

Dose-Finding Based on Efficacy–Toxicity Trade-Offs

Peter F. Thall* and John D. Cook**

Department of Biostatistics and Applied Mathematics, The University of Texas,
M.D. Anderson Cancer Center, Houston, Texas 77030, U.S.A.

*email: rex@mdanderson.org

**email: cook@mdanderson.org

SUMMARY. We present an adaptive Bayesian method for dose-finding in phase I/II clinical trials based on trade-offs between the probabilities of treatment efficacy and toxicity. The method accommodates either trinary or bivariate binary outcomes, as well as efficacy probabilities that possibly are nonmonotone in dose. Doses are selected for successive patient cohorts based on a set of efficacy–toxicity trade-off contours that partition the two-dimensional outcome probability domain. Priors are established by solving for hyperparameters that optimize the fit of the model to elicited mean outcome probabilities. For trinary outcomes, the new algorithm is compared to the method of Thall and Russell (1998, *Biometrics* 54, 251–264) by application to a trial of rapid treatment for ischemic stroke. The bivariate binary outcome case is illustrated by a trial of graft-versus-host disease treatment in allogeneic bone marrow transplantation. Computer simulations show that, under a wide range of dose-outcome scenarios, the new method has high probabilities of making correct decisions and treats most patients at doses with desirable efficacy–toxicity trade-offs.

KEY WORDS: Adaptive design; Bayesian design; Biologic agents; Dose-finding; Phase I clinical trial; Phase II clinical trial.

1. Introduction

We present an outcome-adaptive, model-based Bayesian procedure that chooses doses of an experimental agent for successive patient cohorts in a clinical trial based on both efficacy (E) and toxicity (T). The method accommodates settings with trinary outcomes where E and T are disjoint and it is possible that neither event may occur, and also trials with bivariate binary outcomes, where the patient may experience both events. Denote the probabilities of E and T by $\pi(x, \theta) = \{\pi_E(x, \theta), \pi_T(x, \theta)\}$, where x denotes dose and θ is the model parameter vector. Each time a dose must be chosen based on the current interim data, \mathcal{D} , from the trial, the method evaluates the desirability of each x by using a family of contours characterizing the trade-off between E and T to define a non-Euclidean distance from $E\{\pi(x, \theta) | \mathcal{D}\}$ to the ideal point $\pi = (1, 0)$. The contours are constructed from target values of π elicited from the physician, similarly to the method of Thall, Sung, and Estey (2002). We establish priors by using mean values of $\pi(x, \theta)$ elicited at each x to obtain the prior mean of θ , and calibrating second-order hyperparameters to obtain a vague prior.

This research was motivated by several practical problems that we encountered in applying the method of Thall and Russell (TR, 1998) to design and conduct various dose-finding trials in oncology (cf. deLima et al., 2001; Couriel et al., 2001) and stroke treatment (Warach et al., 2002). TR require E and T to be disjoint, with trinary outcome $Y \in \{E, T, (E \cup T)^c\}$. They assume that $\pi(x, \theta)$ follows a proportional odds (PO) model (McCullagh, 1989), given by $\text{logit}\{\pi_T(x, \theta)\} =$

$\eta_T(x, \theta) = \mu + x\beta$ and $\text{logit}\{\pi_E(x, \theta) + \pi_T(x, \theta)\} = \eta_{E \cup T}(x, \theta) = \mu + \alpha + x\beta$, where $\theta = (\mu, \alpha, \beta)$ has dimension $p = 3$ with entries following rectangular priors subject to $\alpha > 0$ and $\beta > 0$. Given \mathcal{D} , TR define x to be an acceptable dose if

$$\Pr\{\pi_E(x, \theta) > \underline{\pi}_E | \mathcal{D}\} > p_E \quad (1)$$

and

$$\Pr\{\pi_T(x, \theta) < \bar{\pi}_T | \mathcal{D}\} > p_T, \quad (2)$$

where $\underline{\pi}_E$ and $\bar{\pi}_T$ are fixed lower and upper limits specified by the physician, and p_E and p_T are fixed probability cutoffs. Denote the current set of acceptable doses, given \mathcal{D} , by $\mathcal{A}(\mathcal{D})$. TR define the best dose in $\mathcal{A}(\mathcal{D})$, which we denote by $x^*(\mathcal{D})$, to be that maximizing $\Pr\{\pi_E(x, \theta) > \underline{\pi}_E | \mathcal{D}\}$. Thall, Estey, and Sung (1999) modify this definition so that, if this criterion is very close to the largest value for two or more doses, then the dose among these that maximizes $\Pr\{\pi_T(x, \theta) < \bar{\pi}_T | \mathcal{D}\}$ is chosen as best. Subject to the safety constraint that no untried dose may be skipped when escalating, each new cohort is treated at $x^*(\mathcal{D})$ and, at the end of the trial, $x^*(\mathcal{D})$ is selected for future study.

A major limitation of the TR method is that, in cases where all doses have acceptable toxicity but higher dose levels have substantially higher efficacy, it does not escalate to the more desirable doses with high probability. Consequently, in such settings it is likely to fail to select a higher dose level that is safe and provides greater efficacy. More generally, the TR method may fail to reliably determine the best among several

acceptable doses. For trinary outcomes, we also found that the three-parameter model used by TR, while parsimonious, may be overly restrictive in certain cases. Finally, because it is limited to the case of trinary outcomes, the TR method cannot accommodate settings where the physician wishes to allow the possibility that both E and T may occur.

In the trinary outcome case, we illustrate the new method and compare it to the TR method by application to a trial of thrombolytic agents for the treatment of ischemic stroke. The bivariate binary outcome case is illustrated by a trial of a treatment for steroid-refractory graft-versus-host disease in patients with hematologic malignancies who have undergone allogeneic bone marrow transplantation. While it may seem that extension to the bivariate binary outcome case is straightforward, in fact this involves a substantively different probability model and a different domain for the efficacy–toxicity trade-off contours. Additionally, the two illustrative trials have very different sample sizes and numerical values for their efficacy and toxicity probability targets. Consequently, together they provide a more complete illustration of the new method than could be obtained by studying either application alone. Our computer simulation studies show that, compared to the TR method in the ischemic stroke trial, the new method provides a substantial improvement in terms of both the probabilities of correctly selecting higher doses that are safe and have greater efficacy, and the numbers of patients treated at desirable doses. For both the trinary and the bivariate binary outcome cases, in dose-outcome scenarios where no dose is acceptable due to excessive toxicity or poor efficacy the new proposal has a high early stopping probability. In particular, the new method has the ethically desirable property that it is likely to treat few patients in the trial at unacceptably toxic doses and more patients at safe doses having higher efficacy.

Several other authors have proposed methods for dose-finding based on both efficacy and toxicity. Gooley et al. (1994) were perhaps the first to consider two dose-outcome curves, and they also proposed the now common procedure of using computer simulation as a clinical trial design tool. In the context of an HIV trial, analogous to our trinary outcome case, O’Quigley, Hughes, and Fenton (2001) propose a two-stage dose-finding design, assuming CRM models for $\pi_T(x, \theta)$ and $\pi_{E|T^c}(x, \theta) = \Pr(E|x, \text{No Toxicity}, \theta)$. An acceptable level of toxicity is determined in the first stage, starting with a low toxicity target that later may be increased, and a sequential probability ratio test is used in the second stage to compare null and alternative values of $\Pr(E \text{ and } T|x) = \pi_{E|T^c}(x, \theta)\{1 - \pi_T(x, \theta)\}$. Braun (2002) considers a stem cell transplantation trial with two binary outcomes, toxicity and progression, analogous to our bivariate binary case. Using a three-parameter Bayesian model, he chooses doses by minimizing a Euclidean or non-Euclidean distance from $[E\{\pi_E(x, \theta) | \mathcal{D}\}, E\{\pi_T(x, \theta) | \mathcal{D}\}]$ to a fixed two-dimensional target. The methods proposed here differ from these earlier approaches in terms of both the underlying model and the dose-finding algorithm. In each case, trinary or bivariate binary outcomes, we assume a more flexible model having more parameters than the models used in the previous approaches noted above. Moreover, we provide an algorithm for establishing priors based on elicited mean outcome prob-

abilities. The greatest difference between our approach and these earlier methods is that our dose-finding algorithm is based on explicit trade-offs between π_E and π_T .

2. Illustrative Trials

2.1 Rapid Treatment of Acute Ischemic Stroke

In the trinary outcome case, we illustrate and compare the models and methods with a trial, sponsored by the National Institute of Neurological Diseases and Stroke, for rapid treatment of acute ischemic stroke (Warach et al., 2002). Ischemic stroke is the third leading cause of death in the U.S. and the leading cause of disability among the elderly. Each patient receives a fixed dose of abciximab (0.25 mg/kg as a bolus followed by 0.125 $\mu\text{g}/\text{kg}/\text{minute}$ for 12 hours) followed by one of the five doses $\{0.0, 2.5, 5.0, 7.5, 10.0\}$ units (U) of reteplase. The rationale for this combination arises from the fact that the two major components of a blood clot, the cause of ischemic stroke, are platelets and fibrin. To dissolve the patient’s clot, abciximab acts against platelet aggregation, while reteplase, a tissue plasminogen activator, increases the ability of vascular tissues to reestablish blood flow (reperfusion) by degrading fibrin. Patient outcome is evaluated by both physical examination and magnetic resonance imaging. Toxicity is defined as symptomatic intracranial hemorrhage, other severe regimen-related adverse event, or death within 48 hours. Response (efficacy) is defined as reperfusion at 24 hours, without toxicity. Up to $N = 72$ patients will be treated in cohorts of size $c = 3$. Initially, the trial was designed and conduct was begun using the TR method with $\pi_E = 0.50$, $\bar{\pi}_T = 0.10$, and $p_E = p_T = 0.10$ for the criteria (1) and (2). Based on computer simulations validating the trade-off-based method, the stroke trial’s Data Safety Monitoring Board approved switching to the new method for the remainder of the trial. Specifics of this implementation will be given in Section 6 along with the simulation results.

2.2 Treating Graft-versus-Host Disease

An inherent risk in allogeneic blood or marrow stem cell transplantation (allotx) is graft-versus-host disease (GVHD), wherein the transplanted cells from an HLA-matched donor engraft and repopulate the patient’s bone marrow, but the engrafted cells attack the patient’s organs in an autoimmune reaction. Consequently, GVHD prophylaxis is essential in allotx. When GVHD occurs and cannot be brought into remission by conventional steroid therapy, other measures must be taken to save the patient. In an ongoing trial of Pentostatin, a new agent for the treatment of steroid-refractory GVHD, up to $N = 36$ patients are treated in cohorts of size $c = 3$. The scientific goal is to find the best dose among $\{0.25, 0.50, 0.75, 1.00\}$ mg/m². For dose-finding, within a 2-week evaluation period, toxicity is defined as infection that cannot be resolved by antibiotics, or death, and response (efficacy) is defined as a decrease in the GVHD severity level by at least one grade. In particular, a patient who is alive with infection but a decreased GVHD severity has both toxicity and response, $T \cap E$. The trial currently is being conducted using the trade-off-based algorithm described below in Section 4, under the bivariate binary dose-outcome model described below in Section 3.

3. Dose-Outcome Models

Given the J doses d_1, \dots, d_J to be considered in the trial, we code dose as $x_j = \log(d_j) - J^{-1} \sum_{k=1}^J \log(d_k)$ for use in the regression models underlying the dose-finding method. In the ischemic stroke trial, where $0 = d_1 < d_2$, we first added d_2 to each d_j before taking logs. For trinary outcomes, a more flexible but still tractable alternative to the PO model is the continuation ratio (CR) model (cf. McCullagh and Nelder, 1989). We formulate a CR model in this setting by defining $\pi_T(x, \theta)$ as in Section 1, but modeling $\eta_{E|T^c}(x, \theta) = \text{logit}\{\pi_{E|T^c}(x, \theta)\} = \text{logit}\{\text{Pr}(Y = E | Y \neq T, x, \theta)\}$ rather than $\eta_{E \cup T}(x, \theta)$. Specifically, we assume $\eta_{E|T^c}(x, \theta) = \mu_E + x\beta_E$. Thus, $p = 4$ and $\theta = (\mu_T, \beta_T, \mu_E, \beta_E)$, subject to the constraints $\beta_T > 0$ and $\beta_E > 0$. To see that this is a CR model, regard $(E \cup T)^c < E < T$ as ordinal levels of Y and observe that $\eta_T = \log\{\pi_T/\text{Pr}(Y < T)\}$ and $\eta_{E|T^c} = \log\{\pi_E/\text{Pr}(Y < E)\}$. More generally, for either the PO or CR model, the logit may be replaced by any link function $g(\pi) = \eta$, where $g = F^{-1}$ for a continuous monotone cdf F .

For clinical settings where E and T may both occur, let $\mathbf{Y} = (Y_E, Y_T)$ be the indicators of E and T . Denote $\pi_{a,b}(x, \theta) = \text{Pr}(Y_E = a, Y_T = b | x, \theta)$ for $a, b \in \{0, 1\}$. There are numerous bivariate binary regression models (McCullagh, 1989; Gloneck and McCullagh, 1995; Joe, 1997). To facilitate model interpretation and prior elicitation, and ensure tractability, we formulate the model in terms of the marginal probabilities, $\pi_T(x, \theta) = g^{-1}\{\eta_T(x, \theta)\}$ and $\pi_E(x, \theta) = g^{-1}\{\eta_E(x, \theta)\}$, and one association parameter, ψ . For toxicity, we assume $\eta_T(x, \theta) = \mu_T + x\beta_T$, with $\beta_T > 0$ if it is known that $\pi_T(x, \theta) \uparrow$ in x , and real-valued β_T otherwise. For efficacy, to allow a wide variety of possible dose-response relationships, including nonmonotone functions, we use the flexible quadratic form $\eta_E(x, \theta) = \mu_E + x\beta_{E,1} + x^2\beta_{E,2}$. This allows the method to be used for trials of biologic agents, where $\pi_E(x, \theta)$ may initially increase in x and then vary very little or possibly decrease for higher doses. Thus, $\theta = (\mu_T, \beta_T, \mu_E, \beta_{E,1}, \beta_{E,2}, \psi)$ and $p = 6$. For simplicity, temporarily suppress (x, θ) . Any bivariate binary probability model characterized by $\{\pi_E, \pi_T, \pi_{1,1}\}$ must satisfy the consistency constraint $\max\{0, \pi_E + \pi_T - 1\} \leq \pi_{1,1} \leq \min\{\pi_E, \pi_T\}$. We use the model given by

$$\pi_{a,b} = (\pi_E)^a(1 - \pi_E)^{1-a}(\pi_T)^b(1 - \pi_T)^{1-b} + (-1)^{a+b}\pi_E(1 - \pi_E)\pi_T(1 - \pi_T)\left(\frac{e^\psi - 1}{e^\psi + 1}\right), \quad (3)$$

for $a, b \in \{0, 1\}$ and real-valued ψ . Aside from the forms of $\eta_E(x, \theta)$ and $\eta_T(x, \theta)$, this has been called a ‘‘Gumbel’’ (Murtaugh and Fisher, 1990) or a ‘‘Morgenstern’’ distribution.

The likelihood for a single patient treated at dose x is

$$\mathcal{L}(Y, x | \theta) = \prod_{y=E, T, (E \cup T)^c} \{\pi_y(x, \theta)\}^{I(Y=y)}$$

in the trinary outcome case and

$$\mathcal{L}(\mathbf{Y}, x | \theta) = \prod_{a=0}^1 \prod_{b=0}^1 \{\pi_{a,b}(x, \theta)\}^{I(Y=(a,b))}$$

for bivariate binary outcomes. In general, denoting the data for the first n patients in the trial by \mathcal{D}_n for $1 \leq n \leq N$, the likelihood is

$$\mathcal{L}_n(\mathcal{D}_n | \theta) = \prod_{i=1}^n \mathcal{L}(\mathbf{Y}_i, x_{(i)} | \theta), \quad (4)$$

where \mathbf{Y}_i and $x_{(i)}$ denote the i th patient’s outcome and dose. In Section 5, we describe a method for specifying a Gaussian prior on θ , under any model of the type discussed here, that does not contain artificial information. This ensures that the prior is uninformative in the sense that the dose-finding algorithm’s behavior is dominated by the data.

4. A Dose-Finding Algorithm Based on Trade-Offs

4.1 *Efficacy-Toxicity Trade-Off Contours*

The criteria given in (1) and (2) are motivated by the desire to limit the risk of treating patients at a dose with either unacceptably high toxicity or unacceptably low efficacy. A different way to address these concerns is to consider the desirability of the pair $\pi(x, \theta)$ in the two-dimensional domain, $\mathbf{\Pi}$, of possible values of π . For bivariate binary outcomes, $\mathbf{\Pi} = [0, 1]^2$, the unit square. In the trinary outcome case, $\mathbf{\Pi}$ is the lower left triangular subset of $[0, 1]^2$ where $\pi_E + \pi_T \leq 1$. We first construct a target efficacy-toxicity trade-off contour, \mathcal{C} , in $\mathbf{\Pi}$ by fitting a curve to target values of π elicited from the physician. The target contour is then used to construct a family of trade-off contours such that all π on the same contour are equally desirable. Because the family of contours partitions $\mathbf{\Pi}$, this construction provides a basis for comparing doses in terms of their posterior means, $E\{\pi(x, \theta) | \mathcal{D}\}$. This will play a central role in the dose-finding algorithm.

To construct \mathcal{C} , we first elicit three target values, $\{\pi_1^*, \pi_2^*, \pi_3^*\}$, that the physician considers equally desirable. The following procedure works well, although other approaches are possible (Figure 1). First, elicit a desirable trade-off target, $\pi_1^* = (\pi_{1,E}^*, \pi_{1,T}^*) = (\pi_{1,E}^*, 0)$, in the case where toxicity has probability 0. That is, elicit the smallest efficacy probability, $\pi_{1,E}^*$, that the physician would consider desirable if toxicity were impossible. Next, elicit π_2^* having the same desirability as π_1^* by asking the physician what the maximum value of π_T may be if (a) in the bivariate binary outcome case, $\pi_E = 1$ or (b) in the trinary outcome case, $\pi_E = 1 - \pi_T$, equivalently, $\pi_{E|T^c} = 1$. Given these two equally desirable extremes, elicit a third pair, π_3^* , that is equally desirable but is intermediate between π_1^* and π_2^* . Plot each target as it is elicited and draw the target efficacy-toxicity trade-off contour, \mathcal{C} , determined by $\{\pi_1^*, \pi_2^*, \pi_3^*\}$. One may obtain \mathcal{C} by fitting a continuous, strictly increasing function, $\pi_T = f(\pi_E)$ to $\{\pi_1^*, \pi_2^*, \pi_3^*\}$ and defining $\mathcal{C} = \{(\pi_E, \pi_T) : 0 \leq \pi_E \leq 1, \pi_T = f(\pi_E)\} \cap \mathbf{\Pi}$, as illustrated in Figure 1. Any tractable function $\pi_T = f(\pi_E)$ may be used, provided that it increases continuously over the domain $\pi_{1,E}^* \leq \pi_E \leq 1 - \pi_T$ in the trinary case, or $\pi_{1,E}^* \leq \pi_E \leq 1$ in the bivariate binary case. In practice, the elicitation process is iterative, and the physician usually modifies the target points on the basis of the plotted contour.

Once \mathcal{C} is established, we use it to define the desirability of any pair of probabilities $\mathbf{q} = (q_E, q_T)$ in $\mathbf{\Pi}$, as follows. Again,




Figure 1. Efficacy–toxicity trade-off contours for the Pentostatin trial. The target contour \mathcal{C} is given by the solid line, and the three elicited target points that determine \mathcal{C} are given by round dots. The two triangular points illustrate the homotopy $h_{0.75}(\mathbf{p}) = \mathbf{q}$, and all points on the contour containing \mathbf{q} have desirability $\delta = 0.75$.

refer to Figure 1. Draw the straight line, $L(\mathbf{q})$, from \mathbf{q} to $(1,0)$, and find the point \mathbf{p} where $L(\mathbf{q})$ intersects \mathcal{C} . Calculate the Euclidean distances $\rho(\mathbf{p})$ from \mathbf{p} to $(1,0)$, and $\rho(\mathbf{q})$ from \mathbf{q} to $(1,0)$. To reflect the fact that values of \mathbf{q} closer to $(1,0)$ are more desirable, we define the desirability of \mathbf{q} to be $\delta(\mathbf{q}) = z - 1$, where $z = \rho(\mathbf{p})/\rho(\mathbf{q})$. We subtract 1 from z to standardize the desirability of points on the elicited contour, $\mathcal{C} = \mathcal{C}_1$, to equal 0. Thus, $\delta(\mathbf{q}) > 0$ for all \mathbf{q} inside the region in Π between \mathcal{C} and $(1,0)$ where the probability pairs are more desirable than those on \mathcal{C} , and $\delta(\mathbf{q}) < 0$ for all \mathbf{q} outside the region, where the probability pairs are less desirable than those on \mathcal{C} . The following definition exploits this structure to induce an ordering on the set of doses.

DEFINITION: Given \mathcal{D} and x , the desirability, $\delta(x, \mathcal{D})$, of x is the desirability of the posterior mean $E\{\boldsymbol{\pi}(x, \boldsymbol{\theta}) | \mathcal{D}\}$.

To apply this during the trial, after the most recent cohort’s data have been incorporated into \mathcal{D} , for each x , $\mathbf{q} = E\{\boldsymbol{\pi}(x, \boldsymbol{\theta}) | \mathcal{D}\}$ is first computed, then $\mathbf{p} = L(\mathbf{q}) \cap \mathcal{C}$ is obtained algebraically, and the desirability of x is then given by $\delta(x, \mathcal{D}) = \rho(\mathbf{p})/\rho(\mathbf{q}) - 1$. Among the doses with acceptable efficacy and toxicity, the dose that maximizes $\delta(x, \mathcal{D})$ is selected.

This may be used to construct a family of trade-off contours that partition Π , with the points on each contour equally desirable. Given any $\mathbf{p} \in \mathcal{C}$ and $z > 0$, we define the homotopy $h_z(\mathbf{p}) = \mathbf{q}$ if $\mathbf{q} \in L(\mathbf{p})$ and $\rho(\mathbf{p})/\rho(\mathbf{q}) = z$. This implies that $h_z(\mathbf{p}) = (1 - (1 - p_E)/z, p_T/z)$. While the range of h_z may not be a subset of Π for some z , since we are only interested in contours inside Π we define $\mathcal{C}_z = \{h_z(\mathcal{C})\} \cap \Pi$. Thus, \mathcal{C}_z is the contour in Π obtained by shifting each $\mathbf{p} \in \mathcal{C}$ along $L(\mathbf{p})$ to the point \mathbf{q} in Π such that $\rho(\mathbf{p})/\rho(\mathbf{q}) = z$. Since the contours are ordered by their desirabilities and the set $\{\mathcal{C}_z : z > 0\}$ of all contours is a partition of Π , the contours induce an ordering on Π .

We require $f(\pi_E)$ to be strictly increasing to ensure that, given $(\pi_E, \pi_T) \in \mathcal{C}$ and $\epsilon > 0$, provided that $(\pi_E + \epsilon, \pi_T) \in \Pi$, it must be the case that $(\pi_E + \epsilon, \pi_T)$ is on a contour below \mathcal{C} and hence is more desirable than (π_E, π_T) . Similarly, a pair $(\pi_E, \pi_T + \epsilon) \in \Pi$ must be on a contour above \mathcal{C} and hence less desirable than (π_E, π_T) . In particular, e.g., in the bivariate binary outcome case, the rectangular set comprised of the line segments from $(\pi_E, \bar{\pi}_T)$ to $(\pi_E, 0)$ and from $(\pi_E, \bar{\pi}_T)$ to $(1, \bar{\pi}_T)$ does not satisfy this admissibility criterion. We used the convenient form $\pi_T = f(\pi_E) = a + b/\pi_E + c/\pi_E^2$ in the applications described here, fit to the three elicited target pairs subject to the constraint that f be nondecreasing for π_E such that $\{\pi_E, f(\pi_E)\} \in \mathcal{C}$. Other functions, such as an elliptical contour, should work as well.

4.2 The Trade-Off-Based Algorithm

Initially, the physician must provide a set of doses, a starting dose for the first cohort, N , c , and the limits π_E and $\bar{\pi}_T$ used in the acceptability criteria (1) and (2). The trade-off targets $\{\pi_1^*, \pi_2^*, \pi_3^*\}$ then must be elicited in order to construct \mathcal{C} and the family of trade-off contours. The probability cut-offs p_E and p_T in (1) and (2) then may be determined, using preliminary computer simulation results, to obtain a design with desirable operating characteristics. Given this structure, the dose-finding algorithm proceeds as follows:

- (1) Treat the first cohort at the starting dose specified by the physician.
- (2) For each cohort after the first, $x \in \mathcal{A}(\mathcal{D})$ if x satisfies both (1) and (2), or if x is the lowest untried dose above the starting dose and it satisfies (2).
- (3) If $\mathcal{A}(\mathcal{D}) \neq \phi$, then the next cohort is treated at the most desirable $x \in \mathcal{A}(\mathcal{D})$, subject to the constraint that no untried dose may be skipped when escalating.
- (4) If $\mathcal{A}(\mathcal{D}) = \phi$, then the trial is terminated and no dose is selected.
- (5) If the trial is not stopped early and $\mathcal{A}(\mathcal{D}_N) \neq \phi$ at the end of the trial, then the dose $x \in \mathcal{A}(\mathcal{D}_N)$ maximizing $\delta(x, \mathcal{D}_N)$ is selected.

The trade-off-based algorithm differs from that given by TR (1998) and extended by Thall et al. (1999) in two essential ways. The first difference, which may appear subtle but in fact has a large substantive effect, is that the lowest untried dose is considered to be acceptable if it has acceptable toxicity, but the efficacy requirement (1) is not imposed. This says that, if the predicted safety of the lowest untried dose is acceptable, then it may be used to treat patients regardless of its predicted efficacy. The effect of this modification

is that the new algorithm is more likely to treat patients at higher, untried doses, provided that patient safety is still protected, and consequently it is more likely to discover a higher dose having superior efficacy. The second difference is that $\mathcal{A}(\mathcal{D})$ is now ordered in terms of the desirabilities, $\{\delta(x, \mathcal{D}) : x \in \mathcal{A}(\mathcal{D})\}$, with these values defined for each x in terms of the distance of the pair $E\{\pi(x, \theta) | \mathcal{D}\}$ to $(1, 0)$ based on the family of contours generated from the elicited target points. Because $\delta(x, \mathcal{D})$ is much more sensitive to differences between elements of $\mathcal{A}(\mathcal{D})$ than are the posterior probabilities $\Pr\{\pi_E(x, \theta) > \bar{\pi}_E | \mathcal{D}\}$ and $\Pr\{\pi_T(x, \theta) < \bar{\pi}_T | \mathcal{D}\}$, and moreover $\delta(x, \mathcal{D})$ reflects the elicited trade-offs between the probabilities of E and T , this new trade-off-based criterion provides a more reliable basis for choosing $x^*(\mathcal{D})$. Our simulations, described below, show that, under a wide variety of dose-outcome scenarios, the new algorithm is more likely to make correct decisions and, on average, it treats more patients at the more desirable doses.

5. Numerical Methods

5.1 Establishing Priors

In clinical trials with model-based Bayesian adaptive decision making, the prior must give a reasonable representation of the physician’s uncertainty, provide a reliable basis for sensible decisions early in the trial, but be sufficiently vague so that the accumulating data dominate the posterior and hence the decisions as the trial progresses. Except for the original PO model with rectangular prior, for all models considered here we will assume each component θ_l of θ is normally distributed with mean $\tilde{\mu}_l$ and standard deviation (SD) $\tilde{\sigma}_l$, denoted $\theta_l \sim N(\tilde{\mu}_l, \tilde{\sigma}_l)$. Let $\xi = (\tilde{\mu}_1, \tilde{\sigma}_1, \tilde{\mu}_2, \tilde{\sigma}_2, \dots, \tilde{\mu}_p, \tilde{\sigma}_p)$ denote the $2p$ -vector of hyperparameters, with all prior covariances set equal to 0, and let $\phi_p(\theta | \xi)$ denote the p -variate normal prior of θ . For each dose $x_j, j = 1, \dots, J$, and outcome $y = E, T$, let $m_{y,j}(\xi)$ and $s_{y,j}(\xi)$ denote the prior mean and SD of $\pi_y(x_j, \theta)$. These are given explicitly by

$$m_{y,j}(\xi) = \int \pi_y(x_j, \theta) \prod_{\ell=1}^p \frac{1}{\sqrt{2\pi\tilde{\sigma}_\ell}} \exp\left\{-\frac{(\theta_\ell - \tilde{\mu}_\ell)^2}{2\tilde{\sigma}_\ell^2}\right\} d\theta$$

and

$$s_{y,j}^2(\xi) = \int \{\pi_y(x_j, \theta) - m_{y,j}(\xi)\}^2 \times \prod_{\ell=1}^p \frac{1}{\sqrt{2\pi\tilde{\sigma}_\ell}} \exp\left\{-\frac{(\theta_\ell - \tilde{\mu}_\ell)^2}{2\tilde{\sigma}_\ell^2}\right\} d\theta.$$

Initially, for each x_j , we elicit the means of $\pi_E(x_j, \theta)$ and $\pi_T(x_j, \theta)$, denoted by $\hat{m}_{E,j}$ and $\hat{m}_{T,j}$, that the physician expects a priori at that dose. In practice, it is easiest to elicit these prior values at the same time one elicits the targets $\{\pi_1^*, \pi_2^*, \pi_3^*\}$ and the other design parameters. We also specify values of $\hat{s}_{E,j}$ and $\hat{s}_{T,j}$ in the range 0.29–0.50, corresponding to beta distributions with parameters having sum at most 2. That is, we allow the marginal prior of each $\pi_y(x_j, \theta)$ to have roughly as much information as at most two data points. Given $\{\hat{m}_{E,j}, \hat{m}_{T,j}, \hat{s}_{E,j}, \hat{s}_{T,j}, j = 1, \dots, J\}$, we numer-

ically solve for the value of ξ that best fits the target means and variances by minimizing the objective function

$$h(\xi) = \sum_{y=E,T} \sum_{1 \leq j \leq J} [\{m_{y,j}(\xi) - \hat{m}_{y,j}\}^2 + \{s_{y,j}(\xi) - \hat{s}_{y,j}\}^2] + c \sum_{1 \leq j < k \leq J} (\tilde{\sigma}_j - \tilde{\sigma}_k)^2. \tag{5}$$

The second term in $h(\xi)$ is included so that the solution will distribute the prior variance more evenly among the components of θ , with c a small positive constant, for example, $c = 0.15$. We intentionally elicit more pieces of prior information than hyperparameters, with $4J > 2p$, and obtain ξ as the least-squares solution to the $4J$ equations in $2p$ unknowns. If $p = 6$ and $J \leq 3$, then the additional prior values can be obtained for an intermediate dose between the J doses in the trial, or by eliciting values from more than one physician. However, the case of $J = 2$ dose levels, while not degenerate, may constitute an impractical use of resources, and we recommend in general that at least three dose levels be investigated. We minimize $h(\xi)$ using the Nelder–Mead algorithm (Nelder and Mead, 1965). In practice, h is modified so that an arbitrarily large value is returned for negative values of σ_j so that, if the algorithm ventures into a region in which a hyperparameter is invalid, it will simply turn back. Since the Nelder–Mead algorithm is not derivative based, it will not overreact to extremely large values of h .

5.2 Posterior Computation

To compute posteriors, we numerically integrate $\mathcal{L}_n(\mathcal{D}_n | \theta)\phi_p(\theta | \xi)$ with respect to θ using the method of Monahan and Genz (1997). One of the crucial steps in this method is to find the mode of the integrand, which requires an initial guess as to where this mode may be. For the first integration, to obtain the posterior based on the data from the first cohort in the trial, we use the mode of the prior distribution as the initial guess of the posterior mode. Subsequently, as each successive cohort’s data are observed, we use the posterior mode from the previous integration as the starting point in the search for the next mode. This works well in the dose-finding setting considered here because the mode does not move very far with the addition of the data from a single cohort.

6. Simulation Studies

To evaluate the new method and also compare it to the TR method in the trinary outcome case, we simulated the ischemic stroke trial using three combinations of algorithm and model. These were (TR, PO) = the TR algorithm under the PO model, (trade-off, PO) = the trade-off-based algorithm under the PO model, and (trade-off, CR) = the trade-off-based algorithm under the CR model. We include (trade-off, PO) as an intermediate case to help assess how much of the difference in performance between (trade-off, CR) and (TR, PO) is due to the new algorithm and how much is due to using the four-parameter CR model rather than the three-parameter PO model. Both algorithms were implemented with $\bar{\pi}_E = 0.50, \bar{\pi}_T = 0.10, p_E = p_T = 0.10$, starting at dose level 1, with $c = 3$ and $N = 72$. The trade-off contour \mathcal{C} for the new algorithm was determined by fitting the curve

$\pi_T = a + b/\pi_E + c/\pi_E^2$ to the three elicited trade-off points $(\pi_E, \pi_T) = (0.45, 0), (0.55, 0.10), (0.84, 0.16)$ that the physician considered equally desirable, and the family of contours was generated as described in Section 4.1. The rectangular prior for the PO model had domain $[-2.982, 2.756]$ for μ , $[-16.249, 13.963]$ for α , and $[22.120, 12.061]$ for β . The mean and SD of each CR model parameter’s Gaussian prior was $(-1.966, 1.791)$ for μ_T , $(1.05925, 1.79113)$ for β_T , $(0.464, 0.332)$ for μ_E , and $(0.968, 0.333)$ for β_E . The (trade-off, CR) design is that currently being used to conduct the ischemic stroke trial. Each case was simulated 5000 times.

We summarize simulation results under six dose-outcome scenarios, chosen to represent a larger set that we examined. Each scenario is characterized by fixed probabilities (π_E, π_T) at each dose. Table 1 gives the operating characteristics of the methods in terms of the selection percentage and number of patients treated for each dose under each dose-outcome scenario. The selection percentages by dose level for (trade-off, CR) are plotted in Figure 2, which also provides a graphical illustration of the scenarios relative to the elicited trade-off contour. The advantage of the new algorithm, in terms of both

selection percentage and numbers of patients treated at desirable doses, is clearly illustrated by Scenario 1, where toxicity is negligible for all doses but efficacy increases substantially with dose. This case also shows that part of the improvement is due to the new algorithm and part to the CR model. Scenario 2 may be considered the opposite case, since the response rate is high and varies only slightly between doses, but toxicity increases rapidly with dose. Here the three (algorithm, model) combinations have very similar behavior. Under Scenario 4, where all doses are safe and desirable, but toxicity increases only slightly while efficacy increases substantially with dose, (trade-off, CR) has greatly superior behavior. When no dose is acceptable, under Scenarios 5 and 6, all three pairs stop the trial early with high probability, although (trade-off, CR) has a slight advantage in both cases.

Scenario 3 is a difficult case in that intermediate dose levels 3 and 4 are most desirable while level 5 has $\pi_T = 0.20$, slightly above the limit $\bar{\pi}_T^* = 0.10$. Here, (trade-off, PO) performs best, with (trade-off, CR) behaving more aggressively by choosing dose level 5 slightly more often. To investigate the surprising result that the three-parameter PO model




Figure 2. Decision percentages for the ischemic stroke trial using the trade-off-based algorithm, under the continuation ratio model, for each of the dose-outcome scenarios given in Table 1. The doses of reteplase are labeled 1, 2, 3, 4, and 5. The area of each disc equals the probability that the dose having (π_E, π_T) located as its center is selected. The darker shaded area in the lower right portion of each graph is set of (π_E, π_T) that are more desirable than the pairs on the elicited target contour.

Table 1

Selection percentage, followed in parentheses by the number of patients treated, for each dose in the ischemic stroke trial, under three combinations of either the Thall–Russell (TR) or trade-off-based algorithm and either the proportional odds (PO) or continuation-ratio (CR) model

Algorithm	Model	Dose (units) of reteplase					None
		0.0	2.5	5.0	7.5	10.0	
		<i>Scenario 1</i> (π_E, π_T)					
	$\delta =$	(0.05, 0.01) -0.74	(0.20, 0.02) -0.48	(0.35, 0.03) -0.22	(0.60, 0.04) 0.22	(0.80, 0.05) 0.54	
TR	PO	0.0 (3.0)	0.0 (3.0)	6.0 (2.6)	52.9 (23.3)	9.0 (15.0)	32.1
Trade-off	PO	0.0 (3.0)	0.0 (3.4)	0.2 (5.3)	22.5 (23.1)	74.2 (35.3)	3.1
Trade-off	CR	0.0 (3.0)	0.0 (4.0)	0.7 (5.1)	5.8 (7.1)	92.8 (52.2)	0.7
		<i>Scenario 2</i> (π_E, π_T)					
	$\delta =$	(0.57, 0.01) 0.21	(0.58, 0.03) 0.20	(0.60, 0.06) 0.18	(0.62, 0.20) -0.31	(0.64, 0.32) -1.00	
TR	PO	0.3 (3.7)	18.2 (19.0)	65.5 (34.3)	14.0 (13.3)	0.1 (0.7)	1.9
Trade-off	PO	0.7 (6.2)	10.8 (18.3)	72.2 (35.6)	15.2 (10.5)	0.6 (1.1)	0.5
Trade-off	CR	0.1 (5.4)	20.5 (21.6)	61.9 (29.5)	16.1 (11.6)	0.9 (3.6)	0.5
		<i>Scenario 3</i> (π_E, π_T)					
	$\delta =$	(0.20, 0.02) -0.48	(0.40, 0.03) -0.13	(0.60, 0.04) 0.22	(0.68, 0.06) 0.32	(0.74, 0.20) -0.26	
TR	PO	0.0 (3.0)	0.1 (3.9)	37.5 (22.0)	53.8 (32.7)	0.8 (4.4)	7.8
Trade-off	PO	0.0 (3.2)	0.4 (6.6)	19.8 (21.2)	71.6 (34.4)	6.5 (5.7)	1.7
Trade-off	CR	0.0 (3.4)	1.6 (8.8)	32.2 (20.8)	49.4 (22.3)	15.7 (16.0)	1.0
		<i>Scenario 4</i> (π_E, π_T)					
	$\delta =$	(0.52, 0.01) 0.12	(0.62, 0.015) 0.29	(0.71, 0.02) 0.45	(0.79, 0.025) 0.58	(0.86, 0.03) 0.69	
TR	PO	0.2 (3.5)	33.9 (24.4)	62.9 (36.3)	2.6 (6.7)	0.0 (0.7)	0.4
Trade-off	PO	0.1 (4.7)	1.7 (11.9)	10.1 (16.8)	34.3 (19.8)	53.7 (18.8)	0.0
Trade-off	CR	0.0 (3.5)	0.1 (4.3)	1.1 (5.3)	4.6 (6.6)	94.0 (52.2)	0.1
		<i>Scenario 5</i> (π_E, π_T)					
	$\delta =$	(0.05, 0.18) -1.03	(0.20, 0.22) -0.90	(0.35, 0.26) -0.85	(0.47, 0.30) -0.94	(0.58, 0.33) -1.07	
TR	PO	0.1 (3.4)	2.1 (4.6)	3.2 (5.0)	0.2 (2.5)	0.0 (0.5)	94.3
Trade-off	PO	0.0 (3.7)	2.1 (5.7)	1.9 (5.3)	0.3 (1.4)	0.1 (0.1)	95.6
Trade-off	CR	0.1 (3.4)	0.9 (8.3)	1.6 (3.6)	1.4 (0.8)	0.2 (0.3)	97.3
		<i>Scenario 6</i> (π_E, π_T)					
	$\delta =$	(0.15, 0.08) -0.66	(0.38, 0.18) -0.50	(0.52, 0.25) -0.64	(0.59, 0.30) -0.89	(0.62, 0.35) -1.18	
TR	PO	0.5 (3.6)	16.8 (17.5)	3.9 (11.7)	0.0 (2.2)	0.0 (0.1)	78.7
Trade-off	PO	0.2 (4.1)	12.7 (20.6)	1.9 (6.7)	0.0 (0.5)	0.0 (0.0)	85.2
Trade-off	CR	0.4 (5.3)	11.4 (20.1)	1.3 (4.5)	0.0 (1.1)	0.0 (0.4)	86.9

outperforms the four-parameter CR model in this case, we constructed a hypothetical data set of 70 patients. At each dose, 14 patients had counts $(Y_E, Y_T, 14 - Y_E - Y_T)$ set as close as possible to the fixed values $14(\pi_E, \pi_T, 1 - \pi_E - \pi_T)$.

Fitting each model to these data yielded posterior odds $\Pr(\text{data}|\text{PO})/\Pr(\text{data}|\text{CR}) = 50$. Thus, the PO model's superior performance in this case is due to the fact that it happens to fit Scenario 3 much more closely. This case also




Figure 3. Decision percentages for the Pentostatin trial using the trade-off-based algorithm, under the bivariate binary outcome model, for each of the dose-outcome scenarios given in Table 2. The doses of Pentostatin are labeled 1, 2, 3, and 4. The area of each disc equals the probability that the dose having (π_E, π_T) located as its center is selected. The shaded area in the lower right portion of each graph is set of (π_E, π_T) that are more desirable than the pairs on the elicited target contour.

suggests that the trade-off-based method under the CR model is robust.

We illustrate the new algorithm’s behavior in the bivariate binary outcome case in a simulation study of the Pentostatin trial under the six scenarios given in Table 2 and illustrated in Figure 3. In each of these scenarios, we assumed initially that $\psi = 0$, so that E and T occurred independently according to the fixed values of (π_E, π_T) at each dose given in Table 2. The trial design parameters were those actually being used to conduct the trial, with $c = 3, N = 36, \pi_E = 0.20, \bar{\pi}_T = 0.40, p_E = p_T = 0.10$, starting at dose level 1, and \mathcal{C} determined by fitting the curve $\pi_T = a + b/\pi_E + c/\pi_E^2$ to the elicited trade-off points $(\pi_E, \pi_T) = (0.15, 0), (0.25, 0.30), (1, 0.60)$. The means and SDs of the Gaussian prior parameters were $(-0.619, 0.941)$ for μ_T , $(0.587, 1.659)$ for β_T , $(-1.496, 1.113)$ for μ_E , $(1.180, 0.869)$ for $\beta_{E,1}$, $(0.149, 1.192)$ for $\beta_{E,2}$, and

$(0.00, 1.00)$ for ψ . In addition to having bivariate binary outcomes, the Pentostatin trial differs from the stroke trial in that $(\bar{\pi}_E, \bar{\pi}_T)$ and \mathcal{C} are numerically quite different, and moreover the sample size of 36 is half that of the stroke trial.

These simulation results are extremely favorable. In all six cases, the method makes a correct decision, namely selecting an acceptable dose or stopping early when no doses are acceptable, at least 94% of the time, and relatively few patients are treated at undesirable doses. Scenarios 1–4 are graphically similar to analogous cases studied in the stroke trial, and the algorithm shows behavior similar to that of (trade-off, CR) in the trinary case. However, in the bivariate binary case, $\pi_E(x)$ may vary freely with x . Under Scenario 3, where $\pi_E(x)$ increases and then sharply decreases with x , the method very reliably detects this pattern and selects the best dose, 0.50 mg/m². In Scenario 5, toxicity is moderate for

Table 2

Selection percentage, followed by number of patients treated in parentheses, by dose, of the trade-off-based algorithm under the bivariate binary model, for the GVHD treatment trial

Scenario		Pentostatin dose (mg/m ²)				None
		0.25	0.50	0.75	1.00	
1	(π_E, π_T)	(0.02, 0.05)	(0.30, 0.12)	(0.55, 0.30)	(0.65, 0.80)	0.0
	δ	-0.16 0.0 (3.0)	0.16 17.3 (8.7)	0.35 80.3 (20.8)	-0.45 2.4 (3.4)	
2	(π_E, π_T)	(0.02, 0.05)	(0.28, 0.10)	(0.50, 0.16)	(0.80, 0.22)	0.0
	δ	-0.16 0.0 (3.0)	0.14 0.5 (3.6)	0.38 13.6 (7.7)	0.59 86.0 (21.7)	
3	(π_E, π_T)	(0.25, 0.05)	(0.65, 0.15)	(0.50, 0.42)	(0.05, 0.65)	0.2
	δ	0.11 5.2 (5.8)	0.55 81.3 (21.0)	0.18 13.1 (7.8)	-0.38 0.2 (1.3)	
4	(π_E, π_T)	(0.45, 0.05)	(0.50, 0.45)	(0.55, 0.70)	(0.60, 0.85)	0.1
	δ	0.35 72.9 (23.5)	0.14 26.9 (11.2)	-0.28 0.1 (1.3)	-0.54 0.0 (0.1)	
5	(π_E, π_T)	(0.80, 0.05)	(0.50, 0.10)	(0.28, 0.16)	(0.02, 0.22)	0.2
	δ	0.76 92.5 (32.9)	0.40 4.8 (1.2)	0.13 2.1 (1.0)	-0.19 0.4 (0.9)	
6	(π_E, π_T)	(0.05, 0.50)	(0.25, 0.75)	(0.50, 0.85)	(0.70, 0.87)	94.5
	δ	-0.26 3.5 (8.0)	-0.41 1.4 (5.6)	-0.55 0.3 (2.6)	-0.58 0.3 (1.2)	

all dose levels, but response decreases monotonically, and in this case the algorithm very reliably chooses doses level 1. In Scenario 6, where no dose is acceptable, the trial is correctly stopped early with no dose selected 94.5% of the time, and the mean sample size is 17.4 patients.

To assess the method's sensitivity to association between E and T , we simulated extended versions of Scenario 2 with fixed values of ψ equal to $\{-2.049, -0.814, 0, +0.814, +2.0486\}$, which correspond to $\pi_{E|T}$ at 0.50 mg/m² equal to $\{0.14, 0.21, 0.28, 0.35, 0.42\}$. This range of association corresponds to varying $\pi_{E|T}$ at 0.50 mg/m² from 0.5 π_E to 1.5 π_E . For each value of ψ and each dose x , the simulation probabilities $\{\pi_{a,b}(x) : a, b = 0, 1\}$ were obtained from equation (3) using ψ and the fixed marginal values of (π_E, π_T) given in Table 2 for Scenario 2. As a basis for comparison, in the case where E and T are independent ($\psi = 0$) given in Table 2, dose levels 3 or 4 are selected 13.6% and 86.0% of the time. These two percentages become (64.0, 35.6) for $\psi = -2.049$, (69.4, 30.4) for $\psi = -0.814$, (40.4, 41.1) with no dose chosen 17.8% of the time for $\psi = +0.814$, and (82.0, 15.8) for $\psi = +2.049$. These results are not unexpected, for the following reasons. If $\pi_{E|T}(x) < \pi_E(x)$, then the occurrence of T at x reduces the posterior values of the marginal probability $\pi_E(x, \theta)$ and hence makes it less likely that x will satisfy the efficacy acceptability criterion (1). If $\pi_{E|T}(x) > \pi_E(x)$ then, by Bayes's Law, $\pi_{T|E}(x) > \pi_T(x)$ and the occurrence of E at x increases the marginal posterior values of $\pi_T(x, \theta)$ and hence makes it less likely that x will satisfy the toxicity acceptability criterion (2). Thus, either strong positive or strong negative association between E and T reduces the likelihood of acceptability of all doses. Given the dual goals of controlling toxicity while achieving a minimal rate of efficacy, this is

a desirable property of the method, provided that ψ does not vary with x . If it is anticipated that association may in fact vary with dose, then an extended model, say with $\psi = \mu_\psi + \beta_{\psi,1}x + \beta_{\psi,2}x^2$, may be more appropriate.

7. Discussion

The methodology proposed here provides a substantial improvement over the TR method in the trinary case and, using an appropriate model, it also accommodates the bivariate binary outcomes. In both cases, our simulation results indicate that the method has very desirable properties. The method is somewhat more structured than most dose-finding methods, and thus requires the statistician to work harder. We feel that this extra effort is well warranted by the scientific and ethical advantages of accounting for both efficacy and toxicity. Certainly, this methodology requires reliable, user-friendly computer programs for the simulations during the design process and trial conduct. These computer programs are freely available from the second author on request.

In developing this methodology, we investigated numerous other models and algorithms. These are not reported due to space limitations and, more importantly, the fact that none worked better than the method described here, and most were inferior. For example, in the trinary outcome case, one may extend the PO model by defining $\eta_T(x, \theta) = \mu + x\beta_1$ and $\eta_{E \cup T}(x, \theta) = \mu + \alpha + x\beta_2$. However, this requires the constraint $\alpha + (\beta_2 - \beta_1)x > 0$ for all x , which severely reduces the model's tractability. Alternatively, one may let $\eta_{E \cup T}(x, \theta) = \mu + \alpha + x\beta + x^2\gamma$, but this produces an inferior method. For bivariate binary outcomes, defining linear $\eta_E(x, \theta) = \mu_E + x\beta_E$, rather than using the quadratic form $\mu_E + x\beta_{E,1} + x^2\beta_{E,2}$ as we have done, greatly reduces the method's

ability to deal with cases such as Scenarios 1 and 3 given in Table 2 and Figure 3.

The method could be extended to account for patient prognostic covariates, (Z_1, \dots, Z_q) , by including additional terms of the form $Z_1\gamma_1 + \dots + Z_q\gamma_q$ in the linear components. This would allow patient-specific dosing, similarly to the method described by Babb and Rogatko (2001) for a single binary toxicity outcome. Our use of the trade-off contours, $\{C_z : z > 0\}$, which may be considered as utilities, to order doses by their desirabilities suggests that a decision-theoretic approach might yield an even better method. This would involve the computational difficulties of performing a backward induction to optimize the decision, however (cf. Carlin, Kadane, and Gelfand, 1998). We currently are investigating these extensions.

ACKNOWLEDGEMENTS

Peter Thall's research was partially supported by NCI grant CA 83932. We thank an associate editor for very constructive comments on two earlier drafts of the manuscript.

RÉSUMÉ

Pour les choix de doses dans les essais de phase I/II, nous présentons une méthode bayésienne adaptative basée sur un arbitrage entre les probabilités respectives de toxicité ou d'efficacité du traitement. La méthode permet de traiter des critères ternaires ou binaires bivariés, ainsi que des relations dose-efficacité éventuellement non monotones. On étudie les domaines du plan des probabilités d'efficacité et de toxicité et on se base sur les contours d'ensemble d'arbitrages dans ce plan pour définir les doses destinées aux patients successivement enrôlés. Les lois a priori sont établies en optimisant les lois des hyperparamètres pour l'ajustement d'un modèle de la réponse moyenne observée. Pour les critères ternaires, on compare le nouvel algorithme à la méthode de Thall et Russell (1998, *Biometrics* **54**, 251–264) en l'appliquant à un essai de traitement rapide des accidents vasculaires cérébraux ischémiques. Le cas bivarié est illustré par une application à un essai sur les rejets hôte-greffon dans les transplantations allogéniques de moelle osseuse. Des simulations montrent que, pour un large éventail de scénarios dose-événement, la nouvelle méthode a une forte probabilité de conduire à une décision correcte et de traiter la plupart des patients avec des doses ayant un ratio efficacité-toxicité souhaitable.

REFERENCES

- Babb, J. S. and Rogatko, A. (2001). Patient specific dosing in a phase I cancer trial. *Statistics in Medicine* **20**, 2079–2090.
- Braun, T. (2002). The bivariate continual reassessment method: Extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials* **23**, 240–256.
- Carlin, B. P., Kadane, J. B., and Gelfand, A. E. (1998). Approaches for optimal sequential decision analysis in clinical trials. *Biometrics* **54**, 964–975.
- Couriel, D., Champlin, R., Ippoliti, C., et al. (2001). *A sequentially adaptive, open-label, dose finding, phase I/II trial of Pentostatin in the treatment of steroid-refractory acute graft-versus-host disease (aGVHD)*. Clinical Protocol ID 01-327, M.D. Anderson Cancer Center.
- deLima, M., Champlin, R., Kantarjian, H., et al. (2001). *Allogeneic stem cell transplantation using Mylotary plus non-myeloablative chemotherapy in older or medically infirm patients with high-risk acute leukemia, chronic myelogenous leukemia or myelodysplastic syndrome*. Clinical Protocol ID 00-153, M.D. Anderson Cancer Center.
- Gloneck, G. F. V. and McCullagh, P. (1995). Multivariate logistic models. *Journal of the Royal Statistical Society, Series B* **57**, 533–546.
- Gooley, T. A., Martin, P. J., Fisher, L. D., and Pettinger, M. (1994). Simulation as a design tool for phase I/II clinical trials: An example from bone marrow transplantation. *Controlled Clinical Trials* **15**, 450–462.
- Joe, H. (1997). *Multivariate Models and Dependence Concepts*. London: Chapman and Hall.
- McCullagh, P. (1989). Models for discrete multivariate responses. *Bulletin of the International Statistical Institute* **53**, 407–418.
- McCullagh, P. and Nelder, J. A. (1989). *Generalized Linear Models*, 2nd edition. New York: Chapman and Hall.
- Monahan, J. and Genz, A. (1997). Spherical-radial integration rules for Bayesian computation. *Journal of the American Statistical Association* **92**, 664–674.
- Murtaugh, P. A. and Fisher, L. D. (1990) Bivariate binary models of efficacy and toxicity in dose-ranging trials. *Communications in Statistics, Part A—Theory and Methods* **19**, 2003–2020.
- Nelder, J. A. and Mead, R. (1965). A simplex method for function minimization. *Computer Journal* **7**, 308.
- 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.
- Quigley, J., Hughes, M. D., and Fenton, T. (2001). Dose-finding designs for HIV studies. *Biometrics* **57**, 1018–1029.
- Thall, P. F. and Russell, K. T. (1998). A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* **54**, 251–264.
- Thall, P. F., Estey, E. H., and Sung, H.-G. (1999). A new statistical method for dose-finding based on efficacy and toxicity in early phase clinical trials. *Investigational New Drugs* **17**, 155–167.
- Thall, P. F., Sung, H.-G., and Estey, E. H. (2002). Selecting therapeutic strategies based on efficacy and death in multi-course clinical trials. *Journal of the American Statistical Association* **97**, 29–39.
- Warach, S., Lynch, J. K., Davis, L., Haymore, J., Baird, A., Chalela, J., Hallenbeck, J., Todd, J., and Dunn, W. (2002). *ReoPro retavase reperfusion of stroke safety study—Imaging evaluation*. Clinical Protocol, National Institute of Neurological Diseases and Stroke, DHHS.

Received September 2003. Revised January 2004.

Accepted February 2004.