken into account.

4.5 Discussion

One of the main goals of a phase I clinical trial is to arrive at an accurate estimate of the MTD without having too many patients experience a severe toxicity.

Assigning multiple doses to patients ensures that patients are more likely to receive an efficacious dosage before having a toxicity. Keeping track of the graded toxicity outcomes of patients prevents patients from receiving dose levels that could result in more severe toxicities on future cycles and thereby making sure that patients remain in the study for longer and receive dose levels closer to their range of tolerability. The ordinal outcome Markov model presented in this chapter is a novel method of incorporating both the repeated measure information from all the patients and the ordinal toxicity information of the responses on dose levels in the past. The advantage of using the data from all the patients allows making correct dose assignment decisions in the future for the patients.

The benefits of using this model are demonstrated through comparisons with the dichotomized Markov model. There were improved gains in trial properties by incorporating the ordinal outcomes. The number of overall toxicities were lower and the proportion of trials recommending the true dose regimen were also considerably higher. We know of one other method that incorporates individual grades or toxicity scores [Doussau et al., 2013] and having some form of comparison between the two methods could be a great possibility for future work.

The other advantage to using the ordinal Markov model in carrying out the trial is that the dose levels are adjusted based on the past responses of the patients especially if a mild toxicity had occurred. Such an adaptive design ensures that patients with frail dispositions are not exposed to higher doses levels. In this regard finding

an accurate estimate of ϕ instead of setting to a constant would be an advantage. Investigating the estimation of ϕ merits further research in future work.

The results from using the trial in an adaptive trial setting could be more dramatic if we used the probability of mild toxicity in dose assignment decisions. By doing do we could have assigned doses to continuing patients by incorporating their chance of having a mild toxicity on a particular dose level instead of only looking at the probability of a severe toxicity. This would add to the burden of defining another set of target probabilities on the lines of $\mathring{P}(C), \mathring{P}(B)$ and $\mathring{P}(A)$ pertaining to the chance of observing a mild toxicity. Although this has not been done in this chapter the algorithm in its present form could be easily extended to do this.

The Markov model presented in this chapter is most relevant to clinical trials involving cytotoxic drugs where the toxicity is assumed to increase with the cumulative effect. Having non-delayed outcomes is also essential to the study design so that the DLT could be assigned at the end of the cycle to the appropriate dose level for the patient.

There exist various other possibilities for further developments of the method. In this chapter we showed improvements over the dichotomized Markov model. The ordinal Markov model could be further extended to include the un-grouped five classes of toxicity ranging from zero (none) to five (death) as defined by National Cancer Institute [NCI, 2003]. The model currently allows treating patients at the dose level that would not result in any severe toxicity thereby focusing only on the

safety of the patients. More recently much interest is generated in estimating a safe and efficacious dose level [Zhang et al., 2006], such approaches could be extended to the ordinal Markov model with repeated measures from the patients.

Table 4.1: Parameter estimates obtained through simulation of a 500 datasets with $N = 30$ patients under two different scenarios of true parameter values with $\phi = 0.90$.

	True Value	Prior Mean (SD)	Estimate(Bias)	MSD[1]	ESD[2]	Coverage
Scenario 1						
α	1	1 (2)	1.114 (0.114)	0.408	0.375	99
β	0.5	0.5 (1)	0.553 (0.053)	0.325	0.315	96
θ	4	4 (4)	4.387 (0.387)	1.752	1.666	97
ρ	0.8	0.83 (0.14)	0.802 (0.002)	0.111	0.076	99
Scenario 2						
α	1	1 (2)	1.044 (0.044)	0.392	0.405	93
β	0.8	0.5 (1)	0.813 (0.013)	0.405	0.423	95
θ	3	4 (4)	3.440 (0.440)	1.337	1.385	89
ρ	0.8	0.83 (0.14)	0.796 (-0.004)	0.122	0.075	100

[1] MSD is mean of the SD from 500 estimates

[2] ESD is empirical SD of the 500 estimates

Table 4.2: Trial/Patient summary using $\phi = 0.8, 0.90$ while conducting 100 sequential trials with $N = 30$ patients and $\mathring{P}(A_1) = \mathring{P}(A_2)$, $\mathring{P}(B) = 0.3, 0.4$ and $\mathring{P}(C) = 0.4$. True model used for generating responses is Model 4.1 and model used in running the trial is Model 4.1. Columns correspond to the mean dose received per patient, the mean number of severe toxicities over 100 trials, number of trials stopping early, mean number of patients having distance less than two from the recommended regimen, the mean expected dose for the recommended regimen, mean probability of toxicity for the recommended regimen, the proportion of trials having recommended regimen with an overall toxicity under 40% and the proportion of trials having recommended regimen in the green region as defined in Section 4.4.2.3.

	Mean Dose	Mean Tox	Early Stop Trials	P.R.Dist ≤ 2	Mean R R.EDose	Mean R R.Tox	Prop $\leq 40\%$	Prop in Green
$\phi = 0.80$								
$\mathbf{\mathring{P}(C) = 0.4\ \mathring{P}(B) = 0.3}$								
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.05$	10.96	0.17	5	6.31	12.86 (95)	0.23	0.95	0.78
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.1$	14.86	0.28	2	13.53	14.39 (98)	0.27	0.95	0.77
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.2$	17.45	0.40	0	9.36	15.37 (100)	0.29	0.99	0.80
$\mathbf{\mathring{P}(C) = 0.4\ \mathring{P}(B) = 0.4}$								
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.05$	10.96	0.17	5	6.38	12.91 (95)	0.23	0.95	0.77
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.1$	15.14	0.30	2	13.41	14.48 (98)	0.27	0.95	0.79
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.2$	17.64	0.43	0	8.74	15.21 (100)	0.28	0.99	0.86
$\phi = 0.90$								
$\mathbf{\mathring{P}(C) = 0.4\ \mathring{P}(B) = 0.3}$								
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.05$	10.97	0.17	6	6.5	12.75 (94)	0.22	0.94	0.77
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.1$	15.17	0.29	2	14.45	14.82 (98)	0.27	0.95	0.75
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.2$	17.55	0.40	0	9.42	15.47 (100)	0.28	0.97	0.74
$\mathbf{\mathring{P}(C) = 0.4\ \mathring{P}(B) = 0.4}$								
$\mathring{P}(A_1) = 0.05, \mathring{P}(A_2) = 0.05$	10.97	0.17	6	6.1	12.89 (94)	0.23	0.94	0.73
$\mathring{P}(A_1) = 0.1, \mathring{P}(A_2) = 0.1$	15.31	0.29	2	13.78	14.81 (98)	0.27	0.95	0.72
$\mathring{P}(A_1) = 0.2, \mathring{P}(A_2) = 0.2$	17.88	0.42	0	9.02	15.59 (100)	0.28	0.99	0.74

Table 4.3: Trial/Patient summary using $\phi = 0.8, 0.90$ while conducting 100 sequential trials with $N = 30$ patients and $\mathring{P}(A_1) \neq \mathring{P}(A_2)$, $\mathring{P}(B) = 0.3, 0.4$ and $\mathring{P}(C) = 0.4$. True model used for generating responses is Model 4.1 and model used in running the trial is Model 4.1.

	Mean Dose	Mean Tox	Early Stop Trials	P.R.Dist ≤ 2	Mean R R.EDose	Mean R R.Tox	Prop $\leq 40\%$	Prop in Green
$\phi = 0.80$								
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.3}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.1$	13.91	0.29	5	10.32	12.84 (95)	0.23	0.95	0.79
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.15$	15.97	0.36	2	12.04	14.49 (98)	0.26	0.97	0.86
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.05$	13.20	0.20	0	9.49	14.87 (100)	0.27	1	0.81
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.4}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.1$	13.92	0.29	5	9.69	12.86 (95)	0.23	0.95	0.78
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.15$	16.25	0.38	2	10.18	14.53 (98)	0.27	0.97	0.80
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.05$	13.16	0.20	0	9.79	14.87 (100)	0.27	1	0.86
$\phi = 0.90$								
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.3}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.1$	14.09	0.28	5	9.64	12.97 (95)	0.23	0.95	0.74
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.15$	16.20	0.35	2	11.85	14.49 (98)	0.25	0.98	0.82
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.05$	13.32	0.20	0	9.93	15.27 (100)	0.27	0.99	0.73
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.4}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.1$	14.20	0.29	5	8.20	13.10 (95)	0.24	0.95	0.69
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.15$	16.51	0.37	2	10.36	14.79 (98)	0.27	0.95	0.80
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.05$	13.26	0.21	0	10.17	15.32 (100)	0.27	0.99	0.77

Table 4.4: Trial/Patient summary using $\phi = 0.8, 0.90$ while conducting 100 sequential trials with $N = 30$ patients and $\mathring{P}(A_1) = \mathring{P}(A_2)$, $\mathring{P}(B) = 0.3, 0.4$ and $\mathring{P}(C) = 0.4$. True model used for generating responses is Model 4.1 and model used in running the trial is Model 2.1.

	Mean Dose	Mean Tox	Early Stop Trials	P.R.Dist ≤ 2	Mean R R.EDose	Mean R R.Tox	Prop $\leq 40\%$	Prop in Green
$\phi = 0.80$								
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.3}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.05$	11.62	0.22	8	4.59	13.66 (92)	0.23	0.91	0.58
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.1$	14.91	0.34	5	12.92	14.80 (95)	0.26	0.94	0.45
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.2$	15.55	0.34	2	12.51	15.42 (98)	0.27	0.96	0.55
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.4}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.05$	11.71	0.22	9	4.54	13.83 (91)	0.24	0.90	0.54
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.1$	15.30	0.38	5	13.98	14.89 (95)	0.25	0.94	0.63
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.2$	16.88	0.44	2	7.92	15.34 (98)	0.26	0.96	0.58
$\phi = 0.90$								
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.3}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.05$	11.81	0.22	8	4.50	13.98 (92)	0.25	0.90	0.50
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.1$	15.03	0.33	5	12.69	14.91 (95)	0.26	0.94	0.43
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.2$	15.63	0.34	2	13.12	15.50 (98)	0.27	0.96	0.57
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.4}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.05$	11.82	0.22	8	4.32	13.93 (92)	0.24	0.91	0.52
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.1$	15.73	0.38	5	12.80	15.21 (95)	0.27	0.91	0.53
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.2$	17.11	0.44	2	8.78	15.58 (98)	0.27	0.97	0.53

Table 4.5: Trial/Patient summary using $\phi = 0.8, 0.90$ while conducting 100 sequential trials with $N = 30$ patients and $\mathring{P}(A_1) \neq \mathring{P}(A_2)$, $\mathring{P}(B) = 0.3, 0.4$ and $\mathring{P}(C) = 0.4$. True model used for generating responses is Model 4.1 and model used in running the trial is Model 2.1.

	Mean Dose	Mean Tox	Early Stop Trials	P.R.Dist ≤ 2	Mean R R.EDose	Mean R R.Tox	Prop $\leq 40\%$	Prop in Green
$\phi = 0.80$								
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.3}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.1$	14.13	0.34	11	10.91	13.97 (88)	0.24	0.88	0.52
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.15$	15.11	0.36	5	12.52	14.84 (95)	0.25	0.94	0.55
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.05$	14.13	0.26	2	10.09	15.34 (98)	0.27	0.92	0.58
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.4}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.1$	14.34	0.36	9	6.56	13.57 (91)	0.23	0.91	0.56
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.15$	16.11	0.44	5	7.99	14.76 (95)	0.25	0.95	0.58
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.05$	14.35	0.26	2	9.83	15.34 (98)	0.27	0.94	0.60
$\phi = 0.90$								
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.3}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.1$	14.21	0.34	11	10.27	13.81 (88)	0.24	0.88	0.53
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.15$	15.21	0.36	5	12.43	14.92 (95)	0.26	0.95	0.54
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.05$	14.23	0.26	2	10.07	15.46 (98)	0.27	0.90	0.55
$\breve{\boldsymbol{P}}(\boldsymbol{C}) = \mathbf{0.4}\ \breve{\boldsymbol{P}}(\boldsymbol{B}) = \mathbf{0.4}$								
$\breve{P}(A_1) = 0.05, \breve{P}(A_2) = 0.1$	14.46	0.35	9	6.58	13.69 (91)	0.24	0.90	0.55
$\breve{P}(A_1) = 0.1, \breve{P}(A_2) = 0.15$	16.31	0.43	5	8.23	14.84 (95)	0.25	0.93	0.57
$\breve{P}(A_1) = 0.2, \breve{P}(A_2) = 0.05$	14.51	0.26	2	10.04	15.51 (98)	0.27	0.95	0.56

4.6 Appendix

4.6.1 Outline of the code written in JAGS

The code presented below corresponds to applying ordinal Markov Model 4.1 in simulations for parameter estimation from the posterior samples.

```
#Defining the model.bug file
model {
#Define the likelihood for each of the N subjects
for (i in 1:N) {
```

```
prob[i,3]< −1-exp(-alpha* ( dose[i]-rho*phi.flag[i])*step(dose[i]-rho*phi.flag[i])
-beta*dose[i]*cumdose[i])
prob32[i] < − 1- exp((theta+1) *(-alpha* ( dose[i] - rho*phi.flag[i])*step(dose[i] -
rho*phi.flag[i]) - beta*dose[i]*cumdose[i] ))
p[i,2] < − p32[i] - p[i,3]
p[i,1] < −1 - p32[i]
pat.response[i] ∼ dcat (p[i,])
}
#Setting up the priors
#prior on α - E(α) = 1 and V ar(α) = 4
mu1 < − -0.3465736; tau1 < − 1.4426950
alpha ∼ dlnorm(mu1,tau1)
#prior on ρ
a1 < − 5 ; b1 < − 1
rho ∼ dbeta(a1,b1)
#Prior on β - E(β) = 0.5 and V ar(β) = 1
mu2 < − -1.498; tau2 < − 0.621
beta ∼ dlnorm(mu2,tau2) #Prior on θ - E(θ) = 4 and V ar(θ) = 16
mu3 < − 1.356; tau3 < − 16.495
theta ∼ dlnorm(mu3,tau3) }
#Initializing the parameters
inits< −list(list(alpha=1,beta=0.1,rho=0.7,theta=0.1),
list(alpha=0.5,beta=0.8,rho=0.9,theta=2) )
parameters < −c(“alpha”,“beta”, “rho”,“theta”)
#updating the simulations
data < −list(“response”=response,“phi.flag”=phiflag,
“cumdose”=cumdose,“dose”=dose,“N”=N)
jags < − jags.model(file=”prior.bug”, data = data, inits=inits, n.chains = 2, n.adapt
= 10000)
adapt(jags,n.iter=1000)
update(jags,1000000) # burin samples
sim1< −coda.samples(jags, parameters, 500000,thin=5)
#check for convergence
plot(sim1)
gelman.plot(sim1)
gelman.diag(sim1)
geweke.plot(sim1)
geweke.diag(sim1)
```

```
autocorr(sim1)
autocorr.plot(sim1)
#report the mean and quantiles of the posterior distributions
y3< −summary(sim1)
ystat< −data.frame(y3statistics)
yquant = data.frame(y3quantiles)
```

4.6.2 Example of an adaptive trial in progress

Based on the algorithm presented in Section 4.4.2.3 for maximizing the total expected dose received by a patient, an example of a trial in progress is presented in this section for demonstrating the dose assignment in practice. The target probability bounds used during the execution are $\mathring{P}(C) = 0.40, \mathring{P}(B) = 0.30, \mathring{P}(A_1) = 0.10, \mathring{P}(A_2) = 0.10$ and the true values of the parameters are $\alpha = 1, \beta = 0.5, \rho = 0.8, \theta = 4$. The fixed parameter $\phi = 0.80$ in these simulations.

Table 4.6 presents the current patient profile in the trial. The rows correspond to the unique patients added sequentially in the trial. The columns correspond to the six cycles with the dose level assigned to the patient and the response in parenthesis. A one signifies none toxicity while a two and three denote mild and severe toxicity respectively. A cross is placed in all cycles once a severe response is observed for a patient. There are 13 patients in the trial and decisions need to be made for dose assignment to patients 9, 11, 12 and 13 and a new patient 14. Before patient 13 was added to the trial the parameter estimates were $\hat{\alpha} = 1.082(0.436), \hat{\beta} = 0.335(0.277)$, $\hat{\rho} = 0.903(0.102)$ and $\hat{\theta} = 4.111(0.959)$ with the posterior standard deviations in

parenthesis. The updated current estimates of the parameters are $\hat{\alpha} = 1.163(0.441)$ $\hat{\beta} = 0.338(0.279), \hat{\rho} = 0.906(0.099)$ and $\hat{\theta} = 4.241(0.981)$. Notice that the estimate of $\hat{\theta}$ increases in response to the mild toxicity observed by patient 11 on cycle 3.

The probability of toxicity on dose levels 1 through 4 for patient 9 are 0.003,0.023, 0.089,0.168 of which dose level 3 has probability of toxicity $\leq \mathring{P}(A_2) = 0.10$ and is also able to provide a regimen combination that satisfies the $\mathring{P}(C) = 0.40, \mathring{P}(B) = 0.30$ and hence is assigned to patient 9 on cycle 6. Using the true values of the parameters and the current dose assignment the true probability of toxicity is calculated and a Bernoulli response is generated. In a similar fashion the remaining patients are assigned doses and responses and the updated patient profile is presented in Table 4.7. The updated parameter estimates are now $\hat{\alpha} = 1.079(0.407), \hat{\beta} = 0.309(0.252)$, $\hat{\rho} = 0.905(0.099)$ and $\hat{\theta} = 4.214(0.974)$. Notice the slight decrease in estimates $\hat{\alpha}, \hat{\beta}$ and $\hat{\theta}$ since no fresh toxicities are observed.

At the end of the trial the completed patient profile is presented in Table 4.8. The parameter estimates at the conclusion of the trial are $\hat{\alpha} = 0.958(0.291), \hat{\beta} = 0.478(0.302)$, $\hat{\rho} = 0.826(0.124)$ and $\hat{\theta} = 3.398(0.727)$ The probability of toxicity on the first cycle and on any cycle is calculated using the current estimates of the parameters for all the 19 regimens in Table 2.3 to select the recommended regimen and by using the true parameter values to select the target regimen. By setting $P^r(A_1) = \mathring{P}(A_1) = 0.10$ and $P^r(C) = 0.3$ and using the true parameter values the target regimen selected is 333333 while the recommended regimen is 33333 using the

parameter estimates obtained at the conclusion of the trial. If only $P^r(C) = 0.3$ is used the target regimen is 333333 while the recommended regimen is 333333.

Table 4.6: Table showing the dose level assignment and patient responses in parenthesis for an adaptive trial in progress using Model 4.1 with accrual of 13 patients.

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
1	2 (1)	2 (1)	3 (1)	4 (2)	3 (1)	3 (2)
2	2 (1)	3 (3)	X	X	X	X
3	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)
4	3 (2)	2 (1)	2 (1)	2 (1)	2 (1)	2 (1)
5	2 (2)	2 (1)	2 (1)	3 (2)	2 (1)	2 (1)
6	2 (1)	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)
7	2 (1)	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)
8	2 (2)	2 (1)	3 (1)	2 (1)	2 (1)	2 (1)
9	2 (1)	3 (2)	2 (1)	3 (1)	3 (1)	?
10	2 (1)	3 (3)	X	X	X	X
11	2 (1)	3 (1)	4(2)	?	?	?
12	2 (1)	3 (2)	?	?	?	?
13	2 (1)	?	?	?	?	?

X Terminated patient having a severe toxicity (3)
? Continuing patient

Table 4.7: Table showing the dose level assignment and patient responses in parenthesis for an adaptive trial in progress using Model 4.1 with accrual of 14 patients.

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
1	2 (1)	2 (1)	3 (1)	4 (2)	3 (1)	3 (2)
2	2 (1)	3 (3)	X	X	X	X
3	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)
4	3 (2)	2 (1)	2 (1)	2 (1)	2 (1)	2 (1)
5	2 (2)	2 (1)	2 (1)	3 (2)	2 (1)	2 (1)
6	2 (1)	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)
7	2 (1)	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)
8	2 (2)	2 (1)	3 (1)	2 (1)	2 (1)	2 (1)
9	2 (1)	3 (2)	2 (1)	3 (1)	3 (1)	3(1)
10	2 (1)	3 (3)	X	X	X	X
11	2 (1)	3 (1)	4(2)	3(1)	?	?
12	2 (1)	3 (2)	2(1)	?	?	?
13	2 (1)	3(1)	?	?	?	?
14	2 (1)	?	?	?	?	?

X- Terminated patient having a severe toxicity (3)

? Continuing patient

Table 4.8: Table showing the dose level assignment and patient responses in parenthesis for a completed adaptive trial using Model 4.1.

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
1	2 (1)	2 (1)	3 (1)	4 (2)	3 (1)	3 (2)
2	2 (1)	3 (3)	X	X	X	X
3	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)
4	3 (2)	2 (1)	2 (1)	2 (1)	2 (1)	2 (1)
5	2 (2)	2 (1)	2 (1)	3 (2)	2 (1)	2 (1)
6	2 (1)	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)
7	2 (1)	3 (1)	3 (1)	3 (1)	3 (1)	3 (1)
8	2 (2)	2 (1)	3 (1)	2 (1)	2 (1)	2 (1)
9	2 (1)	3 (2)	2 (1)	3 (1)	3 (1)	3(1)
10	2 (1)	3 (3)	X	X	X	X
11	2 (1)	3 (1)	4(2)	3(1)	3 (1)	3 (1)
12	2 (1)	3 (2)	2(1)	3 (1)	3 (1)	3 (3)
13	2 (1)	3(1)	4(1)	4(1)	4(1)	4(3)
14	2 (1)	3(1)	4(1)	4(1)	4(3)	X
15	2 (2)	2 (1)	2(2)	3(3)	X	X
16	2 (1)	3 (1)	4(3)	X	X	X
17	2 (1)	3 (1)	3(2)	3 (1)	3 (1)	3 (1)
18	2 (1)	3 (1)	3(1)	3 (3)	X	X
19	2 (1)	3 (1)	3(1)	3 (3)	X	X
20	2 (1)	3 (1)	3(1)	3 (1)	3 (1)	3 (2)
21	2 (1)	3 (1)	3(1)	3 (1)	3 (1)	3 (1)
22	2 (3)	X	X	X	X	X
23	2 (2)	1 (1)	2 (1)	2 (1)	2 (1)	2 (1)
24	2 (1)	3 (1)	3(1)	3 (1)	3 (1)	3 (1)
25	2 (1)	3 (1)	3(1)	3 (1)	3 (1)	3 (1)
26	2 (2)	2 (1)	2 (1)	2 (1)	2 (1)	2 (1)
27	2 (1)	3 (1)	3(1)	3 (1)	3 (1)	3 (1)
28	2 (1)	3 (1)	3(1)	3 (1)	3 (2)	4 (1)
29	2 (1)	3 (1)	3(2)	3 (1)	3 (1)	3 (1)
30	2 (1)	3 (1)	3(1)	3 (1)	3 (1)	3 (1)

X -Terminated patient having a severe toxicity (3)

CHAPTER 5

Discussion and future work

We have proposed novel models for the conditional probability of toxicity to specify the dose-toxicity relationship in clinical trials in oncology having repeated dose administrations. Using Bayesian methods the models can be fit to data that arises in the conduct of a trial that allows patients to have dose escalation or de-escalation. Allowing for intra-patient dose escalation and de-escalation gives the patient a greater chance to be treated at a therapeutic dose, an advantage over the current trials in oncology that restrict patients to have the same dose over all the cycles.

The first model in Chapter 2 had three parameters to account for the effect of the current dose, the cumulative dose and the effect of dependency between patient responses. Benefits in modeling the data from all cycles were demonstrated. In addition estimation of parameters by allowing patients to vary doses over the course of treatment was presented. The model application in conducting a sequential clinical

trial by assigning doses to patients based on the all the available information was also demonstrated. Chapter 3 demonstrated extensions of the model incorporating dose and gender covariates. An ingenious way to build priors for the three parameters was presented. Application of the model to the sarcoma dataset demonstrated its ability to include covariates in modeling the dose-toxicity relationship. Chapter 4 demonstrated the extension of the Markov model to include ordinal outcomes accounting for none, mild or severe toxicities. Sequential design of a clinical trial using the model was presented. Benefits of using the ordinal model in comparison to the dichotomized, two-state Markov model, were also demonstrated through simulations.

Overall, the methods proposed in this dissertation represent a meaningful contribution to the field of adaptive clinical trial design. Although many statistical methods have been proposed for adaptive clinical trials most are not used in practice [Dent and Eisenhauer, 1996]. We are hopeful that with current interest in the CRM in carrying out clinical trials there will be an eventual shift towards within-patient dose-escalation trials and the methods presented in this dissertation would provide the necessary tools and framework to carry them out. Our models provide a simple way to model a complex data structure parsimoniously.

In general for dose escalation studies the design should be influenced by the steepness of the dose-response curves, interpatient variability in pharmacokinetics and whether the toxicities are reversible [Chevret, 2006]. In our setting when intra-patient dose changes are allowed and repeated measurements of toxicity response are

used it is also necessary that the toxicity can be attributed to the dose received in the current cycle and that it occurs within the time frame of the cycle.

There are numerous pros and cons to using the model based methods presented in this dissertation for adaptive clinical trial designs. It is hard to ignore the complexity of the model design and the algorithm for dose assignment. The need to set up the probability skeleton for the dose levels to be studied, the priors on the different parameters of interest, the safety criteria rules for intra-patient dose escalation/de-escalation, the choice of plans for optimizing the dose given to the patients, the bounds for the various target probabilities and the bounds for the probabilities for the eventual selection of a dosing regimen are a wide array of factors to consider before carrying out the trial. The calculation of the probabilities that are used in deciding the dose for the next cycle for a patient are based on accumulated data that must be available in real time. In comparison the most widely used '3+3' algorithmic design has a simpler approach to arriving at the maximum tolerable dose level in a single dose setting. In actuality the model based approach presented in this dissertation incorporates the safety criteria rules used in the algorithmic '3+3' design but in contrast provides the additional benefit of treating patients close to the safer dose level by incorporating information from all the patients in the trial and additionally allowing patients to receive multiple doses. In situations where there are limited patients to be recruited our model based methods provide efficient estimates and allow the patients to be treated at the best dose level thereby making the extra

effort of setting up the study design worthwhile.

In the evaluation of the methods we considered regimens of six cycles and the only reason a patient would drop out prior to that was if they experienced a DLT. In practice patients may drop out for other reasons. The Bayesian estimation approach can still be used as long as at least a few patients have a long sequence of toxicity measurements, however there would be less precision of the parameter estimates with less accumulated data. In the dichotomous Markov model presented in Chapter 2, three parameters were estimated of which the parameter ρ captured the dependency of the within-patient responses. Estimating this parameter is challenging in the presence of limited data especially at the start of the trial. Others including [Whitehead et al., 2001, 2006] seem to have encountered similar problems when trying to estimate the dependency between patient responses and have resorted to setting it to a constant. We have circumvented this issue by using a tight Beta prior on ρ. These difficulties in estimating ρ were less profound when the sample size increased as was the case in Chapter 3 when the Markov model was used in a Phase II setting. This chapter presented other non-standard priors on ρ offering options to incorporate prior belief in dependency of patient responses. A tight Beta prior was also used on ρ in Chapter 4 since it was difficult to estimate all the four parameters in the ordinal Markov model.

Since the parameter estimation is done via MCMC methods a fairly high understanding of the use of MCMC techniques is crucial and might pose a limitation to

the use of the methods. In the simulation results presented in this dissertation a relatively large fixed number (2000K) of simulations have been used as the burn-in period for the initial 15 patients to ensure convergence of the posterior distributions and a lower burn-in period (500K) for the subsequent patients in the ordinal outcome model. The burn-in periods were determined based on the initial testing phase of the method. A more stringent monitoring of the burn-in and convergence of the posterior distributions is advocated when the method is being used in practice for dose assignment in actual clinical trials.

A better understanding of the operation of the method for different values of the target probabilities and safety criteria through simulations is advocated. The simulation results presented in Chapters 2 and 4 demonstrated the effect of the safety criteria set up, the target probability bounds and the optimization criteria for the dose assignment on the trial properties. Additional safety criteria rules can be incorporated, for example to prevent a new patient from escalating to a higher and new dose level $(r_g + 1)$ on the first cycle given that the previous patient had a DLT on dose level r_g, the new patient could be assigned a dose level of r_g or less. Clinicians often hesitate in de-escalating the dose level when a patient tolerates the dose on previous cycles. The algorithm could be modified to prevent continuing patients in the trial from escalating to a higher dose level thus overriding the dose recommendations made by the model algorithm. Studying the effect on the trial properties by incorporating such safety criteria rules is strongly advocated via simulation studies.

We have focused on using the models since they provide a good estimate of the expected dose for the recommended regimen. In the context of multiple dose administrations per patient once a recommended regimen is selected, the expected dose corresponding to this regimen can be calculated and the probability of surviving the entire regimen without a DLT can also be estimated. Given this value of the recommended expected dose a number of dose level combinations are possible that could match the expected dose and yet have an acceptable level of overall probability of toxicity. Hence having an estimate of the tolerable expected dose gives rise to the possibility of proposing various regimen combinations meeting the expected dose level and the overall toxicity rate on all the cycles and could be used to narrow down the possible choices of regimens for recommending to the next phase of testing.

Simulations are presented for five dose levels in Chapter 2 and 4 and for two dose levels in Chapter 3 but in practice the model could be easily extended to different number of dose level combinations. Also the number of cycles for the regimen are not limited to four or six cycles are presented in Chapter 3 and Chapters 2 and 4 respectively. Currently all simulations demonstrating the sequential operation of the models assumed that the patients complete their cycles simultaneously and that a new patient is ready for dose assignment at the same time. Dose assignment happens for the continuing patients and the new patient based on the parameter estimates available at that stage. This simplistic assumption reduced the computational time of the simulations and also minimized the complexity of the results during comparisons.

In practice patient arrival could be generated using an exponential distribution and the length of individual cycles could also be programmed when considering patient completion. This would entail estimating the parameters more often since patients would not be aligned to complete their cycles simultaneously. Simulations can be done to study the effects of perturbation on differences in accrual rates of new patients, varying cycle duration/length and possibly varying cycle duration/length per patient.

Another obvious extension to the models presented would be to include time to event outcomes. We currently use the information from patients who have completed their ongoing cycle. Using a time to event outcome on the lines of the TITE-CRM [Cheung and Chappell, 2000, Braun, 2006], we could incorporate the partial information from patients currently in a cycle by using weights for the period of time without a DLT. This could provide a more accurate estimate of probability of toxicity rates when deciding the dose level for a new patient or a continuing patient.

Thus far, we have considered incorporating only the safety information through the occurrence of a DLT. There is an increasing use of early clinical trials to demonstrate efficacy in addition to safety. Many authors including [Thall and Cook, 2004, Braun, 2002, Thall et al., 1999] have provided models to be used in the phase I/II setting to simultaneously arrive at a safe and efficacious dose level. In a similar vein extensions to the Markov models can be envisioned that include a bivariate trial design in which the MTD is based jointly on both the toxicity and disease progression

information.

Yet another possibly complex extension of the Markov models could be to study the dose-toxicity profiles of two study drugs simultaneously. [Thall et al., 2003] provided a two stage Bayesian method giving acceptable dose-pairs of two agents in the phase I cancer chemotherapy setting. Similar extensions could be designed for the multiple dose cycles per patient setting using the Markov models either in the binary or the ordinal outcome setting presented in this dissertation.

One of the most promising results from the research in this dissertation is that the use of ordinal responses can lead to improved selection of the recommended regimen at the end of the trial. This is worthy of further research.

BIBLIOGRAPHY

BIBLIOGRAPHY

A. Agresti. *Categorical Data Analysis*. Wiley-Interscience., 2nd ed. edition, 2002.

J. Babb, A. Rogatko, and S. Zacks. Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine*, 17:1103–1120, 1998.

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

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

T.M. Braun. Generalizing the TITE-CRM to adapt for early- and late-onset toxicities. *Statistics in Medicine*, 25(12):2071–83, 2006.

Y. K. Cheung and R. Chappell. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics*, 56(4):1177–82, 2000.

Y.K. Cheung and M.S.V. Elkind. Stochastic approximation with virtual observations for dose-finding on discrete levels. *Biometrika*, 97(1):109–121, 2010.

S. Chevret. *Statistical Methods for Dose-Finding Experiments.* Wiley, 2006.

R. Chugh, T. Wagner, K.A. Griffith, J.M.G. Taylor, D.G. Thomas, F.P. Worden, K.M. Leu, M.M. Zalupski, and L.H. Baker. Assessment of ifosfamide pharmacokinetics, toxicity, and relation to CYP3A4 activity as measured by the erythromycin breath test in patients with sarcoma. *Cancer*, 109(11):2315–22, 2007.

S.F. Dent and E.A. Eisenhauer. Phase I trial design: are new methodologies being put into practice? *Annals of Oncology*, 7(6):561–6, 1996.

A. Doussau, R. Thiébaut, and X. Paoletti. 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.

A.E. Gelfand and A.F.M. Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410):398–409, 1990.

A. Gelman and D.B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472, 1992.

W.R. Gilks, D. G. Clayton, D.J. Spiegelhalter, N. G. Best, and A.J. McNeil. Mod-

elling Complexity: Applications of Gibbs Sampling in Medicine. *Journal of the Royal Statistical Society. Series B (Methodological)*, 55(1):39–52, 1993.

H.W. Greene. *Econometric analysis (5th ed.)*. Prentice Hall, 2003.

A. Iasonos, S. Zohar, and O'Quigley J. Incorporating lower grade toxicity information into dose finding designs. *Clinical Trials*, 8(4):370–379, 2011.

A. Ivanova and S.H. Kim. Dose finding for continuous and ordinal outcomes with a monotone objective function: a unified approach. *Biometrics*, 65(1):307–315, 2009.

N.L. Johnson and S. Kotz. *Continuous univariate distributions-1*. John Wiley and Sons, 1970.

S.M. Lee and Y.K. Cheung. Model calibration in the continual reassessment method. *Clinical Trials*, 6(3):227–238, 2009.

S.M. Lee, B. Cheng, and Y.K. Cheung. Continual reassessment method with multiple toxicity constraints. *Biostatistics*, 12(2):386–98, 2011.

S.M. Lee, D. Hershman, P. Martin, K. Leonard, and Y.K. Cheung. Toxicity burden score: A novel approach for multiple toxicities. *Annals of Oncology*, 23(2):537–541, 2012.

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

National Cancer Institute NCI. Common terminology criteria for adverse events (CTCAE). Version 3.0. Bethesda, 2003. URL `http://ctep.cancer.gov`.

J. O'Quigley and S. Chevret. Methods for dose finding studies in cancer clinical trials: a review and results of a Monte Carlo study. *Statistics in Medicine*, 10(11): 1647–64, 1991.

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

M. Plummer. Discussion of the paper by Spiegelhalter et al. Journal of the Royal Statistical. *Journal of the Royal Statistical Society, Series B*, (64):620, 2002.

M. Plummer. Penalized loss functions for Bayesian model comparison. *Biostatistics*, 9(3):523–39, 2008.

M. Plummer. *rjags: Bayesian graphical models using MCMC*, 2011.

S. Postel-Vinay, C. Gomez-Roca, L.R. Molife, B. Anghan, A. Levy, I. Judson, J. De Bono, J-C. Soria, S. Kaye, and X. Paoletti. Phase I trials of molecularly targeted agents: should we pay more attention to late toxicities? *Journal of Clinical Oncology*, 29(13):1728–35, 2011.

K. Pye and A. Whitehead. A Bayesian approach to dose finiding studies for cancer therapies: incorporating later cycles of therapy [abstract]. page 33. 33rd Annual

Conference of the International Society for Clinical Biostatistics: Bergen Norway, 2012.

R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2011.

C.P. Robert and G. Casella. *Monte Carlo Statistical Methods*. New York: Springer, 1999.

R. Simon, B. Freidlin, L. Rubinstein, S.G. Arbuck, J. Collins, and M.C. Christian. Accelerated titration designs for phase I clinical trials in oncology. *Journal of the National Cancer Institute*, 89(15):1138–47, 1997.

D.J. Spiegelhalter, N.G. Best, B.P. Carlin, and A. van der Linde. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society, Series B*, (64):583–639, 2002.

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

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

P.F. Thall and S.J. Lee. Practical model-based dose-finding in phase I clinical trials: methods based on toxicity. *International Journal of Gynecological Cancer*, 13(3): 251–61, 2003.

P.F. Thall, E.H. Estey, and H.G. Sung. A new statistical method for dose-finding based on efficacy and toxicity in early phase clinical trials. *Invest New Drugs.*, 17 (2):155–67, 1999.

P.F. Thall, R.E. Millikan, P. Mueller, and S.J. Lee. Dose-finding with two agents in phase I oncology trials. *Biometrics*, 59(3):487–496, 2003.

J.R. van Dorp and S. Kotz. Generalized trapezoidal distributions. *Metrika*, 58(1): 85–97, 2003.

J.R. van Dorp, S.C. Rambaud, J.G. Pérez, and R.H. Pleguezuelo. An elicitation procedure for the generalized trapezoidal distribution with uniform central stage. *Decision Analysis*, 4(3):156–166, 2007.

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

E.M. Van Meter, E. Garrett-Mayer, and D. Bandyopadhyay. Dose-finding clinical trial design for ordinal toxicity grades using the continuation ratio model: an extension of the continual reassessment method. *Clinical Trials*, 9(3):303–313, 2012.

C. Wang, T. Chen, and I. Tyan. Designs for phase I cancer clinical trials with differentiation of graded toxicity. *Communications in Statistics-Theory and Methods*,

29:975–987, 2000.

J. Whitehead, S. Patterson, D. Webber, S. Francis, and Y. Zhou. Easy-to-implement Bayesian methods for dose-escalation studies in healthy volunteers. *Biostatistics*, 2(1):47–61, 2001.

J. Whitehead, Y. Zhou, A. Mander, S. Ritchie, A. Sabin, and A. Wright. An evaluation of Bayesian designs for dose-escalation studies in healthy volunteers. *Statistics in Medicine*, 25(3):433–445, 2006.

F.P. Worden, J.M.G. Taylor, J.S. Biermann, V.K. Sondak, K.M. Leu, R. Chugh, C.J. McGinn, M.M. Zalupski, and L.H. Baker. Randomized phase II evaluation of 6 g/m2 of ifosfamide plus doxorubicin and granulocyte colony-stimulating factor (G-CSF) compared with 12 g/m2 of ifosfamide plus doxorubicin and G-CSF in the treatment of poor-prognosis soft tissue sarcoma. *Journal of Clinical Oncology*, 23 (1):105–12, 2005.

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

J. Zhang and T.M. Braun. 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):892–901, 2013.

W. Zhang, D.J. Sargent, and S. Mandrekar. An adaptive dose-finding design incorporating both toxicity and efficacy. *Statistics in Medicine*, 25(14):2365–83, 2006.