

Monitoring late-onset toxicities in phase I trials using predicted risks

B. NEBIYOU BEKELE*, YUAN JI, YU SHEN, PETER F. THALL

Department of Biostatistics, M.D. Anderson Cancer Center, Houston, TX 77030, USA
bkele@mdanderson.org

SUMMARY

Late-onset (LO) toxicities are a serious concern in many phase I trials. Since most dose-limiting toxicities occur soon after therapy begins, most dose-finding methods use a binary indicator of toxicity occurring within a short initial time period. If an agent causes LO toxicities, however, an undesirably large number of patients may be treated at toxic doses before any toxicities are observed. A method addressing this problem is the time-to-event continual reassessment method (TITE-CRM, Cheung and Chappell, 2000). We propose a Bayesian dose-finding method similar to the TITE-CRM in which doses are chosen using time-to-toxicity data. The new aspect of our method is a set of rules, based on predictive probabilities, that temporarily suspend accrual if the risk of toxicity at prospective doses for future patients is unacceptably high. If additional follow-up data reduce the predicted risk of toxicity to an acceptable level, then accrual is restarted, and this process may be repeated several times during the trial. A simulation study shows that the proposed method provides a greater degree of safety than the TITE-CRM, while still reliably choosing the preferred dose. This advantage increases with accrual rate, but the price of this additional safety is that the trial takes longer to complete on average.

Keywords: Adaptive design; Bayesian inference; Dose finding; Isotonic regression; Latent variables; Markov chain Monte Carlo; Ordinal modeling; Predictive probability.

1. INTRODUCTION

Phase I clinical trials focus on treatment-related adverse events so severe that they impose a practical limitation to the delivery of therapy. Such an event is called a dose-limiting toxicity (DLT) or simply “toxicity.” The primary scientific goal in a phase I trial is to determine the maximum tolerable dose (MTD), which is the highest dose with an acceptable risk of toxicity (cf. Storer, 1989; Babb *and others*, 1998). Because most DLTs occur soon after the start of therapy, most phase I methods are based on a binary variable indicating that a DLT has occurred within a time interval of fixed length, t^* , which we will call the assessment window.

The possibility of late-onset (LO) toxicities is an important concern in phase I trials. Basing dose finding on DLTs scored within a fixed window t^* is a practical compromise motivated by algorithms that choose doses sequentially for successive patients based on previous patients’ doses and outcomes.

*To whom correspondence should be addressed.

Conventional methods consider a patient to be fully evaluated if he/she experiences toxicity before t^* or is followed for this length of time without toxicity. A problem with such methods is that they may cause investigators to treat an undesirably large number of patients at toxic doses before LO toxicities are first observed. While a safer approach would be to use a very large t^* and wait for each patient's outcome before choosing the next patient's dose, in most settings this would produce an infeasibly long trial.

A method that deals with LO toxicities is the time-to-event continual reassessment method (TITE-CRM, Cheung and Chappell, 2000), which uses the time to toxicity, T , as the outcome, rather than a binary indicator. Denoting the cumulative distribution function (cdf) of T for dose d by $F_T(t; d, \boldsymbol{\theta}) = \Pr(T \leq t | d, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is the model parameter, the TITE-CRM relies on the fact that $F_T(t; d, \boldsymbol{\theta}) = \Pr(T \leq t | d, \boldsymbol{\theta}, T \leq t^*) \Pr(T \leq t^* | d, \boldsymbol{\theta})$ for $t \leq t^*$, approximates $\Pr(T \leq t | d, \boldsymbol{\theta}, T \leq t^*)$ by a weight function $w(t, t^*)$ such as t/t^* or a sample-based estimator, and uses $\pi_w(t, t^*, d) = w(t, t^*) \Pr(T \leq t^* | d, \boldsymbol{\theta})$ as a weighted version of the usual probability in the CRM (O'Quigley and others, 1990). Indexing patients by $i = 1, \dots, n$ and denoting the time to toxicity or right censoring for the i th patient by T_i^o , the resulting approximate likelihood is

$$\tilde{L}_n(\boldsymbol{\theta}) = \prod_{i=1}^n \pi_w(T_i^o, t^*, d_{(i)})^{I\{T_i \leq \min(T_i^o, t^*)\}} \{1 - \pi_w(T_i^o, t^*, d_{(i)})\}^{I\{T_i > \min(T_i^o, t^*)\}}, \quad (1.1)$$

where $d_{(i)}$ is the dose for patient i . Using this approximation, the TITE-CRM applies the usual CRM criterion by choosing the dose with posterior mean $E\{\pi_w(T, t^*, d) | \text{data}\}$ closest to a given target, π^* . Although the TITE-CRM does not increase the risk of LO toxicities when accrual is slow, if accrual is faster it has a higher risk of treating patients at unsafe doses. This is because, with faster accrual, new patients are more likely to be assigned to doses considered safe but later found to be unsafe if patients who have not been fully evaluated experience LO toxicities.

In this paper, we address the problem of LO toxicities by using predictive probabilities to quantify the risks that prospective doses for future patients will turn out to be excessively toxic. A fundamental difference between our method and existing methods is that we provide rules for suspending accrual if the predicted risk of toxicity (PRT) is unacceptably high. If additional follow-up data reduce the predicted risk at any prospective dose to an acceptable level, then accrual is restarted. We present a simulation study showing that, compared to the TITE-CRM, our proposed method provides an extra measure of safety when toxicities are likely to occur late in the assessment window, and this advantage increases with the patient accrual rate. The price of this improved safety is that, on average, the trial takes longer to complete.

The probability model is presented in Section 2. Priors and posteriors are developed in Section 3. Decision criteria and rules for trial conduct are presented in Section 4. In Section 5, we provide guidelines and an illustration of the method, and in Section 6, we present a computer simulation study, including comparison to the TITE-CRM. We conclude with a discussion in Section 7.

2. PROBABILITY MODEL

To describe the relationship between dose and time to toxicity, we use the sequential ordinal model of Albert and Chib (2001), which provides flexibility in modeling the hazard function while facilitating computation. We first fix a sequence of times $0 = t_0 < t_1 < \dots < t_{C-1} < t_C = \infty$, where $[t_0, t_{C-1}] = [0, t^*]$ is the assessment window, and we define $Y_i = j$ if patient i experiences toxicity in the j th interval, $t_{j-1} \leq T_i < t_j$, for $j = 1, \dots, C$. Thus, Y_i is a discretized ordinal representation of T_i , with $Y_i = C$ if toxicity does not occur within the assessment window. Given the current time T_i^o to toxicity or right censoring, we denote the usual indicator $\delta_i = 1$ if $T_i^o = T_i$ and $\delta_i = 0$ if $T_i^o < T_i$, and we also define $Y_i^o = j$ if $t_{j-1} \leq T_i^o < t_j$ so that Y_i^o is the current interval index for the i th patient. Denote the doses to be studied by $d_1 < \dots < d_K$. For $j = 1, \dots, C - 1$ and $k = 1, \dots, K$, denoting the standard normal

cdf by $\Phi(\cdot)$, we define the discrete-time hazard $\Phi(\beta_{j,k}) = \Pr(Y_1 = j | Y_1 \geq j, d_k)$ such that $\beta_{j,k}$ is the effect of dose k on the hazard of toxicity in the j th interval. It can be shown that the probability of toxicity during the j th interval is $\Pr(Y_1 = j | d_k) = \Phi(\beta_{j,k}) \prod_{h=1}^{j-1} \{1 - \Phi(\beta_{h,k})\}$, and the probability of toxicity not occurring by t_j is $\Pr(Y_1 > j | d_k) = \prod_{h=1}^j \{1 - \Phi(\beta_{h,k})\}$ for $j \leq C - 1$. Let $k(i)$ be the index (level) of the dose administered to the i th patient. At any point in the trial, the discretized data from the current n patients take the form $D_n = \{(Y_i^o, k(i), \delta_i), i = 1, \dots, n\}$ and, denoting $\boldsymbol{\beta} = (\beta_{1,1}, \dots, \beta_{C-1,K})$, the likelihood is

$$L(\boldsymbol{\beta} | D_n) = \prod_{i=1}^n \Phi(\beta_{Y_i^o, k(i)})^{\delta_i} \prod_{h=1}^{Y_i^o-1} \{1 - \Phi(\beta_{h, k(i)})\}. \quad (2.1)$$

Our procedure will be based on the conditional probabilities $\pi(\boldsymbol{\beta}, d_k, j) = \Pr(Y \leq C - 1 | Y \geq j, \boldsymbol{\beta}, d_k)$ for $j = 1, \dots, C - 1$ and $k = 1, \dots, K$. This is motivated by the fact that $\pi(\boldsymbol{\beta}, d_k, Y^o)$ is the probability that a patient who has survived $Y^o - 1$ intervals without toxicity will experience toxicity by $t^* = t_{C-1}$ at dose d_k . Since $\Pr(Y \geq 1) = 1$, the unconditional probability of toxicity within the window $[0, t^*]$ is $\pi(\boldsymbol{\beta}, d_k, 1)$. We will abuse the notation slightly by writing this as $\pi(\boldsymbol{\beta}, d_k)$. It follows from the discrete-time hazard given above that $\pi(\boldsymbol{\beta}, d_k, j) = 1 - \prod_{h=j}^{C-1} \{1 - \Phi(\beta_{h,k})\}$, and in particular $\pi(\boldsymbol{\beta}, d_k) = 1 - \prod_{h=1}^{C-1} \{1 - \Phi(\beta_{h,k})\}$.

To facilitate computation of posterior quantities, following the general computational strategy of Albert and Chib (1993), we express the likelihood (2.1) using a latent variable formulation. For patient i , define the vector of latent variables $\mathbf{Z}_i = (Z_{i,1}, \dots, Z_{i, Y_i^o})$ if $Y_i^o < C$ and $\mathbf{Z}_i = (Z_{i,1}, \dots, Z_{i, C-1})$ if $Y_i^o = C$, with $Z_{i,j} \sim N(\beta_{j,k}, 1)$ if the i th patient received dose d_k . Denote the $N(\mu, \sigma^2)$ density by $\phi(z; \mu, \sigma^2)$. Since $\Phi(\mu) = \int_{-\infty}^0 \phi(z; \mu, 1) dz$, the likelihood (2.1) may be expressed equivalently in terms of the latent variables as

$$L(\boldsymbol{\beta} | D_n) = \prod_{i=1}^n \left\{ \int_{-\infty}^0 \phi(Z_{i, Y_i^o}; \beta_{Y_i^o, k(i)}, 1) dZ_{i, Y_i^o} \right\}^{\delta_i} \prod_{j=1}^{Y_i^o-1} \int_0^{\infty} \phi(Z_{i,j}; \beta_{j, k(i)}, 1) dZ_{i,j}. \quad (2.2)$$

Comparing the representation given in (2.2) to the original form given in (2.1) shows that one can think of $Z_{i,j}$ as a continuous latent variable with mean equal to the patient's risk of toxicity during the j th interval given the dose administered. The general strategy is to reformulate (2.2) by including $\mathbf{Z} = (\mathbf{Z}_1, \dots, \mathbf{Z}_n)$ as arguments of the likelihood rather than integrating them out. The advantage of this "data-augmentation" technique, in terms of latent quantities, is that it provides full conditional distributions for the Gibbs sampling algorithm when computing the posterior. Without this device, computing posteriors would be very inefficient and time consuming. Thus, the augmented likelihood may be expressed as

$$L(\boldsymbol{\beta} | D_n, \mathbf{Z}) = \prod_{i=1}^n \left\{ \phi(Z_{i, Y_i^o}; \beta_{Y_i^o, k(i)}, 1) I(Z_{i, Y_i^o} < 0) \right\}^{\delta_i} \prod_{j=1}^{Y_i^o-1} \phi(Z_{i,j}; \beta_{j, k(i)}, 1) I(Z_{i,j} > 0). \quad (2.3)$$

Using this augmented likelihood, the latent variable vector \mathbf{Z} is sampled repeatedly from its full conditional distribution as a step in the Albert–Chib Markov chain Monte Carlo (MCMC) framework.

3. PRIORS AND POSTERIORS

For each interval $[t_{j-1}, t_j)$, $j = 2, \dots, C - 1$, we assume a state-space model for the prior of $\boldsymbol{\beta}_j = (\beta_{j,1}, \dots, \beta_{j,K})$, defined by the recursive relationship $\beta_{j,k} | \beta_{j,k-1} \sim N(\beta_{j,k-1}, \sigma_\beta^2)$ for $k = 2, \dots, K$,

with $\beta_{j,1} \sim N(\beta_{j,0}, \sigma_\beta^2)$ and $\beta_{j,0}$ fixed to ensure identifiability. Under this prior, $E(\beta_{j,k}) = E\{E(\beta_{j,k} | \beta_{j,k-1})\} = E(\beta_{j,k-1})$, which by recursion implies that $E(\beta_{j,k}) = \beta_{j,0}$ for all $k = 1, \dots, K$. However, $\text{var}(\beta_{j,k}) = k\sigma_\beta^2$ so that, with this prior, the uncertainty about the risk of toxicity increases with dose.

This model does not assume a parametric dose–toxicity function. Rather, $\boldsymbol{\beta}$ is used to characterize the probability of toxicity $\pi(\boldsymbol{\beta}, d_k, j)$ for $j = 1, \dots, C - 1$ and $k = 1, \dots, K$. Although the model can flexibly characterize the relationship between dose and toxicity, it does not guarantee that $\pi(\boldsymbol{\beta}, d_k, j)$ increases with d_k . Although a prior constraint on $\Phi(\beta_{j,k-1}) < \Phi(\beta_{j,k})$ for all k and j would guarantee that the $\pi(\boldsymbol{\beta}, d_k, j)$ are monotone increasing in d_k , imposing such a constraint would lead to highly biased posterior estimates of the toxicity probabilities and would also lead to major computational difficulties. Thus, we take the alternative approach of applying the Bayesian isotonic regression method of Dunson and Neelon (2003). This method produces transformed versions of $\pi(\boldsymbol{\beta}, d_k, j)$ that are monotone increasing in d_k by taking weighted averages of subvectors of $\boldsymbol{\pi}(\boldsymbol{\beta}, j) = (\pi(\boldsymbol{\beta}, d_1, j), \dots, \pi(\boldsymbol{\beta}, d_K, j))$. In the present setting, this ensures that the probability of toxicity by $t^* = t_{C-1}$ is monotonically increasing in dose, while retaining the flexibility and tractability of the state-space model on $\boldsymbol{\beta}_k$ without introducing bias into posterior estimates for toxicity. We use the following 3-step algorithm to compute posteriors, applied each time a new patient is enrolled. These steps combine the Albert–Chib method for exploiting latent variables and the Dunson–Neelon method for enforcing monotonicity, in the context of our dose–toxicity model. Denote $A_0 = (-\infty, 0]$, $A_1 = (0, \infty)$, and $\tilde{A}_{i,j} = A_{1-I(Y_i^o=j, \delta_i=1)}$. The following process is initialized using the prior mean of $\boldsymbol{\beta}$; Steps 1 and 2 are iterated until convergence and then Step 3 is applied.

Step 1: Generation of the latent variables. Generate each $Z_{i,j}$ independently from the full conditional which follows a truncated normal distribution $\phi(z; \beta_{j,k(i)}, 1)I(z \in \tilde{A}_{i,j})$.

Step 2: Generation of $\boldsymbol{\beta}$. Denote

$$S_{j,k} = \sum_{i=1}^n \sum_{h=1}^{Y_i^o} I\{k[i] = k, h = j\} \sigma_\beta^2 \quad \text{and} \quad Z_{j,k}^+ = \sum_{i=1}^n \sum_{h=1}^{Y_i^o} Z_{i,h} I\{k[i] = k, h = j\}.$$

Given \mathbf{Z} and the current data, generate $\boldsymbol{\beta}$ from its full conditional distribution under which, for $k = 1, \dots, K$, $\beta_{j,k}$ is normal with mean $\tilde{\beta}_{j,k} = \{\sigma_\beta^2 Z_{j,k}^+ + \beta_{j,k-1} + I(k < K)\beta_{j,k+1}\} / \{1 + I(k < K) + S_{j,k}\sigma_\beta^2\}$ and variance $\tilde{\sigma}_{\beta,k}^2 = \sigma_\beta^2 / \{1 + I(k < K) + S_{j,k}\sigma_\beta^2\}$. This step exploits the fact that $\boldsymbol{\beta}$ has a closed-form full conditional distribution under the augmented likelihood, which is not the case using the original likelihood (2.1).

Step 3: Isotonic regression transformation. Apply the Dunson–Neelon algorithm to $\boldsymbol{\pi}(\boldsymbol{\beta}, j)$ as follows. Denote the posterior variance–covariance matrix of $\boldsymbol{\pi}(\boldsymbol{\beta}, j)$ by \mathbf{V}_j and, for each pair of indices $k_1, k_2 \in \{1, \dots, K\}$ with $k_1 \leq k_2$, denote the subvector $\boldsymbol{\pi}(\boldsymbol{\beta}, j)_{[k_1:k_2]} = (\pi(\boldsymbol{\beta}, d_{k_1}, j), \dots, \pi(\boldsymbol{\beta}, d_{k_2}, j))$. Let $\mathbf{V}_{j,[k_1:k_2]}$ be the submatrix of \mathbf{V}_j of dimension $(k_2 - k_1 + 1) \times (k_2 - k_1 + 1)$ corresponding to $\boldsymbol{\pi}(\boldsymbol{\beta}, j)_{[k_1:k_2]}$. The vector obtained by the Dunson–Neelon Bayesian isotonic regression transformation is

$$\tilde{\boldsymbol{\pi}}(\boldsymbol{\beta}, d_k, j) = \min_{k_2 \in U_k} \max_{k_1 \in L_k} \left(\frac{\mathbf{1}'_{k_2-k_1+1} \mathbf{V}_{j,[k_1:k_2]}^{-1} \boldsymbol{\pi}(\boldsymbol{\beta}, j)_{[k_1:k_2]}}{\mathbf{1}'_{k_2-k_1+1} \mathbf{V}_{j,[k_1:k_2]}^{-1} \mathbf{1}_{k_2-k_1+1}} \right),$$

where $L_k = \{s: s \leq k\}$, $U_k = \{s: s \geq k\}$, and $\mathbf{1}_k$ is the k -vector with all entries 1. This transformation ensures that $\tilde{\boldsymbol{\pi}}(\boldsymbol{\beta}, d_1, j) \leq \tilde{\boldsymbol{\pi}}(\boldsymbol{\beta}, d_2, j) \leq \dots \leq \tilde{\boldsymbol{\pi}}(\boldsymbol{\beta}, d_K, j)$ for all j . Despite the apparent complexity of the above notation, this transformation is straightforward to apply in our setting and does not introduce any substantive computational difficulties. In summary, our approach is to first compute the posterior of $\boldsymbol{\pi}(\boldsymbol{\beta}, j)$ without requiring monotonicity in dose (Steps 1 and 2) and then apply the Bayesian isotonic regression algorithm (Step 3) to obtain $\tilde{\boldsymbol{\pi}}(\boldsymbol{\beta}, d_k, j)$, which increases in dose.

4. DECISION RULES

4.1 *Specifying design parameters*

Once a formal definition of what constitutes a DLT has been established, the physicians should be asked to choose t^* as the smallest value that includes possible LO toxicities based on their clinical judgment. For example, if one course of treatment lasts 6 weeks but there is substantial concern about DLTs occurring up to 12 weeks from the start of treatment, then t^* should be set equal to 12 weeks. Given this, they should then choose the corresponding target π^* for the posterior mean of $\Pr(T \leq t^* | d, \boldsymbol{\theta})$.

The decision rules for our method rely on 2 different types of Bayesian criteria, the posterior probabilities $\zeta_k(D_n) = \Pr\{\tilde{\pi}(\boldsymbol{\beta}, d_k) > \pi^* | D_n\}$ for $k = 1, \dots, K$, and predictive probabilities based on approximate values of the ζ_k 's that involve both D_n and future outcomes. To characterize the level of risk associated with each dose, we partition the unit interval into 3 subintervals determined by the probability cutoffs $0 < \underline{\zeta} < \bar{\zeta} < 1$.

DEFINITION Given data D_n , the probability of toxicity at dose d_k is “negligible” if $\zeta_k(D_n) < \underline{\zeta}$, “acceptable” if $\underline{\zeta} \leq \zeta_k(D_n) \leq \bar{\zeta}$, and “excessive” if $\zeta_k(D_n) > \bar{\zeta}$.

Reasonable values for these cutoffs are $0.05 \leq \underline{\zeta} \leq 0.30$ and $0.70 \leq \bar{\zeta} \leq 0.95$, with the particular values chosen to reflect how aggressively or conservatively the physicians wish the algorithm to behave. This decision should be guided by preliminary computer simulations, with $\underline{\zeta}$ and $\bar{\zeta}$ adjusted to obtain a design with desirable properties. In Section 6, below, we will illustrate how one may conduct such a sensitivity analysis in $(\underline{\zeta}, \bar{\zeta})$.

4.2 *Criteria for suspending accrual*

Our goal is to construct a method that protects patient safety by controlling accrual without sacrificing the dose-finding method's reliability. To help motivate how we will address this problem, we first describe the following crude method. When a new patient arrives and a dose must be chosen, one may perform a deterministic look ahead (DLA) by assuming that, of the patients treated but not yet fully evaluated, either (i) all will have toxicity by t^* or (ii) none will have toxicity by t^* . If the next assigned dose is the same under both these extreme assumptions, then the next patient is accrued and treated at that dose. Otherwise, accrual is suspended (Thall *and others*, 1999). The problem with the DLA algorithm is that it does not account for the follow-up times of patients who have not been fully evaluated. For example, if $t^* = 100$ days and 2 patients have been followed without toxicity for 10 days and 90 days, respectively, the DLA algorithm treats them identically.

To address this problem, we define criteria for choosing doses and deciding whether to suspend accrual based on the “PRT”. Our approach may be considered as an extension of the DLA that utilizes predicted outcomes of patients treated but not yet fully evaluated, thus accounting for the times they have gone without toxicity. Let d_k denote the most recently assigned (current) dose. The basic idea is to use predictive probabilities of either negligible or excessive toxicity at the current dose to decide whether to treat the next cohort of patients at d_k , treat them at d_{k+1} , or suspend accrual until subsequent follow-up data lead to a decision to either assign a dose or stop the trial.

Let $p_k(a_k, b_k)$ be the beta(a_k, b_k) random variable defined to have the same mean and variance as $\tilde{\pi}(\boldsymbol{\beta}, d_k)$ under the current posterior. Specifically, $E\{p_k(a_k, b_k)\} = a_k/(a_k + b_k) = E\{\tilde{\pi}(\boldsymbol{\beta}, d_k) | D_n\}$ and $\text{var}\{p_k(a_k, b_k)\} = E(p_k)\{1 - E(p_k)\}/(a_k + b_k + 1) = \text{var}\{\tilde{\pi}(\boldsymbol{\beta}, d_k) | D_n\}$. Thus, $p_k(a_k, b_k)$ is a tractable approximation to $\tilde{\pi}(\boldsymbol{\beta}, d_k)$. Let m_k be the number of patients treated at d_k who have not been fully evaluated, indexed by i_1, \dots, i_{m_k} . For each $r = 1, \dots, m_k$, the available outcome data are $\delta_{i_r} = 0$ and the

follow-up time $T_{i_r}^o$ without toxicity. We define the indicator $W_{i_r} = I(Y_{i_r} \leq C - 1)$ that patient i_r will eventually have toxicity within his/her assessment window. Thus, $\mathbf{W}_k = (W_{i_1}, \dots, W_{i_{m_k}})$ is a vector of future observations, and $S(\mathbf{W}_k) = W_{i_1} + \dots + W_{i_{m_k}}$ is the number of patients among the m_k not yet fully evaluated who will have toxicity by t^* . Let $p_k(\mathbf{W}_k, m_k, a_k, b_k)$ denote the probability following a beta($a_k + S(\mathbf{W}_k), b_k + m_k - S(\mathbf{W}_k)$) distribution. This is an approximation of $\tilde{\pi}(\boldsymbol{\beta}, d_k)$ that is obtained by updating a beta(a_k, b_k) prior with the future toxicity indicators for all patients treated at d_k . The PRT criteria are defined as follows:

$$\text{PN}_k(D_n) = \sum_{\mathbf{w}} I[\Pr\{p_k(\mathbf{w}, m_k, a_k, b_k) > \pi^*\} \leq \underline{\xi}] \Pr(\mathbf{W}_k = \mathbf{w} | D_n), \quad (4.1)$$

and

$$\text{PE}_k(D_n) = \sum_{\mathbf{w}} I[\Pr\{p_k(\mathbf{w}, m_k, a_k, b_k) > \pi^*\} \geq \bar{\xi}] \Pr(\mathbf{W}_k = \mathbf{w} | D_n), \quad (4.2)$$

where the sums in (4.1) and (4.2) are over the 2^{m_k} possible realizations of \mathbf{W}_k . Thus, $\text{PN}_k(D_n)$ and $\text{PE}_k(D_n)$ are approximate predictive probabilities that d_k has negligible or excessive toxicity, respectively, and $\text{PA}_k(D_n) = 1 - \text{PN}_k(D_n) - \text{PE}_k(D_n)$ is the predictive probability that d_k has acceptable toxicity. The PRT criteria will be used during the trial as a basis for deciding whether to suspend accrual or choose a dose for the next cohort. Computing the indicators in the sums in (4.1) and (4.2) is straightforward since they are determined by beta probabilities. Denoting the marginal posterior of $\boldsymbol{\beta}$ by $f(\boldsymbol{\beta} | D_n)$, since $E(W_{i_r} | \boldsymbol{\beta}, D_n) = \tilde{\pi}(\boldsymbol{\beta}, d_k, Y_{i_r}^o)$, the predictive probabilities in the sums in (4.1) and (4.2) may be expressed as

$$\Pr(\mathbf{W}_k = \mathbf{w} | D_n) = \int \prod_{r=1}^{m_k} \{\tilde{\pi}(\boldsymbol{\beta}, d_k, Y_{i_r}^o)\}^{w_r} \{1 - \tilde{\pi}(\boldsymbol{\beta}, d_k, Y_{i_r}^o)\}^{1-w_r} f(\boldsymbol{\beta} | D_n) d\boldsymbol{\beta}, \quad (4.3)$$

which are computed as a by-product of the MCMC algorithm described earlier.

4.3 Rules for trial conduct

The following rules will be used to decide whether or not to suspend accrual and, if accrual is not suspended, determine an appropriate dose for the next cohort of patients. These rules will be applied after the last patient of the current cohort has been enrolled. If the rules indicate that accrual should not be suspended, the investigators should enroll the next cohort of patients and treat them at the dose indicated by the rules. If the rules indicate accrual should be suspended, the investigators should reapply these rules each time a new patient arrives at the clinic or a patient previously not allowed to enroll due to suspended accrual is reconsidered for enrollment. A patient previously not allowed to enroll may be reconsidered whenever the current trial data are updated, which occurs when a toxicity is observed or a currently enrolled patient advances from one interval to the next. From a practical perspective, the investigators may wish to define a maximum waiting time after which a patient would receive an alternative treatment. Reapplication of these rules may result in continued suspension or reopening of accrual. If accrual is reopened, then a new cohort of patients will be enrolled and the rules will be applied again after the last patient of this new cohort has entered the trial.

Index the current dose by k , let n_k be the total number of patients who have received that dose, and recall that m_k is the total number of patients out of the n_k who have not been fully evaluated. Now, define $\widetilde{\text{PN}}_k(D_n)$ and $\widetilde{\text{PE}}_k(D_n)$ such that

$$\widetilde{\text{PN}}_k(D_n) = \begin{cases} 1 & \text{if } m_k = 0 \\ \text{PN}_k(D_n) & \text{if } m_k > 0 \end{cases} \quad \text{and} \quad \widetilde{\text{PE}}_k(D_n) = \begin{cases} 0 & \text{if } m_k = 0, \\ \text{PE}_k(D_n) & \text{if } m_k > 0. \end{cases}$$

That is, $\widetilde{\text{PN}}_k(D_n) = 1$ if all n_k patients have been fully evaluated, and $\widetilde{\text{PN}}_k(D_n) = \text{PN}_k(D_n)$ otherwise. Similarly, $\widetilde{\text{PE}}_k(D_n) = 0$ if all n_k patients have been fully evaluated, and $\widetilde{\text{PE}}_k(D_n) = \text{PE}_k(D_n)$ otherwise. Recall that the current dose d_k has negligible toxicity if $\check{\zeta}_k(D_n) < \underline{\zeta}$, and it is excessively toxic if $\check{\zeta}_k(D_n) > \bar{\zeta}$. The trial is conducted as follows:

- 1) The first cohort of patients are treated at a starting dose chosen by the physicians.
- 2) No untried dose may be skipped when escalating.
- 3) At any point in the trial, if $\check{\zeta}_1(D_n) > \bar{\zeta}$, then stop the trial and conclude that none of the doses are acceptably safe.
- 4) If $\check{\zeta}_k(D_n) > \bar{\zeta}$ and $k > 1$, then de-escalate to the highest dose $k' < k$ such that $\check{\zeta}_{k'}(D_n) \leq \bar{\zeta}$.
- 5) For lower probability cutoff ϵ , if $\underline{\zeta} \leq \check{\zeta}_k(D_n) \leq \bar{\zeta}$ and
 - 5.1) $\widetilde{\text{PE}}_k(D_n) \leq \epsilon$, then treat next cohort at d_k ;
 - 5.2) $\widetilde{\text{PE}}_k(D_n) > \epsilon$, then suspend accrual.
- 6) If $\check{\zeta}_k(D_n) < \underline{\zeta}$ and
 - 6.1) $k = K$ or $\check{\zeta}_{k+1}(D_n) > \bar{\zeta}$, then apply rules (5.1) or (5.2);
 - 6.2) $k < K$, $\widetilde{\text{PN}}_k(D_n) \geq 1 - \epsilon$, and $\widetilde{\text{PE}}_{k+1}(D_n) \leq \epsilon$, then treat next cohort at d_{k+1} ;
 - 6.3) $k < K$, $\widetilde{\text{PN}}_k(D_n) \geq 1 - \epsilon$, and $\widetilde{\text{PE}}_{k+1}(D_n) > \epsilon$, then suspend accrual;
 - 6.4) for $k < K$ and $\widetilde{\text{PN}}_k(D_n) < 1 - \epsilon$, then suspend accrual.
- 7) At the end of the trial, among set of acceptable doses $\{j: \check{\zeta}_j(D_{N_{\max}}) \leq \bar{\zeta}, j = 1, \dots, K\}$, select the dose minimizing $|E\{\tilde{\pi}(\boldsymbol{\beta}, d_j)|D_n\} - \pi^*|$.

Rules (5) and (6) utilize the PRT criteria when the risk of toxicity at the current dose based on the current data is either acceptable or negligible. These rules exploit the fact that predictive probabilities provide information about the risk of future toxicities that cannot be obtained from posterior probabilities alone. In a large sample evaluation of the method's asymptotic "dose-selection" properties, we relax rules (3) and (6.1) because these rules impose a constraint on the number of patients being exposed to unsafe doses, thus potentially preventing sample sizes at these doses from becoming arbitrarily large. Similar to our assumptions, published studies of the asymptotic properties of the CRM and the TITE-CRM do not include early stopping rules (Shen and O'Quigley, 1996; Cheung and Chappell, 2000), even though when used in a clinical trial setting early stopping is appropriate. Similarly, evaluations of the asymptotic properties of an algorithmic competitor to the CRM (the up-and-down design) also assume no early stopping rules (Gezmu and Flournoy, 2006), although, in practice, rules similar to (3) and (6.1) are typically used. Appendix A of the supplementary material available at *Biostatistics* online (<http://www.biostatistics.oxfordjournals.org>) provides consistency results and explains their practical relationship to the method's decision rules.

Computer programs for implementation of the method and running simulations, respectively, can be obtained at the MD Anderson software download site <http://biostatistics.mdanderson.org/SoftwareDownload/>, under the name PRT.

5. GUIDELINES AND ILLUSTRATION

5.1 Guidelines

Because our method is complex, we provide the following guidelines. Initially, the model and design parameters must be established in collaboration with the physicians. These include the definition of toxicity,

anticipated accrual rate, number of discrete-time intervals, C , length of the assessment window t^* , target toxicity probability π^* , maximum sample size N , and decision cutoffs $\underline{\zeta}$, $\bar{\zeta}$, and ϵ . The choice of t^* and C should be guided by practical considerations including the particular disease, agent, and how often patients are evaluated. The target toxicity rate π^* should be chosen to correspond to the event $T \leq t^*$. Because $C = 1$ reduces toxicity to a binary outcome while each interval requires K additional parameters, we recommend $4 \leq C \leq 10$. Although $\pi^* = 0.20$ – 0.30 are commonly used with the CRM or TITE-CRM, targets as high as 0.50 or as low as 0.10 may be appropriate. A rule of thumb is to set $N = 6K$, although resource limitations must also be considered. The design's sensitivity to $(\epsilon, \underline{\zeta}, \bar{\zeta})$ is discussed in Section 6. We found that $\epsilon=0.05$, $0.20 \leq \underline{\zeta} \leq 0.30$, and $0.90 \leq \bar{\zeta} \leq 0.95$ gave reasonable designs with good operating characteristics.

At trial initiation, the first cohort of patients who enter the trial are enrolled immediately upon arriving to the clinic. Subsequently, after the last patient of any cohort has been enrolled, one should perform the computations in Steps 1–3 given in Section 3 and follow the rules for trial conduct in Section 4.3. If trial conduct rule (3) has not stopped the trial, there are two general cases.

Case 1. All currently enrolled patients have been fully evaluated, that is, have either experienced toxicity before t^* or been followed until time t^* without toxicity. Based on $\{\zeta_k(D_n), k = 1, \dots, K\}$, if toxicity at the current dose is negligible, acceptable, or excessive, then “escalate, stay, or de-escalate,” respectively.

Case 2. Some currently enrolled patients have not been fully evaluated, that is, have been followed for less than time t^* and have not had toxicity. Based on $\{\zeta_k(D_n), k = 1, \dots, K\}$ and $\{\widetilde{\text{PN}}_k(D_n), \widetilde{\text{PE}}_k(D_n), \widetilde{\text{PE}}_{k+1}(D_n), k = 1, \dots, K\}$ “escalate, stay, de-escalate, or suspend accrual.”

5.2 Illustration

We illustrate our method with a clinical trial of an experimental radio-sensitizing agent given in combination with radiation to patients with brain cancer or metastases to the brain due to other cancers. A known severe adverse side effect of this treatment is LO neurotoxicity, including cognitive problems, dementia and neuropathy. Typically, patients receiving a radio-sensitizing agent which has not been studied previously in humans will be followed for at least 3 months to assess these LO toxicities.

To illustrate implementation of our method, we consider a case exploring 6 dose levels where doses 1–4 are safe but severe LO neurotoxicity occurs with this agent at dose levels 5 and 6, first appearing roughly 5–6 weeks from the start of therapy. A reasonable approach for the agent and disease described above would be to design a trial with start at dose level 2, using 3 patients per cohort, a maximum sample size of 36 patients, toxicity window $t^* = 90$ days, $\pi^* = 0.30$, $\epsilon = 0.05$, $\underline{\zeta} = 0.30$, and $\bar{\zeta} = 0.90$. Assuming that patients arrive in approximately 6-day intervals (4 per month), an example of how such a phase I trial would proceed is illustrated by the event chart in Figure 1, which shows results for the first 18 patients. The event chart is complemented by Table 1, which provides numerical values of the quantities ζ_k , $\widetilde{\text{PN}}_k(D_n)$, $\widetilde{\text{PE}}_k(D_n)$, and $\widetilde{\text{PE}}_{k+1}(D_n)$ used in the decision rules provided in Section 4.3. Each of the interim times at which a decision is made based on one or more of these quantities is denoted by a vertical arrow on the horizontal axis of Figure 1. The numerical values of these interim decision times also are given in Table 1 along with the corresponding decision criteria values and the decisions made. For example, on day 192 when the current dose is $k = 3$, $\widetilde{\text{PN}}_3(D_n) = 0.999$ which is larger than $(1 - \epsilon) = 0.95$, and $\widetilde{\text{PE}}_4(D_n)$ is set equal to 0 since there is no data at dose level 4. Therefore, according to rule (6.2), we escalate and treat the next cohort at dose 4. In this illustration, the PRT method would have suspended accrual for a total of 210 days but resulted in only 3 toxicities in 36 patients. Figure 1 illustrates the general property of the algorithm that, early in the trial, it tends to suspend accrual for several weeks after each cohort is accrued, although patients still are enrolled at a reasonable rate. Once data at higher doses have been obtained, thereafter accrual tends to remain open. If desired, the method can be modified

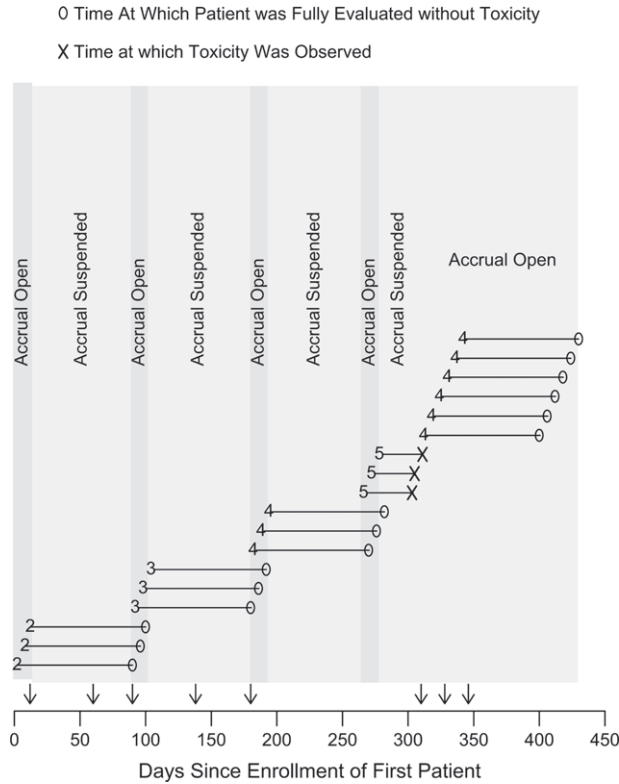


Fig. 1. Illustration of the PRT method for a hypothetical phase I trial. We assume that patients arrive in 6-day intervals. The horizontal axis represents the days since first patient enrolled into the trial. Each horizontal line is the time course of a patient. Numerical values preceding each time course represent the dose administered to that patient. Open circles indicate times at which patients were fully evaluated without toxicity and x’s indicate toxicity times. Darker gray denotes periods when accrual was open. Lighter gray denotes periods when accrual was suspended.

Table 1. Each “time of interim decision” given here corresponds to a time denoted by “↓” on the horizontal (time) axis of Figure 1. For each time point, specific values of the quantities ζ_k , $\text{PN}_k(D_n)$, $\text{PE}_k(D_n)$, and $\text{PE}_{k+1}(D_n)$ are given along with the rule in Section 4.3 triggered, and the decision to suspend accrual (SA), escalate to a higher dose (Esc), stay at the current dose (Stay), or de-escalate (D-Esc). Quantities not relevant to the decision made are denoted by NA

Current dose	Time of interim decision (days from 1st patient enrollment)							
	18	67	90	138	192	312	324	342
	2	2	2	3	3	5	4	4
ζ_k	0.170	0.150	0.020	0.122	0.015	0.997	0.041	0.032
ζ_{k+1}	NA	NA	0.256	NA	0.260	NA	0.998	0.998
$\widetilde{\text{PN}}_k(D_n)$	0.620	0.66	0.988	0.740	0.999	NA	NA	NA
$\widetilde{\text{PE}}_k(D_n)$	NA	NA	NA	NA	NA	NA	0.001	0.001
$\widetilde{\text{PE}}_{k+1}(D_n)$	NA	NA	0.0	NA	0.0	NA	NA	NA
Rule triggered	6.4	6.4	6.2	6.4	6.2	4.0	6.1	6.1
Decision	SA	SA	Esc	SA	Esc	D-Esc	Stay	Stay

to suspend accrual for shorter amounts of time. For example, if we set $\epsilon = 0.20$, accrual suspension can be shortened early in the trial, but at the cost of an increased average number of toxicities. If the method had been implemented with $\epsilon = 1$ (i.e. disallowing any accrual suspension), we would have observed at least 6 toxicities but never suspended accrual, similar to what the TITE-CRM would have done. Alternatively, if $\epsilon = 0$ (i.e. suspend accrual until all patients in a cohort have been fully evaluated for toxicity, as would be the case with a conventional CRM method), we would have suspended accrual for 1080 days and seen 3 toxicities. We believe that our approach provides a reasonable compromise between these two extremes, with an average savings of 1 toxicity for every 70 days of accrual suspension when compared to the design with no accrual suspension and a total of 870 fewer suspended days when compared to the conventional design.

6. SIMULATIONS

6.1 Simulation design

To assess the average behavior of the PRT method, we performed a simulation study under a wide variety of scenarios. We designed the study to address several key issues related to LO toxicities, including the shape of the hazard of toxicity within the window $[0, t^*]$ and the accrual rate, as well as the usual considerations of how the true values of $F_T(t^*|d_k)$ vary with d_k . Since addressing these issues requires a model sufficiently flexible to include a wide range of hazard functions, we simulated toxicity times from a piecewise exponential distribution, described in detail in Appendix B of the supplementary material available at *Biostatistics* online. In all scenarios studied, there were $K = 6$ doses, a maximum of $N = 36$ patients, the first patient was treated at the lowest dose level, and we assumed an accrual rate of 4 patients per month. Since patients are evaluated for toxicity on a weekly or biweekly basis in many oncology settings, for our simulations we chose biweekly assessments and assumed a $t^* = 3$ -month assessment window, with target $\pi^* = 0.30$. Based on our preliminary simulations, we used $\underline{\zeta} = 0.30$, $\bar{\zeta} = 0.90$, and $\epsilon = 0.05$. For the scenarios studied, Figure 2 illustrates the shapes of the hazard functions used in the simulations, in particular the distinction between late- and early-onset toxicities. Scenarios 1–5 have LO toxicities, while toxicities occur early in scenario 6. Since the state-space model is uniquely parameterized by β_{0k} and σ_β , for the prior we set $\beta_{0k} = -14$ and $\sigma_\beta = 28$, which is a sufficiently large variance that the data dominate the method's decisions early in the trial.

In additional preliminary simulations, we assessed the MCMC algorithm's convergence using standard diagnostics. Based on these results, we chose a burn-in of 1000 iterations with a chain of length 1000. Although the posterior sample size is constrained by the need to generate many replications in the simulations, one may use larger posterior sample sizes when used in a clinical trial setting. We simulated 1000 trials for each scenario. The computer program was written using Visual Fortran, and the simulations were carried out on an IBM compatible personal computer with dual 3.06-GHz math coprocessors and 1-GB memory. For each scenario, 1000 simulations were completed in approximately 1 h.

For the TITE-CRM, the model parameters are $\theta = (\hat{p}_1, \dots, \hat{p}_K, \alpha)$ with \hat{p}_k^α , the probability of toxicity by time t^* at dose level k , where α has exponential prior with mean 1 and each \hat{p}_k is fixed with $0 < \hat{p}_1 < \hat{p}_2 < \dots < \hat{p}_K < 1$. The TITE-CRM is implemented with $\pi_w(T_i^o, t^*, d_{(i)}) = w(T_i^o, t^*)\hat{p}_k^\alpha$, where $\{w(T_i^o, t^*), i = 1, \dots, n\}$ are patient-specific weights calculated using the adaptive sample-based weighting scheme given in Display (3) of Cheung and Chappell (2000). For the simulations, we used $(\hat{p}_1, \hat{p}_2, \dots, \hat{p}_6) = (0.05, 0.10, 0.20, 0.30, 0.5, 0.70)$, with each successive dose chosen to minimize $|E(p_k^\alpha | \theta, D_n) - \pi^*|$. To ensure a fair comparison, we modified the TITE-CRM to stop early for excessive toxicity since our method does this. The additional rule is to terminate the trial if $\Pr(\hat{p}_1^\alpha > \pi^* | \theta, D_n) > 0.95$.

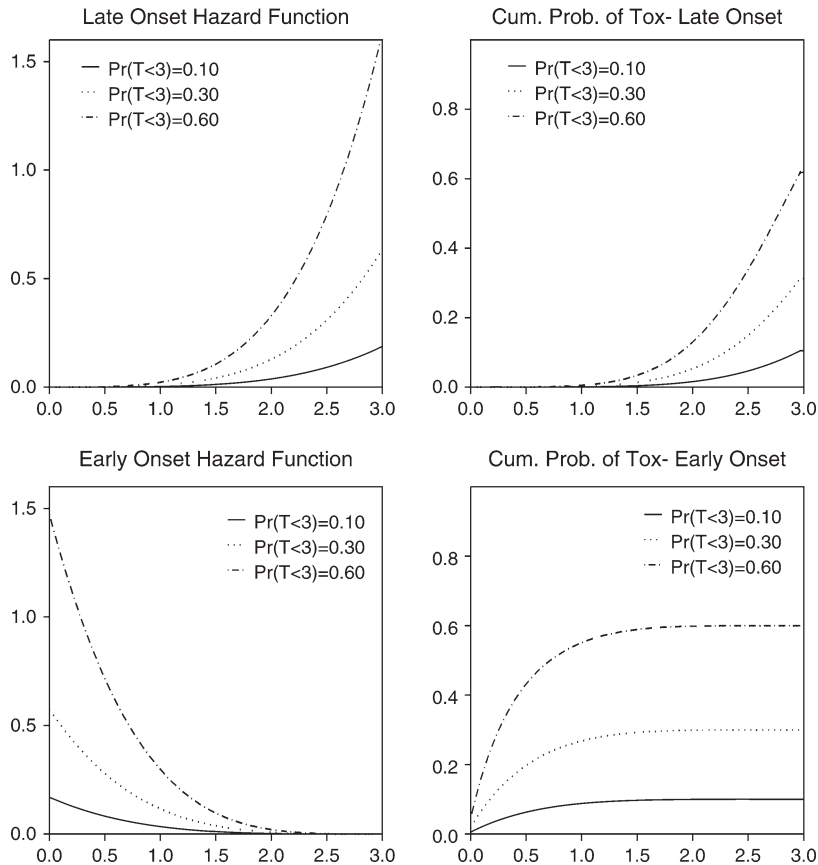


Fig. 2. Hazard functions and cumulative toxicity probabilities of the piecewise exponential distribution used in the simulations.

6.2 Simulation results

Table 2 summarizes the operating characteristics of the PRT method and TITE-CRM, including the selection probability, average number of patients treated, average number of toxicities at each dose, average total number of toxicities, and average overall trial duration. For all 6 scenarios, both methods had comparable correct selection probabilities. While both the average numbers of toxicities and the average numbers of patients treated at toxic doses were lower for the PRT method, the lengths of the trials using the PRT method were typically longer due to the accrual suspension rule. These simulations show that the PRT method has a trade-off between trial safety and trial duration. In addition, we varied the monthly accrual rate from 1 to 6 under scenario 1 (Figure 3) and showed that the percentage of toxicities allowed by the PRT method does not increase with accrual rate, while this is not the case for the TITE-CRM.

Table 3 summarizes a sensitivity analysis in the parameters $\underline{\zeta}$ and $\bar{\zeta}$, under each of scenarios 1, 2, and 3. In this table, elasticity measures the total number of patients “spent” at doses with probability of toxicity more than 10% less than π^* . Based on these simulations, we see that setting $\underline{\zeta} = 0.10$ results in relatively inelastic escalation while setting $\bar{\zeta} = 0.85$ results in low correct selection probability and relatively high early stopping even when all doses are safe under scenario 1. These simulations suggest setting $\underline{\zeta} \in (0.20, 0.30)$ and $\bar{\zeta} \in (0.90, 0.95)$.

Table 2. Operating characteristics of the PRT and TITE-CRM designs

Scenario 1 (late onset)	True Probability($T < 3 \text{ months} d_k$)						None	Total	Duration
	d_1	d_2	d_3	d_4	d_5	d_6			
	0.25	0.35	0.50	0.60	0.68	0.73			
PRT									
% Selected	41	43	4	0	0	0	12	—	1.5
No. of patients	13.1	15.2	4.4	0.5	0.0	0.0		33.2	—
No. of toxicities	3.3	5.3	2.2	0.3	0.0	0.0		11.1	—
TITE-CRM									
% Selected	46	42	8	0	0	0	4	—	1.0
No. of patients	10.6	11.0	7.8	4.0	1.4	0.6		35.5	—
No. of toxicities	2.6	3.8	3.9	2.4	1.0	0.5		14.1	—
<hr/>									
Scenario 2 (late onset)	0.03	0.05	0.10	0.30	0.50	0.60			
PRT									
% Selected	0	1	26	64	9	0	0	—	1.8
No. of patients	3.5	4.8	9.8	12.5	4.6	0.8		36.0	—
No. of toxicities	0.1	0.3	1.0	3.7	2.2	0.5		7.8	—
TITE-CRM									
% Selected	0	0	12	70	19	0	0	—	1.0
No. of patients	3.0	3.2	4.5	11.2	9.4	4.7		36.0	—
No. of toxicities	0.1	0.2	0.4	3.4	4.7	2.8		11.6	—
<hr/>									
Scenario 3 (late onset)	0.50	0.60	0.68	0.73	0.76	0.78			
PRT									
% Selected	11	0	0	0	0	0	89	—	0.8
No. of patients	13.0	5.1	0.5	0.0	0.0	0.0		18.6	—
No. of toxicities	5.6	2.6	0.3	0.0	0.0	0.0		8.6	—
TITE-CRM									
% Selected	18	0	0	0	0	0	82	—	0.7
No. of patients	15.1	5.6	3.7	1.8	0.7	0.3		27.2	—
No. of toxicities	5.7	3.1	2.2	1.2	0.5	0.2		12.8	—
<hr/>									
Scenario 4 (late onset)	0.01	0.02	0.03	0.05	0.50	0.60			
PRT									
% Selected	0	0	0	67	32	1	0	—	1.8
No. of patients	3.1	3.5	4.5	13.1	9.4	2.4		36.0	—
No. of toxicities	0.0	0.1	0.1	0.6	4.7	1.5		7.0	—
TITE-CRM									
% Selected	0	0	1	45	54	0	0	—	1.0
No. of patients	3.0	3.0	3.1	5.9	13.1	7.8		36.0	—
No. of toxicities	0.0	0.1	0.1	0.3	6.6	4.7		11.8	—

Continued...

Table 2. *Continued*

Scenario 5 (late onset)	True Probability($T < 3$ months d_k)						None	Total	Duration
	d_1	d_2	d_3	d_4	d_5	d_6			
	0.05	0.06	0.08	0.11	0.19	0.34			
PRT									
Selected	0	0	5	16	49	30	0	—	1.8
Patients	3.5	4.8	6.2	7.5	7.8	6.2		36.0	—
Toxicities	0.2	0.3	0.5	0.8	1.5	2.1		5.4	—
TITE-CRM									
Selected	0	0	0	13	58	29	0	—	1.0
Patients	3.1	3.3	3.9	7.1	9.9	8.6		36.0	—
Toxicities	0.1	0.2	0.3	0.8	1.9	2.8		6.2	—
Scenario 6 (early onset)	0.250	0.350	0.500	0.600	0.680	0.730			
PRT									
% Selected	44	38	3	0	0	0	14	—	1.0
No. of patients	16.6	12.6	2.7	0.3	0.0			32.1	—
No. of toxicities	4.4	4.7	1.6	0.2	0.0			10.9	—
TITE-CRM									
% Selected	40	41	6	1	0	0	12	—	0.9
No. of patients	17.0	11.4	3.8	0.6	0.0	0.0		32.8	—
No. of toxicities	4.2	4.0	1.9	0.3	0.0	0.0		10.5	—

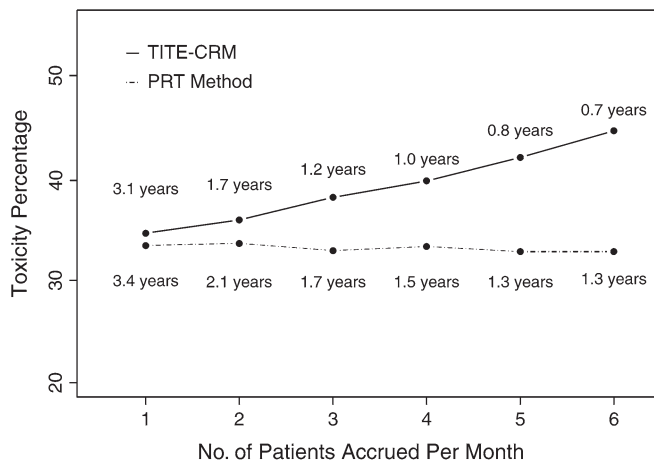


Fig. 3. Comparison of toxicities observed using the PRT method and TITE-CRM under scenario 1 with varying accrual rates. Each toxicity percentage is labeled by the mean trial duration in years.

We also simulated various other scenarios which are presented in Appendix C of the supplementary material available at *Biostatistics* online. These simulations show the method is robust to nonconstant accrual rate, performs worse if instead of suspending accrual patients are assigned to lower doses, and performs worse if the isotonic regression transformation is not used.

Table 3. Sensitivity analyses of the PRT method to $\bar{\xi}$ and $\underline{\xi}$, under scenarios 1, 2, and 3 (S1, S2, S3). % Correct decision denotes the percentage of times the method selects a dose with true toxicity probability $0.25 \leq \Pr(T < 3) \leq 0.35$, within 0.05 of the target $\pi^* = 0.30$ under scenarios 1 and 2, and the early stopping probability under scenario 3. Elasticity denotes the number of patients assigned to doses with true toxicity probability ≤ 0.20

	$\underline{\xi} = 0.10$			$\underline{\xi} = 0.20$			$\underline{\xi} = 0.30$		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
$\bar{\xi} = 0.95$									
% Correct decision	91.8	63.7	75.1	89.1	70.9	75.8	87.3	69.0	78.9
Total no. of toxicities	11.4	5.1	11.9	11.7	7.3	11.2	11.9	8.1	10.4
Elasticity	NA	25.4	NA	NA	19.2	NA	NA	16.6	NA
$\bar{\xi} = 0.90$									
Correct decision	87.4	56.0	86.7	86.0	64.3	87.4	84.0	63.7	88.7
Total no. of toxicities	10.6	4.9	9.0	10.9	6.7	8.6	11.1	7.8	8.6
Elasticity	NA	27.0	NA	NA	21.0	NA	NA	18.5	NA
$\bar{\xi} = 0.85$									
% Correct decision	82.2	52.6	91.6	81.8	60.3	94.2	81.8	59.0	89.6
Total no. of toxicities	10.0	5.3	8.1	10.0	7.2	7.3	10.4	7.6	7.7
Elasticity	NA	26.2	NA	NA	19.8	NA	NA	19.0	NA

6.3 Prior specification

Sensitivity to prior specification is an important concern. In most trials, the only prior information the principal investigator has available is that the probability of toxicity increases with dose, and it is reasonable to assume that the starting dose is relatively safe with low levels of toxicity. Apart from these considerations, physicians are generally most concerned that the method should reliably select the MTD or stop the trial early when all doses are too toxic. Before using the method, it is important to understand how changes in the hyperparameters β_{0k} and σ_β^2 can affect the performance of the method. Increasing β_{0k} results in higher $E\{\tilde{\pi}(\boldsymbol{\beta}, d_k)\}$ (for $k = 1, \dots, K$), the expected prior probabilities of toxicity at each dose, while higher values of σ_β^2 increase the uncertainty in the prior distributions of the $\tilde{\pi}(\boldsymbol{\beta}, d_k)$'s. Preliminary simulations (not provided) showed that setting σ_β^2 between 20 and 40 works well for our method.

Given these provisions, a first step is to calculate $E\{\tilde{\pi}(\boldsymbol{\beta}, d_k)\}$. For our prior, these values were (0.03, 0.18, 0.33, 0.45, 0.53, 0.60) for doses 1 through 6, respectively. As with our parameter settings, we recommend that $E\{\tilde{\pi}(\boldsymbol{\beta}, d_1)\}$ be less than the target π^* . Model parameters may be calibrated by starting with an initial guess of the appropriate prior hyperparameters and then running a small number of simulations (50–100 repetitions) for scenarios in which the lowest dose is too toxic, or has toxicity close to the target, and also a case in which the target toxicity rate is achieved at a relatively high dose. These simulations allow one to assess the effect of the prior on the method's ability to stop appropriately when all doses are too toxic, not when the first dose is close to the target, and escalate appropriately without sticking too rigidly to doses that are not too toxic. In practice, it takes only a few iterations of this process to find a parameterization that works well. As a final step, a few parameter settings with reasonable operating characteristics should be presented to the principal investigators to allow them to choose a prior with which they are comfortable.

One also may carry out an analysis of the method's sensitivity to prior variability by extending the state-space model so that $\beta_{j,k}|\beta_{j,k-1} \sim N(\beta_{j,k-1}, q\sigma_\beta^2)$ for $k = 2, \dots, K$ and $0 < q < 1$. Since $\text{var}(\beta_{j,k}) = \sum_{k=1}^K \sigma_\beta^2 q^{k-1}$, with this prior q may be used as a tuning parameter to adjust the rate at which the variability of the risk of toxicity increases with dose.

7. DISCUSSION

We have proposed a new method for controlling the risk of LO toxicities in a phase I clinical trial by using the PRT at prospective doses to decide whether to temporarily suspend accrual. Our main conclusions are that, if patient accrual is rapid and toxicities occur at the targeted rate in the toxicity evaluation window $[0, t^*]$ but are likely to occur late in this interval, then, on average, (i) the PRT method gives a trial with fewer toxicities but a longer duration compared to the TITE-CRM, (ii) the percentage of toxicities allowed by the PRT method does not increase with accrual rate, while this is not the case for the TITE-CRM, and (iii) in some cases (scenarios 2 and 4 of our study), the TITE-CRM is much more likely than the PRT method to select a final dose having excessive toxicity probability. The additional safety provided by the PRT method compared to the TITE-CRM, in terms of percent reduction in number of toxicities and reduced probability of selecting a final dose that is excessively toxic, is greatest in the most dangerous cases, namely when the probability of toxicity increases sharply between dose levels, when toxicities occur late in the assessment interval, or when accrual is rapid. In the presence of LO toxicities, the lower toxicity rate of the PRT method is due to the fact that it treats fewer patients at excessively toxic doses, which may be attributed to the rules for delaying accrual based on predictive probabilities. In contrast, when toxicities occur early (scenario 5), the two methods perform nearly identically.

Our model and decision rules are complex. A less complex model or method, possibly under a frequentist approach, would be desirable, provided that it maintains the advantages of our method. Implementation of a frequentist version of our method, that is, with accrual suspension, stopping for excessive toxicity, and not escalating to excessively toxic doses, although theoretically possible, would pose certain challenges. One possible likelihood-based approach would employ a penalized time-to-event model to avoid singularities in estimation of model parameters, especially early in the trial before toxicities have been observed. Two possible models that might work in this context are a version of the TITE-CRM model in which the prior is treated as a penalty and a Cox model with Firth's penalty (Heinze and Schemper, 2001). For both models, dose finding could be based on maximum penalized likelihood estimation. In any case, estimation must be constrained to ensure that the probability of toxicity increases with dose. Moreover, the exact sampling distribution of the parameter estimates would be needed, possibly obtained using bootstrap methods, since asymptotic approximations would likely be inaccurate for use in decision making. A difficult challenge in developing such a method would be constructing reasonable frequentist rules for escalating, de-escalating, staying at the same dose, or suspending accrual.

There are some limitations to our method. For example, although in many trial settings toxicity is measured at discrete time points, there are other situations in which toxicity is evaluated continuously. Clearly, for such settings, using our method would be suboptimal due to loss of information. In this context, a model with nonconstant hazard such as a Weibull distribution may be better suited than our method. The main considerations in developing such a method would be model parameterization and prior calibration. Another limitation is that our method is not suited for choosing a best dose from a continuum since it was developed with the goal of selecting a best dose from a finite set of doses.

To address this problem, one could assume a model for the probability of toxicity as a smooth monotone function of dose and estimate dose using standard inverse dose-response methods. If LO toxicity is not a concern, a version of the TITE-CRM based on a logistic regression model, rather than the 1-parameter model given in (1.1), could be used for this purpose.

ACKNOWLEDGMENTS

The authors thank the associate editor, reviewers, and editor for constructive comments that improved the article. *Conflict of Interest*: None declared.

FUNDING

National Institutes of Health (CA-83932, CA-79466).

REFERENCES

- ALBERT, J. H. AND CHIB, S. (1993). Bayesian analysis of binary and polytomous response data. *Journal of the American Statistical Association* **88**, 669–679.
- ALBERT, J. H. AND CHIB, S. (2001). Sequential ordinal modeling with applications to survival data. *Biometrics* **57**, 829–836.
- BABB, J., ROGATKO, A. AND ZACKS, S. (1998). Cancer phase I clinical trials; escalation with overdose control. *Statistics in Medicine* **17**, 1103–1120.
- CHEUNG, Y-K. AND CHAPPELL, R. (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56**, 1177–1182.
- DUNSON, D. B. AND NEELON, B. (2003). Bayesian inference on order-constrained parameters in generalized linear models. *Biometrics* **59**, 286–295.
- GEZMU, M. AND FLOURNOY, N. (2006). Group up-and-down designs for dose-finding. *Journal of Statistical Planning and Inference* **136**, 1749–1764.
- HEINZE, G. AND SCHEMPER, M. (2001). A solution to the problem of monotone likelihood in Cox regression. *Biometrics* **57**, 114–119.
- O’QUIGLEY, J., PEPE, M. AND FISHER, L. (1990). Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* **46**, 33–48.
- SHEN, L. Z. AND O’QUIGLEY, J. (1996). Consistency of continual reassessment method under model misspecification. *Biometrika* **83**, 395–405.
- STORER, B. E. (1989). Design and analysis of phase I clinical trials. *Biometrics* **45**, 925–937.
- THALL, P. F., LEE, J. J., TSENG, C-H. AND ESTEY, E. H. (1999). Accrual strategies for phase I trials with delayed patient outcome. *Statistics in Medicine* **18**, 1155–1169.

[Received May 14, 2007; first revision September 6, 2007; second revision October 18, 2007;
accepted for publication November 1, 2007]