Published in final edited form as:

J Biopharm Stat. 2016 ; 26(3): 475–498. doi:10.1080/10543406.2015.1052492.

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

Laura L. Fernandes, Jeremy M.G. Taylor, and Susan Murray*

*Department of Biostatistics, University of Michigan, Ann Arbor, Michigan, USA.

Abstract

Many phase I trials in oncology involve multiple dose administrations on the same patient over multiple cycles, with a typical cycle lasting three weeks and having about six cycles per patient with a goal to find the maximum tolerated dose (MTD) and study the dose-toxicity relationship. A patient's dose is unchanged over the cycles and the data is reduced to a binary end point, the occurrence of a toxicity and analyzed either by considering the toxicity from the first dose or from any cycle on the study. In this paper an alternative approach allowing an assessment of toxicity from each cycle and dose variations for patient over cycles is presented. A Markov model for the conditional probability of toxicity on any cycle given no toxicity in previous cycles is formulated as a function of the current and previous doses. The extra information from each cycle provides more precise estimation of the dose-toxicity relationship. Simulation results demonstrating gains in using the Markov model as compared to analyses of a single binary outcome are presented. Methods for utilizing the Markov model to conduct a phase I study, including choices for selecting doses for the next cycle for each patient, are developed and presented via simulation.

Keywords

Markov; repeated measures; Phase I; dose-finding; clinical trials; oncology; adaptive; Bayesian

1 Introduction/Background

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. An algorithmic approach could use the '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 & Chevret, 1991; Thall & 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. Anti-cancer drugs are typically administered in cycles and it is usual to only consider the probability of toxicity from the first cycle.

Recent publications (Postel-Vinay et al., 2014) that analyzed data from multiple phase I trials, found that the first occurrence of serious toxicity was frequently after the first cycle. In a survey of clinical trialists (Paoletti et al., 2014) more than 85% were in favor of considering serious side-effects after the first cycle as a dose limiting toxicity (DLT). This provides a rationale for considering toxicities from all cycles, and considering dose modifications in later cycles if severe toxicities are expected. The first few patients in a phase I study usually receive a cautiously low dose, which is not increased in later cycles, and hence probably receive limited benefit from the treatment. Simon et al. (1997) recognized the potential of increasing the dose in later cycles for early patients in the trial, to increase the potential benefit to the patients, as well as to learn more about the dose toxicity profile. This provided the rationale for the accelerated titration design. To evaluate this design a random effects models was used to simulate data. The model included a latent continuous toxic response W_{ij} for person i at time period j as,

$$W_{i,j} = \beta_i^S + \epsilon_{i,j} + \log \left(d_{i,j}^S + \alpha^S D_{i,j}^S \right) \quad (1)$$

where $d_{i,j}^S$ is the dose for person i at time j , $D_{i,j}^S$ is the i th person's cumulative dose prior to time j and the two random terms, $\beta_i^S \sim N \left(\mu_{\beta^S}, \sigma_{\beta^S}^2 \right)$ accounting for inter-patient variability or frailty and $\epsilon_{i,j} \sim N \left(0, \sigma_\epsilon^2 \right)$ representing the intra-patient variability. This model was used to generate data from a trial using a pre-defined escalation plan with the continuous toxicity response categorized into different levels using pre-defined thresholds.

Motivated by pharmacokinetic considerations (Legedza & Ibrahim, 2000) developed similar models for repeated toxicity measures. They proposed a model with three parameters of the form

$$\text{logit} \left(p_{i,j}^L \right) = \nu^L + \beta^L \log \left\{ d_{i,j}^L + D_{i,j}^L \exp \left(-\lambda^L \right) \right\} \quad (2)$$

where $p_{i,j}^L$ and $d_{i,j}^L$ are the probability of toxicity and the dose for person i at time j respectively, $D_{i,j}^L$ is the i th person's cumulative dose prior to time j . The ν^L parameter is not patient or cycle dependent and thus there is no correlation between responses on different cycles. A simpler model excluding the ν^L term was also considered. Due to the in-feasibility in estimating λ^L with small sample sizes it was assumed to be a constant, $\lambda^L = \log(2)$. Fitting Legedza's model requires the estimation of only two parameters and in the absence of the ν^L term a single parameter β^L . Notice that Legedza's models have increasing probability of toxicity with dose on subsequent cycles with no concession given to a patient for surviving a higher dose level on the first cycle.

If a patient does experience a toxicity on any cycle they are typically taken out of 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, such as β_i^S in the Simon et al. (1997) model, is an appealing way to incorporate concepts of frailty, fitting models with random effects to such data is challenging.

A recently developed cure rate model (Zhang & Braun, 2013) for estimating the cumulative effect of multiple administrations of the study drug 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_j at t_j is given by the formula, $h(t) = \theta_{ij} R(v_{t-t_j} | \phi)$, where $R(\cdot)$ is a distribution function. Thus the probability of having a DLT in an interval (t_k, t_{k+1}) can be calculated from $p_k = 1 - \exp\left(-\int_{t_k}^{t_{k+1}} h(u) du\right)$. The hazard from multiple administrations of a drug is obtained by summing the hazards. The model allows for delayed toxicities but it does not include any terms that would induce a correlation between repeated measures of toxicity for each cycle. Thus the model does not allow the probability of toxicity on the second cycle to be lower than that on the first 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 the trend in the 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 paper as an alternative we develop a two-state Markov model, with the states being 0 for no toxicity and 1 for toxicity. 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. The dose-toxicity relationship is described using a total of three parameters, capturing the effect of the current dose, the cumulative dose and the dependency on past responses through the maximum of the past doses.

Since we envision that the model would be fit during the conduct of the trial, an estimation method is needed for the three parameters 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 are 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 model 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 we describe the Markov model providing intuition on model features. The MCMC estimation method is described, with consideration given to the selection of the prior distributions. In Section 3 we evaluate properties of the estimation method in a static situation of a small and a moderate sample size. In Section 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 via simulations in Section 5 and end with a discussion in Section 6.

2 Methodology

2.1 Notation and data structure

We assume that there are G dose levels of an experimental study drug represented by d_g , $g = 1, \dots, G$, that will be studied in $i = 1 \dots N$ patients. We only consider $G = 5$, but the model is applicable to other values of $G > 1$. Each patient i completes $K_i - 6$ 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 $k = 1, \dots, K_i$, 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 $k = 1, \dots, K_i$ is $D_{i,k} = \sum_{j=1}^k d_{i,j-1}$, so that $D_{i,1} = 0$. We also use the notation that $d_{i,k1-} = 0$ for cycle $k = 1$ and

$d_{i,k}^{\dagger} = \max(d_{i,1}, \dots, d_{i,k-1})$, 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, thus 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 $\{(Y_{i,1}, \dots, Y_{i,K_i}, d_{i,1}, \dots, d_{i,K_i}), i = 1, \dots, n\}$.

2.2 Proposed Markov model

We propose Model 3 for the conditional probability of toxicity for patient i on cycle k given that patient i has experienced no previous DLTs as, $p_{i,k} = P(Y_{i,k} = 1 | Y_{i,k-1} = 0, \dots, Y_{i,1} = 0)$,

$$\ln(1 - p_{i,k}) = -\alpha(d_{i,k} - \rho d_{i,k}^{\dagger})^+ - \beta D_{i,k} d_{i,k} \quad (3)$$

or equivalently,

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

The term $(d_{i,k} - \rho d_{i,k}^{\dagger})^+$ is, equal to $(d_{i,k} - \rho d_{i,k}^{\dagger})$ if $d_{i,k} > \rho d_{i,k}^{\dagger}$, and is zero otherwise. Intuition behind Model 3 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, $\alpha \in [0, \phi]$ 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 $(d_{i,k} - \rho d_{i,k}^{\dagger})^+$ accounts for difference between the current assigned dose $d_{i,k}$ and a factor (ρ) of the maximum of the previous doses $d_{i,k}^{\dagger}$, 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 $\rho = 1$, $(d_{i,k} - \rho d_{i,k}^{\dagger})^+$ reduces to $(d_{i,k} - d_{i,k}^{\dagger})^+$ 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 $(d_{i,k} - \rho d_{i,k}^{\dagger})^+$ 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 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 3. 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 .

2.3 Prior and posterior distribution

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 & Cheung, 2009) in the context of the CRM and have been used by many other authors in other contexts (Lee et al., 2011; Cheung & Elkind, 2010). The probability skeleton information can be elicited from prior animal studies or from the clinicians.

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 DLT on cycle K_i gives data ($Y_{i,1} = 0, Y_{i,K_i-1} = 0, \dots, Y_{i,K_i} = 1, d_{i,1} \dots d_{i,K_i}$) and the contribution to the likelihood is,

$$P(Y_{i,1}=0, \dots, Y_{i,K_i-1}=0, Y_{i,K_i}=1) = p_{i,K_i} \prod_{j=1}^{K_i-1} (1 - p_{i,j}). \quad (4)$$

A person completing K cycles without DLT gives data ($Y_{i,1} = 0, \dots, Y_{i,K-1} = 0, Y_{i,K} = 0, d_{i,1} \dots d_{i,K}$) with likelihood contribution as,

$$P\left(Y_{i,1}=0, \dots, Y_{i,K-1}=0, Y_{i,K}=0\right)=\prod_{j=1}^K (1-p_{i,j}). \quad (5)$$

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 3 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 $p_{i,k}$, $k = 1, \dots, K$ 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

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

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 $Var(\alpha|\mu, \sigma) = \exp\{2(\mu + \sigma^2)\} - \exp\{2(\mu + \sigma^2)\} - \exp(2\mu + \sigma^2)$.

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 $\rho \in [0, 1]$ 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. We assume a high prior mean for ρ to represent situations where there is considerable memory about tolerability of previous doses.

The lognormal density is used as the prior on $\beta < 0$. In setting the prior mean for β two approaches could be considered. In general we believe that the contributions of the $\beta D_{i,k} d_{i,k}$

term to the probability of toxicity is likely to be much smaller than that of the $\alpha (d_{i,k} - \rho d_{i,k}^{\dagger})$ 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. Different choices of the prior distribution were considered in other research (L. Fernandes, 2014, PhD dissertation, University of Michigan). In this paper we use the prior distributions that were deemed to be most effective.

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

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

The posterior distribution of α , β and ρ from Model 3 is estimated via Markov Chain Monte Carlo (MCMC) methods (Robert & Casella, 1999) using just another Gibbs sampler (JAGS) *rjags* (Plummer, 2011) package through (R Development Core Team, 2011). Then the estimated probability of toxicity $p_{i,k}$ is obtained from equation 3 using the posterior means for α , β and ρ .

3 Operating characteristics/Results

Simulation results in this section study 1) the estimation of parameters when there may be intra-patient variation in the dose received and 2) the benefits in using the Markov Model 3 in comparison to two different models with single binary endpoints, using simulated data from completed trials.

3.1 Properties of parameter estimates with intra-patient dose variability

Two different cases are considered for comparison of estimation of parameters from a static data set of $N = 30$ patients. In the first case patients were distributed equally between the five dose levels, and the dose did not vary across cycles for each patient. In the second case the patients in the trial are assigned to pre-specified regimens, listed in Table A.1 for reference, which includes examples of constant doses and dose variations over the six cycles. A total of 500 datasets were generated under the two cases using the skeleton probabilities $q_g = (0.02, 0.05, 0.10, 0.15, 0.23)$. Conditional probability of toxicity, $p_{i,k}$, for each patient i at each cycle k was calculated using the Markov Model 3 for fixed values of $\alpha = 1$ and $\beta = 0.5$ and $\rho = 0.8$.

The results of the parameter estimates from the simulations are presented in Table 1 under Case 1 for patients with dose variations and under Case 2 for patients with constant dose.

The five columns report (1) the prior mean and SD, (2) mean of the estimated values from 500 datasets and the mean bias from the true value in parenthesis, (3) the mean SD (MSD) of the estimates from 500 datasets, (4) the empirical SD (ESD) of the 500 estimates, (5) the true value coverage rate in the 95% credible interval across the 500 datasets. We notice that the bias and mean SD of α is slightly lower in Case 1 but that of β is slightly higher. We notice a reduction in the MSD in comparison to the prior SD in the case of α and β indicating that these estimates are driven by the data, while there is less of a reduction in the SD of ρ , indicating a limited amount of information provided by the data for this parameter.

Table 2 presents the corresponding probability estimates. The columns present the probability of toxicity estimates on the first, the second, the sixth cycle and the overall probability of toxicity on any of the cycles, with the bias from the true value in parenthesis and the empirical SD (ESD) of the estimates from the 500 replicates. The results suggest that the model performs suitably in estimating the probabilities and that the bias is comparable in both cases and that the variability across simulation goes down slightly when patients have the same dose. We conclude that there were no problems in fitting the model to the data in which patients have dose variability and that the results do not have major deviations with regard to the bias and efficiency.

3.2 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 3 versus models with a single binary endpoint per patient. We consider the special case with no dose variation across cycles.

Simulation results are presented based on 500 datasets each having either $N = 10$ or $N = 30$ 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 3 and known values of $\alpha = 1$, $\beta = 0.5$ and $\rho = 0.8$. A DLT response $Y_{i,k}$ is assigned based on a Bernoulli($p_{i,k}$) 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 Section 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 3, in the first instance the data is reduced to a single binary outcome by defining $\dot{Y}_i = 1$ for patient i if $Y_{i,1} = 1$ and $\dot{Y}_i = 0$ if $Y_{i,1} = 0$. The probability of toxicity, p_i , on the first cycle is given by Model 6 as follows,

$$\ln(1 - p_i) = -\gamma d_i. \quad (6)$$

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 7 used in estimating the probability of toxicity on any cycle, p'_i in this case is,

$$\ln(1 - p'_i) = -\delta d'_i. \quad (7)$$

Where the probability of toxicity on any of the six cycles (m_j) corresponds to Markov Model 3 via

$$m_j = 1 - \prod_{j=1}^K (1 - p_{i,j}), \quad (8)$$

with $p_{i,j}$ as defined in equation 5. 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 3 indicates adequate model fit for the parameters from Markov Model 3 and Models 6 and 7. The MSD is lower than the prior SD indicating that the estimates of the posterior means are driven by the data. Table 4 presents the simulation results from comparing the Markov Model 3 to the two alternatives, Model 6 and Model 7. The rows are grouped based on the comparison with Model 6 or Model 7. 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 3 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 in fitting a larger model. The slight gain in efficiency in comparison to Model 7 could be attributed to fact the Markov Model 3 incorporates the cycle specific information in the process of estimating the parameters and hence provides better overall estimates of the probability of toxicity. This is useful because improved estimates of the probability of toxicity will lead to a higher chance of selecting the correct MTD.

4 Implementation of a clinical trial

This section describes the application of the Markov Model 3 in designing a sequential clinical trial. The safety criteria for dose assignment, a dose maximization plan and regimen selection at conclusion of the trial are presented.

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 $r_{g,k} = g$ as one of the five dose levels on cycle k corresponding to d_g , $g = 1 \dots 5$. Let $r_{g,k+1}^{max}$ denote the maximum allowed dose level that could be assigned on cycle $k + 1$ determined by incorporating commonly used dose escalation rules while carrying out an adaptive clinical trial based on Markov Model 3. Some of the safety rules prevent patients from jumping dose levels, assigning patients new dose levels only when the lower levels have been assigned previously with further details in Section A.3 of the Appendix.

4.2 Defining the eligible regimen set, $R_{i,k}^{regimen}$

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 $\vec{P(A)} = \{P(A_1), \dots, P(A_K)\}$ be a vector of acceptable toxicity limits for each cycle $1, \dots, K$. It is convenient to restrict limits for cycles $2, \dots, K$ to be equivalent and equal to $P(A_2)$ rather than justify different acceptable toxicity levels at each cycle. Define $P(C)$ 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 $P(B)$ be the acceptable probability of toxicity limit on all the remaining cycles. These bounds are pre-specified by the clinician before the start of the trial.

Monitoring the safety of the patients is ensured by assigning doses that satisfy the set of safety criteria as well as satisfy the bounds $P(A_1)$, $P(A_2)$, $P(B)$ and $P(C)$. In general denote dosing regimens by the vector of doses across the K cycles $(r_{g,1}, \dots, r_{g,K})$. As each patient progresses through cycles $k = 1 \dots K$, 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 $(o_{i,1}, \dots, o_{i,k-1}, r_{g,k}, r_{g,6})$ where $o_{i,k}$ denote previously tolerated doses for patient i on cycle k and future assigned doses $(r_{g,k}, \dots, r_{g,6})$ must not exceed $r_{g,l}^{max}$ for $l = k, \dots, 6$ and must not conflict with bounds defined by $P(A_1)$, $P(A_2)$ and $P(B)$ and $P(C)$. For a patient on cycle 1, it is convenient to limit members of $R_{i,1}^{regimen}$ to reduce computation. Table 5 lists a set of desirable regimens that can be used to construct a limited version of $R_{i,1}^{regimen}$ satisfying safety constraints. Regimen examples in Table 5 include constant dose over all the cycles, single dose variations on the fourth cycle, or double dose variations on the third and fifth cycle.

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 for that patient. Hence $P(A_{i,k,j}) = p_{i,k}$, where $p_{i,k}$ is calculated using the Markov Model 3 for dose level j , and current estimates of $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\rho}$ can be used to define its corresponding

estimate, $P(\hat{A}_{i,k,j}) = p_{i,k}$. 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 $P(B_{i,k,m}) = 1 - \prod_{l=k}^K (1 - p_{i,l})$ 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 $P(\hat{C}_{i,k,m}) = \hat{P}(\hat{B}_{i,1,m})$. During the course of the trial $P(\hat{B}_{i,k,m})$ estimates the current best guess of patient toxicity probability on the remaining cycles while $P(\hat{C}_{i,k,m})$ estimates the best guess of the toxicity probability profile for future patients undergoing regimen m .

4.3 Expected dose

The total dose received by a patient can be considered as a surrogate for efficacy, thus larger doses are better. However, 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 3 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 4 and 5 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,

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

4.4 Running the trial

In running a trial the decisions about what dose each patient should receive next are determined by the Markov Model 3 and the safety criteria laid out in section 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 $P(\hat{A}_{i,k,j})$, $P(\hat{B}_{i,k,m})$ and $P(\hat{C}_{i,k,m})$. Hence, development of an automated procedure is recommended for this trial design.

At the start of the trial two patients are assigned 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 recommendation 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 $\hat{P}(A_{i,1,j})$, $j=1 \dots r_{g,1}^{max}$. Eligible choices for dose level j must satisfy $\hat{P}(A_{i,1,j}) \leq \hat{P}(A_1)$ on cycle one and $\hat{P}(A_{i,k,j}) \leq \hat{P}(A_2)$ 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 $\hat{P}(A)$ safety criterion, and we consider not just the subsequent dose, but all remaining doses in making a recommendation. That is, 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 $R_{i,k}^{regimen}$. Maximizing the expected dose over the entire regimen for a patient is one of the strategies employed in choosing the dose levels. For a patient being dosed on cycle k , the expected dose is calculated using expressions in Section 4.3 for each of the regimens in $R_{i,k}^{regimen}$. The dose level that maximizes the expected dose and satisfies $\hat{P}(B_{i,k,m}) \leq \hat{P}(B)$ and $\hat{P}(C_{i,k,m}) \leq \hat{P}(C)$ is selected for cycle k .

4.5 Recommending a regimen

At the conclusion of the study, posterior estimates of $\hat{\alpha}$, $\hat{\beta}$, and $\hat{\rho}$, may be used to calculate characteristics of the $j = 1 \dots 19$ regimens listed in Table 5. In particular we calculate (1) the overall regimen toxicity probability $\hat{P}(C_j)$, $j = 1 \dots 19$, (2) the regimen's probability of toxicity on the first cycle $\hat{P}(A_{1,j})$, $j = 1 \dots 19$ and (3) the expected dose received for the regimen $\hat{E}_j$, $j = 1 \dots 19$.

Decisions of the best regimen can be chosen based on these quantities individually or in combination. In Table 6, we present an example of $\hat{P}(C_j)$, $\hat{P}(A_{1,j})$ and $\hat{E}_j$, $j = 1 \dots 19$ that one might see at the end of a particular clinical trial. (For the purposes of this example, these were calculated from Model 3, with parameters $\alpha = 1$, $\beta = 0.2$, $\rho = 0.8$). The rightmost 4 columns show example selection criteria based on meeting acceptable thresholds for $\hat{P}(A_{1,j})$ and $\hat{P}(C_j)$, with 1 indicating an acceptable regimen and 0 a regimen not meeting the criteria. The upper bounds for selecting a regimen are denoted by $\hat{P}(A_1)$ and $\hat{P}(C)$. When regimens are restricted to those satisfying $\hat{P}(C_j) \leq 0.30$, the maximum expected dose is attributed to regimen 12, with $(r_{g,1}, \dots, r_{g,K}) = (4, 4, 4, 3, 3, 3)$. However, $\hat{P}(A_{1,12})$ is a bit high at 16% for this regimen. When regimens are further restricted to those also satisfying $\hat{P}(A_{1,j}) \leq 0.05$, regimen 15 with $(r_{g,1}, \dots, r_{g,K}) = (2, 2, 3, 3, 4, 4)$ has the maximum expected dose of the acceptable regimens. Using these criteria, regimens that offered dose variation were found to be helpful in maximizing expected dose on the trial.

Regimens 1-5, 10-13 and 17-19 are examples that 'front-load' the higher doses during the earlier cycles, which could be desirable if the maximum impact of the treatment is thought to be associated with high dose challenges early on. When restricted to these regimens, and also requiring $\hat{P}(A_{1,j}) \leq 0.20$ and $\hat{P}(C_j) \leq 0.30$, then regimen 12 is once again selected.

5 Simulation and results of adaptive clinical trial

This section focuses on designing a clinical trial implementing the various safety criteria through simulations. The effect of using different values for $\hat{P}(A)$, $\hat{P}(B)$ and $\hat{P}(C)$ during trial

conduct were studied. A clinical trial recruiting a maximum of $N = 30$ patients and each patient having a maximum of six cycles was conducted over 500 replicates with the skeleton probability $q_g = (0.02, 0.05, 0.10, 0.23)$. The values of $\alpha = 1$, $\beta = 0.2$, $\rho = 0.8$ were used as the true values in generating the patient response at each of the cycles. An example of a clinical trial in progress is given in Section A.6 of the Appendix.

5.1 Properties of the trial design

For the purposes of evaluating the properties of the clinical trials with different target probabilities over multiple replications various statistics will be calculated that are explained below.

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 (R) 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. The target regimen is identified based on the true values of the parameters α , β , ρ used during the trial conduct and is one of the regimens in Table 5. A higher proportion indicates that patients in the trial received dose quantities considered safe by the target regimen and is desirable.
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 who followed the recommended regimen are averaged all the replications. Higher values are desirable indicating that patients in the trial had experience with the recommended regimen.
6. Mean 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 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 of the patients having distance 2 from the target (T) regimen. Higher values are desirable.

For statistics 7 and 8 the distance between two regimens is the sum over the 6 cycles of the absolute difference in dose levels.

For different combinations of $\vec{P(A)}$, $\vec{P(B)}$ and $\vec{P(C)}$ a total of 500 trials were simulated. Table 7 lists the results for the patient characteristics with the rows grouped by the option used to select the regimen, option 1 considers both $\vec{P(A_1)}$ and $\vec{P(C)}$ while option 2 considers only $\vec{P(C)}$ at the conclusion of the trials which are executed by maximizing the expected dose. The rows are further grouped by $\vec{P(B)} = (0.30, 0.40)$ within which each row correspond to cases of $\vec{P(A_1)} = (0.05, 0.10, 0.20)$ and $\vec{P(A_2)} = (0.09, 0.08, 0.06)$. The columns list the eight criteria presented above.

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 $\vec{P(A_1)} = 0.05$ or 0.10 imply stringent rules for selecting doses with higher toxicities but also results in less stringent rules on subsequent cycles with $\vec{P(A_2)} = 0.09$ or 0.08 . Based on results in Table 7 we notice that as the values of $\vec{P(A_1)}$ increase the mean dose received by the patients in the trial also increases. Secondly the observed mean toxicities tend to be low when $\vec{P(A_1)}$ is high, this could be because the toxicities on subsequent cycles are averted due to the low $\vec{P(A_2)}$ 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. The number of trials stopping early also decreases with increase in $\vec{P(A_1)}$ suggesting that higher values of $\vec{P(A_1)}$ allow patients to have dose assignment.

$\vec{P(A_1)}$ and $\vec{P(A_2)}$ flags are concerned with the selection of the recommended (R) and the target (T) regimen at the end of the trial. Consider results within Table 7 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 $\vec{P(A_1)}$ and $\vec{P(C)}$ are considered. The results in both the options are comparable in the instance when $\vec{P(A_1)} = 0.20$ (the third row). The number of trials that cannot offer a recommended regimen reduces when using a single condition for the regimen selection and these results indicate that having a restrictive condition on the first cycle plays a huge role in running of the trial and definitely affects the regimen selection.

Comparing results within Table 7 for differences in using $\vec{P(B)} = (0.30, 0.40)$, the results of the mean dose and toxicity of the trials does not seem to have huge variations. There does not seem to be any effect of using $\vec{P(B)}$ during the trial conduct. To further understand the effect of using both $\vec{P(C)}$ and $\vec{P(B)}$ versus only $\vec{P(C)}$ 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 only the $\vec{P(C)}$ condition. Simulations were carried out for a single setting of $\vec{P(A_1)} = 0.20$ and $\vec{P(C)} = 0.40$ and are presented in Table 8 which should be compared to the case when $\vec{P(A_1)} = 0.20$ from results in Table 7. The mean dose and toxicities are comparable in both the cases. The numbers are different in the case of patients having dose greater than the recommended dose, the patients with distance 2 is higher. Based on Table 8 it was seen

that having higher value of $P(\hat{C})$ 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 $P(\hat{A}_1)$ is higher.

6 Discussion

The Markov Model 3 presented in this manuscript 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 3 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 3 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 manuscript 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 . The simulation results for carrying out an adaptive clinical trial are affected by multiple factors, including parameters values, safety rules, target probabilities, the dose selection criteria and the regimen set. Further simulations need to be done to understand the performance of the method and its robustness in a broader set of situations. 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, and thus represent a first step in the direction of implementing trials that use toxicity data from multiple cycles.

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 Markov model presented, 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.

Simulations are presented for five dose levels 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 six cycles. 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. This simplistic assumption reduced the computational time of the simulations and also minimized the complexity of the results during comparisons. Alternatively, patient arrival could be generated in other ways which 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.

An alternative to maximizing the expected dose received by the patients is to gain the most possible experience with regimens that would be recommended at the end of the trial. In doing this one could “nudge” the current patient's regimen toward one of the $q = 1, \dots, 19$ suggested dosing regimens in Table 5, some subset of these or completely different user defined regimens not included in this table at every instance when more than one dose level qualifies for dose assignment in continuing patients. This approach was evaluated in simulations and results (not presented in this paper) indicated a higher number of patients matching the recommended regimen with lower toxicities and mean dose per patient. The set of recommendable regimens at the end of the trial does need to be specified. In Table 5 we present an example. Another example would be the six regimens which did not have any intra-patient dose variation; this might be appealing for reasons of simplicity.

The Markov model provides 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 give a similar 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.

The Markov model presented in this manuscript 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 attributed at the end of the cycle to the dose level for the patient in that cycle. In the case with delayed toxicities other models such as the model proposed by (Zhang & Braun, 2013) might be more appropriate.

Phase I studies typically have a small sample size, in the range of 15 to 35 patients. This small sample size leads to considerable uncertainty about whether the estimated MTD is really the best one. Statistical methods that make efficient use of the data can help reduce this uncertainty. The results in section 3.2 which demonstrate some efficiency gains compared to considering a simple binary toxicity measure, together with the imperative to

consider toxicities from all cycles (Postel-Vinay et al., 2014) indicate the need to develop and evaluate models such as the one in this paper.

To implement a phase I trial that uses information from multiple cycles requires not only the type of statistical models described in this paper, but also requires the logistical infrastructure to be able to gather the data and analyze it in real time. While this may be challenging, if the benefits of using the additional information to run a better trial are shown to be meaningful, then this challenge should be overcome.

A Appendix

A.1 Outline of the code written in JAGS

The code presented below corresponds to applying Model 3 in simulations for parameter estimation from the posterior samples.

```

1.      # Defining the model.bug file
2.      model {
3.      # Define the likelihood for each of N subjects
4.      for (i in 1:N) {
5.      prob[i] <- 1- exp (-alpha* (dose[i] - rho*maxprevdose[i])*step(dose[i] -
      rho*maxprevdose[i]) -
6.      beta*dose[i]*cumdose[i])
7.      response[i] ~ dbern (prob[i]
8.      }
9.      # Setting up the priors
10.     #prior on  $\alpha$  -  $E(\alpha) = 1$  and  $Var(\alpha) = 2$ 
11.     mul <- -0.8047190; tau1 <- 0.6213349
12.     alpha ~ dlnorm(mu1,tau1)
13.     #prior on  $\rho$ 
14.     a1 <- 5; b1 <- 1
15.     rho ~ dbeta (a1, b1)
16.     #Prior on  $\beta$  -  $E(\beta) = 0.5$  and  $Var(\beta) = 1$ 
17.     mu2 <- -1.498; tau2 <- 0.621
18.     beta ~ dlnorm (mu2,tau2)
19.     }
20.     #Initializing the parameters
21.     inits<-list(list(alpha=1, beta=0.1, rho=0.2),

```

```

22. list(alpha=0.5, beta=0.8, rho=0.9))
23. parameters <- c("alpha", "beta", "rho")
24. #updating the simulations
25. data <- list("response" = response, "maxprevdose" = maxprevdose,
26. "cumdose"=cumdose, "dose"=dose, "N"=N)
27. jags <- jags.model(file="prior.bug", data = data, inits=inits, n.chains = 2,
  n.adapt = 5000)
28. adapt(jags.model(file="prior.bug", data = data, inits=inits, n.chains = 2, n.adapt
  = 5000)
29. update(jags, 10000) # burn samples
30. sim1 <- coda.samples(jags, parameters, 100000, thin=20)
31. #check for convergence
32. plot(sim1)
33. gelman.plot(sim1)
34. geweke.plot(sim1)
35. geweke.plot(sim1)
36. geweke.diag(sim1)
37. autocorr(sim1)
38. autocorr.plot(sim1)
39. #report the mean and quantiles of the posterior distributions
40. y3 <- summary(sim1)
41. ystat <- data.frame(y3$statistics)
42. yquant = data.frame(y3$quantiles)

```

A.2 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 A.1

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

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
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

A.3 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 $r_{g,k+1}^{max}$ 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 3.

- 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 deescalation 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 $r_{g,k+1}^{max} = \min(r_{g,1} + 2, r_{g,k} + 1, 5)$ 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 $r_{g,1}^{max} = \max(r_{g,1}^{\dagger} + 1, r_{g,k}^{\dagger})$, where $r_{g,1}^{\dagger}$ is the maximum of all the past dose levels assigned to the patients on cycle $k = 1$ and $r_{g,k}^{\dagger}$ 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.

A.4 Algorithmic form of the dose maximization plan

Maximize the expected dose on the study for each patient i .

- For new patient $\hat{i}$ on cycle 1
 1. Estimate $P(\hat{A}_{\hat{i},1,j})$ using current estimates of $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\rho}$ at each of the $j=1 \dots r_{g,1}^{max}$ dose levels. Where $r_{g,1}^{max} = \max(r_{g,1}^{\dagger} + 1, r_{g,k}^{\dagger})$, $r_{g,1}^{\dagger}$ is the maximum of all the past dose levels assigned to the patients on cycle $k = 1$ and $r_{g,k}^{\dagger}$ 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 $P(\hat{A}_{\hat{i},1,j}) \geq P(\hat{A}_1)$ over all dose levels.
 3. For the dose levels satisfying $P(\hat{A}_{\hat{i},1,j}) \geq P(\hat{A}_1)$, subset the list of possible regimens from Table 5 and calculate the overall probability of toxicity $P(\hat{C}_{\hat{i},1,j})$ using the current estimates of $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\rho}$.
 4. Select the dose level that has an overall probability of toxicity $P(\hat{C}_{\hat{i},1,j}) \geq P(\hat{C})$ and maximum expected dose.

5. If none of the doses satisfy $P(\hat{A}_{i,1,j}) - P(\hat{A}_1)$ 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 $P(\hat{A}_{i,k+1,j}) - P(\hat{A}_2)$ from $r_{g,k+1}^{max} = \min(r_{g,1} + 2, r_{g,k} + 1, 5)$ possible choices.
 2. If there is more than one satisfying dose level then list the possible dose regimen set $R_{i,k+1}^{regimen}$.
 3. Calculate the probability of toxicity $P(\hat{B}_{i,k+1,m})$ on the remainder of the cycles for each of the regimens m in $R_{i,k+1}^{regimen}$ and the corresponding expected dose using the current estimates of $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\rho}$.
 4. Select the dose level that has probability to toxicity $P(\hat{B}_{i,k+1,m}) - P(\hat{B})$ and maximizes the expected dose.
 5. If none of the doses satisfy $P(\hat{A}_{i,k+1,j}) - P(\hat{A}_2)$ 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.

A.5 Algorithmic form of the regimen matching plan

Observing a favorable dosing regimen by the end of the study

- For new patient i on cycle 1
 1. Estimate $P(\hat{A}_{i,1,j})$ at each of the $j=1 \dots r_{g,1}^{max}$ dose levels. Where $r_{g,1}^{max} = \max(r_{g,1}^\dagger + 1, r_{g,k}^\dagger)$, $r_{z,1}^\dagger$ is the maximum of all the past dose levels assigned to the patients on cycle $k = 1$ and $r_{g,k}^\dagger$ 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 $P(\hat{A}_{i,1,j}) - P(\hat{A}_1)$ over all dose levels.
 3. For the dose levels satisfying $P(\hat{A}_{i,1,j}) - P(\hat{A}_1)$ subset the list of possible regimens from Table 5 and calculate the overall probability of toxicity $P(\hat{C}_{i,1,j})$ using the current estimates of $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\rho}$.
 4. Select the dose level that has an overall probability of toxicity $P(\hat{C}_{i,1,j})$. Select the highest dose if more than one satisfying dose level.

5. If none of the doses satisfy $P(\hat{A}_{i,1,j}) \leq P(A_1)$ 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 $P(\hat{A}_{i,k+1,j}) \leq P(A_2)$ from $r_{g,k+1}^{max} = \min(r_{g,1} + 2, r_{g,k} + 1, 5)$ possible choices.
 2. If there is more than one satisfying dose level then list the possible dose regimen $R_{i,k+1}^{regimen}$ set.
 3. Subset the regimens satisfying probability of toxicity $P(\hat{B}_{i,k+1,m}) \leq P(B)$ from the $R_{i,k+1}^{regimen}$ set.
 4. For each of the regimens l in the subset calculate the distance R_l from the favorable dosing regimens based on Table 5.
 5. Select the highest dose level that has R_l 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 $P(\hat{A}_{i,k+1,j}) \leq P(A_2)$ 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.

A.6 Example of an adaptive trial in progress

Based on the algorithm presented in Section A.4 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 $P(C) = 0.40$, $P(B) = 0.30$, $P(A_1) = 0.05$, $P(A_2) = 0.09$ and the true values of the parameters are $\alpha = 1$, $\beta = 0.2$, $\rho = 0.8$.

Table A.2 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 $P(A_2) = 0.09$ and is also able to provide a regimen combination that satisfies the $P(C) = 0.40$, $P(B) = 0.30$ 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 A.3. 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 A.4. 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 5 to select the recommended regimen and by using the true parameter values to select the target regimen. By setting $P(A_1) = P(\bar{A}_1) = 0.05$ and $P(C) = 0.3$ 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 $P(C) = 0.3$ is used the target regimen is 444333 while the recommended regimen is 333333.

Table A.2

Table showing the dose level assignment and patient responses in parenthesis for an adaptive trial in progress.

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

Table A.3

Table showing the dose level assignment and patient responses in parenthesis for an adaptive trial in progress.

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

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
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)	?	?	?	?	?

Table A.4

Table showing the dose level assignment and patient responses in parenthesis for an adaptive trial in progress.

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)

Patient ID	Cycle 1	Cycle 2	Cycle 3	Cycle 4	Cycle 5	Cycle 6
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

References

- Babb J, Rogatko A, Zacks S. Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine*. 1998; 17:1103–1120. [PubMed: 9618772]
- Cheung Y, Elkind M. Stochastic approximation with virtual observations for dose-finding on discrete levels. *Biometrika*. 2010; 97:109–121. [PubMed: 23049118]
- Doussau A, Thiébaud R, Paoletti X. Dose-finding design using mixed-effect proportional odds model for longitudinal graded toxicity data in phase I oncology clinical trials. *Statistics in Medicine*. 2013; 32:5430–5447. [PubMed: 24018535]
- Lee S, Cheng B, Cheung Y. Continual reassessment method with multiple toxicity constraints. *Biostatistics*. 2011; 12:386–98. [PubMed: 20876664]
- Lee S, Cheung Y. Model calibration in the continual reassessment method. *Clinical Trials*. 2009; 6:227–238. [PubMed: 19528132]
- Lagedza A, Ibrahim J. Longitudinal design for phase I clinical trials using the continual reassessment method. *Controlled Clinical Trials*. 2000; 21:574 – 588. [PubMed: 11146151]
- O'Quigley J, Chevret S. Methods for dose finding studies in cancer clinical trials: a review and results of a Monte Carlo study. *Statistics in Medicine*. 1991; 10:1647–64. [PubMed: 1792460]
- O'Quigley J, Pepe M, Fisher L. Continual reassessment method: A practical design for phase I clinical trials in cancer. *Biometrics*. 1990; 46:33–48. [PubMed: 2350571]
- Paoletti X, Le Tourneau C, Verweij J, Siu LL, Seymour L, Postel-Vinay S, Collette L, Rizzo E, Ivy P, Olmos D, Massard C, Lacombe D, Kaye SB, Soria J-C. Defining dose-limiting toxicity for phase I trials of molecularly targeted agents: Results of a DLT-TARGETT international survey. *European Journal of Cancer*. 2014; 50:2050–2056. [PubMed: 24928189]
- Plummer M. *rjags: Bayesian graphical models using MCMC*. 2011
- Postel-Vinay S, Collette L, Paoletti X, Rizzo E, Massard C, Olmos D, Fowst C, Levy B, Mancini P, Lacombe D, Ivy P, Seymour L, Tourneau CL, Siu LL, Kaye SB, Verweij J, Soria J-C. Towards new methods for the determination of dose limiting toxicities and the assessment of the recommended dose for further studies of molecularly targeted agents – Dose-Limiting Toxicity and Toxicity Assessment Recommendation Group for Early Trials of Targeted therapies, an European Organisation for Research and Treatment of Cancer-led study. *European Journal of Cancer*. 2014; 50:2040–2049. [PubMed: 24880774]
- R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing; Vienna, Austria: 2011.
- Robert, C.; Casella, G. *Monte Carlo Statistical Methods*. Springer; New York: 1999.
- Simon R, Freidlin B, Rubinstein L, Arbuck S, Collins J, Christian M. Accelerated titration designs for phase I clinical trials in oncology. *Journal of the National Cancer Institute*. 1997; 89:1138–47. [PubMed: 9262252]

- Storer B. Design and analysis of phase I clinical trials. *Biometrics*. 1989; 45:925–937. [PubMed: 2790129]
- Thall P, Lee S. Practical model-based dose-finding in phase I clinical trials: methods based on toxicity. *International Journal of Gynecological Cancer*. 2003; 13:251–61. [PubMed: 12801254]
- Whitehead J, Patterson S, Webber D, Francis S, Zhou Y. Easy-to- implement Bayesian methods for dose-escalation studies in healthy volunteers. *Biostatistics*. 2001; 2:47–61. [PubMed: 12933556]
- Whitehead J, Zhou Y, Mander A, Ritchie S, Sabin A, Wright A. An evaluation of Bayesian designs for dose-escalation studies in healthy volunteers. *Statistics in Medicine*. 2006; 25:433–445. [PubMed: 16220506]
- Zhang J, Braun T. A phase I Bayesian adaptive design to simultaneously optimize dose and schedule assignments both between and within patients. *Journal of the American Statistical Association*. 2013; 108:892–901.

Figure 1.

Conditional P(toxicity) with cycle 1 as the reference based on Model 3. 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 1

Parameter estimates from two different patient profiles, based on Model 3 from 500 simulated datasets containing $N=30$ patients with dose variability in Case 1 and same dose level in Case 2, receiving one of five dose groups $d_1 \dots d_5$ over six cycles. Results presented for $\alpha = 1$, $\beta = 0.5$ and $\rho = 0.8$

	Prior Mean (SD)	Estimate(Bias)	MSD ¹	ESD ²	Coverage
Case 1					
α	1 (2)	0.989 (−0.011)	0.446	0.425	95.2
β	0.5 (1)	0.561 (0.061)	0.361	0.322	98
ρ	0.83 (0.14)	0.808 (0.008)	0.134	0.056	100
Case 2					
α	1 (2)	0.957 (−0.043)	0.458	0.417	96.6
β	0.5 (1)	0.511 (0.011)	0.338	0.296	97.8
ρ	0.83 (0.14)	0.806 (0.006)	0.132	0.056	100

¹MSD is mean of the SD from 500 estimates

²ESD is empirical SD of the 500 estimates

Table 2

Estimates with bias from the true value of conditional probability of toxicity on each of the six cycles estimated using Model 3 from 500 simulated datasets containing $N=30$ patients with dose variability in Case 1 and same dose level in Case 2, 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	
Case 1	Est(bias)	SD(est)	Est(bias)	SD(est)	Est(bias)	SD(est)	Est(bias)	SD(est)
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
Case 2								
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

Table 3

Parameter estimates from 500 simulated datasets containing $N = 10$ and $N = 30$ patients distributed equally to receive one of five dose groups $d_1 \dots d_5$ over six cycles comparing estimates based on Model 3 (α , β , ρ) that use multiple toxicity responses to estimates based on Models 6 (γ) and 7 (δ) that use a single binary toxicity response.

	Prior Mean (SD)	Estimate(Bias)	MSD ¹	ESD ²	Coverage
N=10 patients					
α	1 (2)	0.945 (0.055)	0.682	0.585	99.6
β	0.5 (1)	0.515 (−0.015)	0.464	0.308	100
ρ	0.83 (0.14)	0.820 (−0.020)	0.137	0.040	100
γ	1 (2)	0.916 (0.084)	0.703	0.630	99.8
δ	1 (2)	0.976 (0.024)	0.548	0.536	92.6
N=30 patients					
α	1 (2)	0.989 (0.011)	0.466	0.408	96.0
β	0.5 (1)	0.507 (−0.007)	0.343	0.285	98.8
ρ	0.83 (0.14)	0.802 (−0.002)	0.133	0.058	99.8
γ	1 (2)	0.940 (0.060)	0.474	0.428	96.6
δ	1 (2)	0.998 (0.002)	0.333	0.331	94

¹MSD is mean of the SD from 500 estimates

²ESD is empirical SD of the 500 estimates

Table 4

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 3 to estimates based on Models 6 and 7. Results presented for $\alpha = 1$, $\beta = 0.5$ and $\rho = 0.8$ and with patients distributed equally on the five dose levels.

	$N = 10$		$N = 30$	
	Estimate(Bias)	ESD ^I	Estimate(Bias)	ESD ^I
First cycle				
Markov Model 3				
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 6				
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 3				
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 7				
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

^IESD is empirical SD of the 500 estimates

Table 5

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 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, $P'(C) = 0.30$ and three instances of using option 2, $P'(C) = 0.30$ and $P'(A_1) = (0.05, 0.10, 0.20)$ corresponding to the 19 favorable dose regimen listed in Table 5. $\alpha = 1$, $\beta = 0.2$, $\rho = 0.8$ are the true values of the parameters.

J Biopharm Stat. Author manuscript; available in PMC 2017 January 01.

Table 7

Trial/Patient summary results for maximizing expected dose over 500 simulated clinical trials. Columns present the 1) mean dose received per patient, 2) mean toxicities over the trials with trial stopping early in parenthesis, 3) average patients having dose the recommended expected dose, 4) average patients having dose 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 to the recommended regimen and 8) average patients having regimen to the target regimen.

	Mean Dose	Mean Toxicities	Patients R.Edose	Patients T.Edose	Patients =R.Dist	Patients =T.Dist	P.R.Dist 2	P.T.Dist 2
Using $P^r(A_1) \& P^r(C) = 0.30$								
$P(\dot{B}) = 0.3,$ $P(\dot{C}) = 0.4$								
$R(\dot{A}_1) = 0.05,$ $R(\dot{A}_2) = 0.09$	15.48	0.31 (48)	20.27	15.02	0.62	0.83	5.66	5.35
$R(\dot{A}_1) = 0.1,$ $R(\dot{A}_2) = 0.08$	15.91	0.3 (25)	18.09	18.09	1.64	0.64	11.75	4.89
$R(\dot{A}_1) = 0.2,$ $R(\dot{A}_2) = 0.06$	16.21	0.27 (11)	16.73	17.22	3.96	1.83	10.86	8.57
$P(\dot{B}) = 0.4,$ $P(\dot{C}) = 0.4$								
$R(\dot{A}_1) = 0.05,$ $R(\dot{A}_2) = 0.09$	15.53	0.31 (48)	20.28	14.96	0.6	0.83	5.4	5.24
$R(\dot{A}_1) = 0.1,$ $R(\dot{A}_2) = 0.08$	15.99	0.31 (25)	17.98	17.68	1.59	0.69	11.3	4.67
$R(\dot{A}_1) = 0.2,$ $R(\dot{A}_2) = 0.06$	16.33	0.28 (11)	16.82	17.38	4.46	2.16	11.26	9.31
Using $P^r(C) = 0.30$								
$P(\dot{B}) = 0.3,$ $P(\dot{C}) = 0.4$								
$R(\dot{A}_1) = 0.05,$ $R(\dot{A}_2) = 0.09$	15.48	0.31 (48)	12.84	12.71	0.45	0.13	8.15	2.09
$R(\dot{A}_1) = 0.1,$ $R(\dot{A}_2) = 0.08$	15.91	0.3 (25)	15.46	13.51	1.81	0.68	11.09	4.95
$R(\dot{A}_1) = 0.2,$ $R(\dot{A}_2) = 0.06$	16.21	0.27 (11)	16.73	17.22	3.95	1.83	10.81	8.57
$P(\dot{B}) = 0.4,$ $P(\dot{C}) = 0.4$								
$R(\dot{A}_1) = 0.05,$ $R(\dot{A}_2) = 0.09$	15.53	0.31 (48)	12.81	12.77	0.43	0.12	7.68	2.22
$R(\dot{A}_1) = 0.1,$ $R(\dot{A}_2) = 0.08$	15.99	0.31 (25)	15.29	13.37	1.72	0.65	10.71	5.44

	Mean Dose	Mean Toxicities	Patients R.Edose	Patients T.Edose	Patients =R.Dist	Patients =T.Dist	P.R.Dist 2	P.T.Dist 2
$P(A_1) = 0.2,$ $P(A_2) = 0.06$	16.33	0.28 (11)	16.82	17.38	4.47	2.16	11.24	9.31

T/R.Edose - Target/Recommended expected dose.

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

Table 8

Trial/Patient summary -using only $P(\hat{C})$ flag in carrying out 500 simulated trials.

	Mean Dose	Mean Toxicities	Patients R.Edose	Patients T.Edose	Patients =R.Dist	Patients =T.Dist	P.R.Dist 2	P.T.Dist 2
$P^*(A_1) \& P^*(C) = 0.30$								
$P(\hat{C}) = 0.30$ $P(\hat{A}_1) = 0.20$, $P(\hat{A}_2) = 0.03$	15.78	0.24 (2)	13.03	15.12	2.39	1.44	10.21	6.28
$P(\hat{C}) = 0.40$ $P(\hat{A}_1) = 0.20$, $P(\hat{A}_2) = 0.06$	16.67	0.27 (2)	16.35	18.19	4.86	2.18	11.92	8.95
Using $P^*(C) = 0.30$ in regimen selection								
$P(\hat{C}) = 0.30$ $P(\hat{A}_1) = 0.20$, $P(\hat{A}_2) = 0.03$	15.78	0.24 (2)	13.02	15.12	2.39	1.44	10.05	6.28
$P(\hat{C}) = 0.40$ $P(\hat{A}_1) = 0.20$, $P(\hat{A}_2) = 0.06$	16.67	0.27 (2)	16.34	18.19	4.94	2.18	11.83	8.95

T/R.Edose - Target/Recommended expected dose.

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