

Bayesian Models and Decision Algorithms for Complex Early Phase Clinical Trials

Peter F. Thall

Abstract. An early phase clinical trial is the first step in evaluating the effects in humans of a potential new anti-disease agent or combination of agents. Usually called “phase I” or “phase I/II” trials, these experiments typically have the nominal scientific goal of determining an acceptable dose, most often based on adverse event probabilities. This arose from a tradition of phase I trials to evaluate cytotoxic agents for treating cancer, although some methods may be applied in other medical settings, such as treatment of stroke or immunological diseases. Most modern statistical designs for early phase trials include model-based, outcome-adaptive decision rules that choose doses for successive patient cohorts based on data from previous patients in the trial. Such designs have seen limited use in clinical practice, however, due to their complexity, the requirement of intensive, computer-based data monitoring, and the medical community’s resistance to change. Still, many actual applications of model-based outcome-adaptive designs have been remarkably successful in terms of both patient benefit and scientific outcome. In this paper I will review several Bayesian early phase trial designs that were tailored to accommodate specific complexities of the treatment regime and patient outcomes in particular clinical settings.

Key words and phrases: Adaptive design, Bayesian design, clinical trial, dose-finding, phase I trial, phase I/II trial.

1. INTRODUCTION

1.1 An Early Phase Trial

Clinical trials are much more complex than typical statistical designs may indicate. An example is a phase I stem cell transplantation (SCT) trial in which the continual reassessment method (CRM, O’Quigley, Pepe and Fisher, 1990; O’Quigley, 1990) was applied to optimize the per-administration dose (PAD) of gemcitabine, d_G , when added to an established two-agent preparative regimen consisting of intravenous busulfan and melphalan (Andersson et al., 2002). The design was used for each of two separate, parallel trials, one for allogeneic transplant (allotx), which uses stem cells from a matched donor, and one for autologous transplant (autotx), which uses the patient’s own stem cells.

Peter F. Thall is Professor, Department of Biostatistics, University of Texas, M. D. Anderson Cancer Center, Houston, Texas, USA (e-mail: rex@mdanderson.org).

For each patient, during the period from day -10 to day -1 preceding the SCT on day 0, each of the three agents was given on two or more days using a particular schedule and PAD. Previously, a six-day schedule of additional gemcitabine had been tried, but it was found to be too toxic, so in this trial each patient’s assigned d_G was given in a two-day schedule, on each of days -8 and -3 , for total gemcitabine dose $2d_G$. “Toxicity” was defined to be any regimen-related grade 4 or 5 adverse event (AE) occurring within 30 days post transplant and affecting a vital organ, but excluding AEs that occur routinely in SCT, such as marrow suppression and, in allotx, graft-versus-host disease (GVHD). Using the usual CRM criterion, the design’s nominal goal was to find d_G from a predetermined set of 10 PADs ranging from 225 to 3675 mg/m² having toxicity probability, $\pi(d_G, \theta)$, with posterior mean $E_\theta\{\pi(d_G, \theta)|data\}$ closest to the target 0.10, where θ denotes the model parameters. The principal investigator (PI) specified the conservatively low target 0.10

in part due to the previous negative experience with the six-day schedule, and also because 0.10 is consistent with the toxicity rate of the established two-agent regimen. In each subgroup, gemcitabine doses were to be chosen for successive cohorts of 3 patients, up to a maximum of 36 patients, with the safety rules that no untried dose could be skipped when escalating and accrual to the subgroup would be stopped if the lowest dose $d_G = 255$ was unacceptably toxic, formally if $\Pr\{\pi(225, \theta) > 0.10 | \text{data}\} > 0.80$.

Because clinical trials are medical experiments with human subjects, they often do not play out precisely as designed. In the course of this trial: (i) when no toxicity was seen in the first 24 patients the PI decided to change the two-day gemcitabine schedule $(-8, -3)$ to the three-day schedule $(-8, -6, -3)$ while maintaining the same total dose by giving $\frac{2}{3}d_G$ on each day, (ii) this three-day schedule was quickly found to cause severe skin toxicity in the first few patients who received it and accrual was suspended, and (iii) we redesigned the trial again by returning to the two-day schedule, but (iv) at the PI's request we also expanded the set of possible d_G values. After 28 allotx patients had been treated and fully evaluated, however, (v) concern about observed grade 3 mucositis and skin toxicities seen at higher dose levels caused the physicians to expand the definition of "toxicity" to include these events, which previously had been excluded if they could be resolved therapeutically within two weeks. Along with this change, they also decided to change the CRM target from 0.10 to 0.15. These last changes had the combined effect of substantially increasing the numbers of patients with "toxicity" among those treated at higher dose levels and greatly reducing the value of d_G recommended by the CRM. Per standard regulatory procedure, it was necessary to obtain institutional review board approval for each change in the design. So, this trial actually was designed five times, it evaluated effects of a combination of three agents given in overlapping pre-transplant schedules, both the dose and schedule of gemcitabine were varied adaptively during the trial using three different formulations of the CRM to choose gemcitabine PADs and ad hoc decisions for changing schedules, there were two simultaneous trials involving different SCT modalities, the dominant effects on toxicity were both d_G and the schedule of gemcitabine, and the definition of toxicity was changed near the end of the trial to be more inclusive and thus obtain a more protective dose selection criterion. Overshadowing all of this were the actual goals, which were not only to control toxicity but also

to reduce the rate and severity of GVHD in the allotx patients and to improve the rates of engraftment and 100-day survival, compared to the established preparative regimen.

1.2 Some Generalities

Denote the treatment administered to a given patient by x . In the designs discussed here, x will be the dose of an agent, the dose pair of two agents given together, or a (schedule, dose) combination consisting of a finite sequence of administration times and corresponding doses. Actual patient outcome in oncology trials is very complex, often including numerous different types of toxicity scored on ordinal scales of severity (grade), disease status scored as a binary or ordinal variable, with each often recorded at several successive evaluations, as well as the times of delay or discontinuation of treatment, drop-out or death. In sharp contrast, the outcome Y used for statistical decision-making during the trial usually is defined to be a single variable or possibly a vector of two variables. In the examples given here, the designs assume that Y is, respectively, a single binary toxicity indicator, a vector of two binary indicators of toxicity and efficacy, a vector of ordinal toxicities or a time-to-toxicity variable subject to right censoring. The model consists of a probability density function (p.d.f.) or mass function (p.m.f.) $f(y|x, \theta)$ of Y for a patient who receives treatment x , and a prior $p(\theta|\xi)$, where θ is the model parameter vector and ξ are fixed hyperparameters. The data observed from the first n patients in the trial are $\mathcal{D}_n = \{(x^{(1)}, Y^{(1)}), \dots, (x^{(n)}, Y^{(n)})\}$, with likelihood $\mathcal{L}_n(\mathcal{D}_n|\theta) = \prod_{i=1}^n f(Y^{(i)}|x^{(i)}, \theta)$ and posterior $p_n(\theta|\mathcal{D}_n, \xi) \propto \mathcal{L}_n(\mathcal{D}_n|\theta)p(\theta|\xi)$.

All of the designs that I will discuss here utilize Bayesian "learn-as-you-go" decision rules to choose x from the set of possible treatments, \mathcal{X} , based on the posterior $p_n(\theta|\mathcal{D}_n)$ computed from the most recent data available when a new patient is enrolled. Such a sequentially adaptive decision algorithm may be expressed as a sequence $\alpha = \{\alpha_n\}$ of functions $\alpha_n : \mathcal{D}_n \rightarrow \mathcal{X} \cup \phi$, where ϕ denotes the empty set and $\alpha_n(\mathcal{D}_n) = \phi$ means "Do not treat with any $x \in \mathcal{X}$." In general, $\{\alpha_n\}$ may include several adaptive decision rules used together, such as rules for choosing a dose, a dose pair or a (schedule, dose) combination, for temporarily suspending accrual to wait for additional data on previously treated patients, or for stopping the trial early because no $x \in \mathcal{X}$ is acceptable. The $(n + 1)$ st iteration of the Bayesian medical decision-making process may

be described by the sequence of mappings

$$(1) \quad \begin{aligned} \mathcal{D}_n &\xrightarrow{(f,p)} p_n(\theta|\mathcal{D}_n, \xi) \xrightarrow{\alpha_n} x_{n+1} \\ &\longrightarrow Y_{n+1} \longrightarrow \mathcal{D}_{n+1} \end{aligned}$$

in which Bayes' Theorem uses the assumed probability model (f, p) to map the observed data into a posterior, the decision rules α_n use this to choose the next treatment x_{n+1} , the patient is treated, the outcome Y_{n+1} is observed, and (x_{n+1}, Y_{n+1}) are incorporated into the data. This process is repeated until the end of the trial, which may be when either a maximum sample size N_{max} or trial duration T_{max} is reached, or because the trial is stopped early. The process whereby the expanding data set is used by applying Bayes' Theorem repeatedly to turn $p_n(\theta|\mathcal{D}_n, \xi)$ into $p_{n+1}(\theta|\mathcal{D}_{n+1}, \xi)$ may be called "iterative Bayesian learning," in that one learns about θ as additional data are observed during the trial. In the sequel, to simplify notation I will suppress dependence of the posterior on the prior hyperparameters ξ .

A design consists of the trial's entry criteria, treatments \mathcal{X} , set of possible patient outcomes, probability model (f, p) , decision rules α , N_{max} or T_{max} , and possibly a cohort size, c . Since α_n acts on \mathcal{D}_n indirectly through $p_n(\theta|\mathcal{D}_n)$ in Bayesian adaptive designs, evaluation of a design's properties must account for the fact that $\{p_n(\theta|\mathcal{D}_n)\}$ is a sequence of statistics. The complexity of the process summarized by (1), even for binary Y and a single dose x , has motivated the routine use of computer simulation under each of a set of assumed "true" f as a tool to evaluate the frequentist operating characteristics (OCs) of the design for various α 's. This is used as a basis for choosing decision rules, calibrating design parameters, and possibly calibrating the prior. There is nothing "non-Bayesian" about using frequentist OCs of a Bayesian design to adjust the prior and design parameters. On the contrary, because simulating a trial that is based on a Bayesian design allows the physician to better understand the consequences of particular prior values, simulation provides a tool for the physician to modify his/her prior so that it more accurately reflects what the physician actually believes. It also is important to examine the prior's properties in the natural parameter domain, such as the probability of toxicity at dose x , $\pi(x, \theta)$, rather than in terms of elements of θ that may have no intuitive meaning to a physician. One should also examine the first few decisions α_n for each of several possible configurations of data, in order to avoid a prior that does not make

sense. This is especially important for evaluating decisions that must be made early in the trial based on very little data, such as choosing the second cohort's dose based on data from the first cohort of three patients. The prior always has consequences in an early phase trial, regardless of how "uninformative" it may appear to be.

If one does not wish to use simulation as a design tool, the most common alternative approach is to first specify a formal optimality criterion and solve for α mathematically (cf. Haines, Perevozskaya and Rosenberger, 2003; Dette et al., 2008). However, the simulation-based OCs of a design obtained using a particular optimality criterion are often surprising, essentially because such a design's properties are a consequence of the optimality criterion used. One may maximize information, minimize the variance of a particular estimated quantity, minimize mean or maximum sample size, control false positive or other incorrect decision probabilities, minimize expected financial costs, minimize expected trial duration, optimize outcomes for patients in the trial or for future patients, etc. Since, unavoidably, such goals often are at odds with each other, use of the word "optimal" without qualification may be very misleading.

1.3 Some Practical Issues

Actual clinical trial logistics can be quite complex. While the CRM adaptively chooses a new dose from a continuum for each new patient, Goodman, Zahurak and Piantadosi (1995) proposed the practical modifications of choosing doses for successive cohorts of several patients and limiting doses to a finite set. In most of the outcome adaptive dose-finding applications that I have seen, each newly chosen x is given to a cohort. Moreover, a "do not skip" safety rule often is imposed that does not allow an untried dose to be skipped when escalating. While limiting doses to a finite set of discrete values usually does not allow the exact MTD to be chosen, the difference between the chosen dose and the true MTD may be small, provided that the chosen sizes doses are reasonable. For example, using the dose set $\{100, 200, 400, 800, 1600\}$ must miss an actual MTD of 600 by at least 200, which is larger than the difference between the first two doses. Moreover, escalation from 400 to 800 or from 800 to 1600 may be unsafe, regardless of whether a do-not-skip rule is imposed.

For example, if a cohort size $c = 3$ is used, then α_n chooses $x_{n+1} = x_{n+2} = x_{n+3}$ and, formally, $Y_{n+1},$

Y_{n+2} and Y_{n+3} all must be observed before updating the data and making the next decision. However, if Y_{n+1} has been fully evaluated at the time patient $n + 2$ is accrued, then x_{n+2} may be chosen more reliably by applying the decision criterion based on the updated posterior incorporating the data (x_{n+1}, Y_{n+1}) from patient $n + 1$; similarly, if Y_{n+1} and Y_{n+2} are known when patient $n + 3$ is accrued, their data may be included to choose x_{n+3} . A simple approach that works surprisingly well is to use the “look ahead” rule: If the possible outcomes of treated patients for whom Y has not yet been fully evaluated will not alter the chosen x for the next patient, then treat the next patient with x without delay (Thall et al., 1999). This is closely related to the general fact that the time window required to evaluate Y per its definition and the accrual rate together play critically important roles in trial conduct. For example, if $Y = I$ (toxicity within 3 months from start of therapy) and the accrual rate is 6 patients per month, then any outcome-adaptive rule based on $\Pr(Y = 1|x, \theta)$ is virtually useless, since a large number of patients will be treated before the rule may be applied. Some possible ways to implement an outcome-adaptive design in such settings are as follows: (i) use $c = 1$ but enroll only a very small proportion of eligible patients in the trial, (ii) use $c = 3$ or larger with accrual suspended between cohorts, but use the look-ahead rule to improve logistical feasibility, or (iii) redefine the outcome to be time-to-toxicity, but use a safety rule that may delay accrual interimly to allow the data from previously treated patients to mature (cf. Bekele et al., 2008).

At the start of the trial, when $n = 0$, the first treatment x_1 may be chosen by applying the decision rule α_0 based on the prior $p(\theta|\xi)$. Methods for choosing a starting dose have been proposed by Goodman, Zahurak and Piantadosi (1995) and Cheung (2005). Most commonly, x_1 is chosen by the physician based on the nature of x , the definition of Y , the trial’s entry criteria and clinical experience treating the disease. For example, a trial enrolling prostate cancer patients with a life expectancy of six years is very different from a trial enrolling brain tumor patients with a life expectancy of six months. Similarly, depending on the trial’s entry criteria and treatment, “toxicity” may be defined as anything from severe fatigue to regimen-related death. It thus makes sense, during the prior elicitation process, to calibrate $p(\theta)$ so that α_0 agrees with the physician’s x_1 , since the motivation for choosing a particular x_1 is based on prior experience.

In this paper I will review several designs that focus on the problem of reflecting more fully particular complexities of (x, Y) . Each design addresses some, but not all, of the issues in the SCT trial described earlier. These methods were motivated by problems that I have encountered during the process of designing early phase trials over the past 19 years. Each design was developed by a collaborative team including one or more physicians, one or more statisticians and a computer programmer. Each may be called a “phase I” design in that dose-finding is based on toxicity, or a “phase I/II” design in that dose-finding is based on both efficacy and toxicity, with the exception of the design described in Section 5 that jointly optimizes schedule and dose (Braun et al., 2007). While it is tempting to think that a “one-size-fits-all” design encompassing all phase I/II possibilities may be constructed, in my experience clinical research is far too complex to do this, and each new trial design problem often has unique aspects that require a new model or method. The particular data structure, probability model and decision rules that should be used to design a clinical trial are best determined through careful discussion with the physicians planning the trial, and must strike a compromise between the desire to accurately reflect the medical process and address scientific goals while accommodating the practical realities of trial conduct.

I will not discuss methods for eliciting and calibrating priors, since this topic could easily fill an additional manuscript. I will not explore the ethical aspects of adaptive decision rules either, since they also are quite complex (cf. Palmer, 2002). Early phase trial design and conduct are difficult and complicated in large part due to the tension between optimizing the benefit and safety of patients treated in the trial, and learning about the effects of each x on Y to benefit future patients, as well as economic constraints and regulatory requirements. In this regard, a statistician constructing an adaptive design should be mindful of the ethical issues regarding what happens, for example, to patient number 7 because (s)he was treated with x_6 based on how α_6 acted on $p_6(\theta|\mathcal{D}_6)$.

2. DOSE-FINDING FOR TWO-AGENT COMBINATIONS

2.1 Outcomes and Models

Thall et al. (TMML, 2003) proposed a method for determining one or more acceptable dose pairs $x = (x_1, x_2)$ of two cytotoxic agents given together, based on a binary indicator Y of toxicity. To stabilize the model numerically, each dose is standardized so that $0 \leq x_1, x_2 \leq 1$, for example, by dividing each raw dose

by some maximum value, so that each $x \in \mathcal{X} = [0, 1]^2$. The probability model for toxicity is

$$\begin{aligned} \Pr(Y = 1|x, \theta) &\equiv \pi(x, \theta) \\ &= \frac{\alpha_1 x_1^{\beta_1} + \alpha_2 x_2^{\beta_2} + \alpha_3 x_1^{\beta_1 \beta_3} x_2^{\beta_2 \beta_3}}{1 + \alpha_1 x_1^{\beta_1} + \alpha_2 x_2^{\beta_2} + \alpha_3 x_1^{\beta_1 \beta_3} x_2^{\beta_2 \beta_3}}, \end{aligned} \tag{2}$$

where $\theta = (\alpha_1, \beta_1, \alpha_2, \beta_2, \alpha_3, \beta_3)$. All elements of θ are positive valued, which ensures that $\pi(x, \theta)$ is a probability and that it is increasing in each entry of x . Denoting the subvectors $\theta_j = (\alpha_j, \beta_j)$ for $j = 1, 2, 3$, so that $\theta = (\theta_1, \theta_2, \theta_3)$, the model (2) contains the submodels

$$\begin{aligned} \pi_1(x_1, \theta_1) &= \Pr\{Y = 1|x = (x_1, 0), \theta_1\} \\ &= \frac{\alpha_1 x_1^{\beta_1}}{1 + \alpha_1 x_1^{\beta_1}}, \end{aligned} \tag{3}$$

which is the probability of toxicity when agent 1 is given alone at dose x_1 and, similarly,

$$\begin{aligned} \pi_2(x_2, \theta_2) &= \Pr\{Y = 1|x = (0, x_2), \theta_2\} \\ &= \frac{\alpha_2 x_2^{\beta_2}}{1 + \alpha_2 x_2^{\beta_2}} \end{aligned} \tag{4}$$

for agent 2 given alone at dose x_2 . Since $\alpha_j x_j^{\beta_j} = \exp\{\log(\alpha_j) + \beta_j \log(x_j)\}$, (3) and (4) are logistic models in a log standardized dose. TMML assume that there is clinical experience with each single agent when used alone, since this often is a requirement before investigating a combination in humans. Since θ_j parameterizes $\pi_j(x_j, \theta_j)$ for $j = 1, 2$ and θ_3 parameterizes interaction between the two agents, a key element of TMML's approach is that the priors $p(\theta_1|\xi_1)$ and $p(\theta_2|\xi_2)$ are informative while $p(\theta_3|\xi_3)$ is vague. Assuming gamma priors on the elements of θ for tractability, TMML provide a detailed algorithm for eliciting $p(\theta_1|\xi_1)$ and $p(\theta_2|\xi_2)$, although if historical data are available, the posteriors from preliminary fits of such data may be used as these priors for trial design and conduct. Considering $\pi(x, \theta)$ geometrically as a response surface over the domain $[0, 1]^2$, this says that there is substantial prior knowledge about each of the two lines $\{x : x_2 = 0\}$ and $\{x : x_1 = 0\}$ on the edges of the response surface, but otherwise little is known about the surface, so it is like a sheet tied down at two edges but otherwise varying freely. In particular, the meaning of θ_1 in the model $\pi(x, \theta) = \pi(x, (\theta_1, \theta_2, \theta_3))$ is very different from its meaning in the submodel

$\pi_1(x_1, \theta_1)$. This is underscored by the prior effective sample sizes computed by Morita, Thall and Mueller (2008) for the gamma priors given by TMML (2003, Section 3), which are 547.3 for $p_1(\theta_1|\xi_1)$, 756.8 for $p_2(\theta_2|\xi_2)$, 0.01 for $p_3(\theta_3|\xi_3)$ and 1.5 for $p(\theta|\xi)$. This says that, with respect to toxicity, a priori a lot is known about how each agent behaves when used alone, but almost nothing is known about how the two agents behave together.

2.2 Decision Criteria

The dose-finding method exploits the following geometric structure on $\mathcal{X} = [0, 1]^2$. For each $p \in (0, 1)$, the set $\mathcal{X}_p(\theta) = \{x : \pi(x, \theta) = p\}$ is the isocontour of all dose pairs having toxicity probability p . Several isocontours for a particular fixed θ are illustrated in Figure 1. Since $\mathcal{X}_p(\theta) \cap \mathcal{X}_q(\theta) = \emptyset$ if $p \neq q$ and $\bigcup_{0 \leq p \leq 1} \mathcal{X}_p(\theta) = [0, 1]^2$, every pair x falls on a unique $\mathcal{X}_p(\theta)$ for some p . The interaction term $\alpha_3 x_1^{\beta_1 \beta_3} x_2^{\beta_2 \beta_3}$ in (2) is used instead of the simpler term $\alpha_3 x_1 x_2$ in order to give the model sufficient flexibility to allow ‘‘S’’ shaped isocontours, as shown in Figure 1.

The design proceeds in two stages. In stage 1, doses are chosen for successive cohorts from a finite set of values on the predetermined fixed diagonal line L_1 , shown as the straight line at approximately 45° in Figure 1. The design is robust to the particular angle of L_1 , as long as it is not too far from 45°. Since the response surface $\pi(x, \theta)$ increases in each argument x_1 and x_2 , $\pi(x, \theta)$ must increase as x moves up L_1 from lower left to upper right. Given target toxicity probability π^* , dose-finding in stage 1 proceeds using the CRM

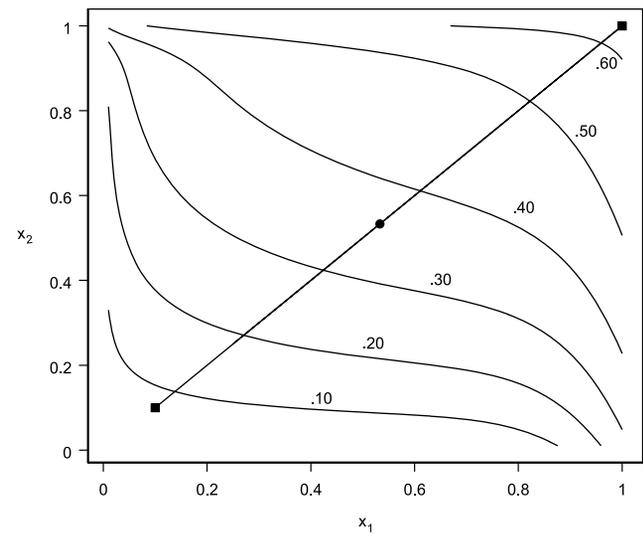


FIG. 1. Isocontours of a dose-toxicity probability surface for two dimensional dose $x = (x_1, x_2)$.

criterion of choosing x for each cohort from the set on L_1 to minimize $|\mathbb{E}\{\pi(x, \theta)|\mathcal{D}_n\} - \pi^*|$, starting at the lowest dose in the set, not skipping untried doses when escalating, and adding additional doses to the set once the first toxicity is observed. That is, restricting x to L_1 in stage 1 reduces dose selection to a one-dimensional problem, and a conventional CRM algorithm may be applied. Geometrically, one may think of stage 1 as walking up and down the toxicity surface, along L_1 , looking for dose pairs with toxicity probability close to π^* .

In stage 2, x is chosen for successive cohorts from the random isocontour

$$(5) \quad \mathcal{X}_{\pi^*}(\mathcal{D}_n) = \{x : E\{\pi(x, \theta)|\mathcal{D}_n\} = \pi^*\}.$$

Since $\mathcal{X}_{\pi^*}(\mathcal{D}_n)$ contains infinitely many x , an additional criterion is needed to choose one x for each cohort in stage 2. TMML suggest two criteria, one based on the clinical criterion of “cancer killing potential” and the other the more usual statistical goal of maximizing Fisher Information. Denoting the elements of θ by $\theta_1, \dots, \theta_6$ for convenience, the Fisher Information matrix $I(x, \theta)$ for dose x has (j, k) entry $\{\partial\pi(x, \theta)/\partial\theta_j\}\{\partial\pi(x, \theta)/\partial\theta_k\}/[\pi(x, \theta)\{1 - \pi(x, \theta)\}]$. Under the Bayesian model, x is chosen to maximize the posterior mean log determinant of $I(x, \theta)$ given the current data, $\mathbb{E}[\log\{\det I(x, \theta)\}|\mathcal{D}_n]$. Doses are chosen for successive cohorts in stage 2 by alternating between the two subsets of $\mathcal{X}_{\pi^*}(\mathcal{D}_n)$ to the left and right of L_1 . For each subset, the x optimizing cancer killing is determined, the x maximizing Fisher Information is determined, the average of these two dose pairs is computed, and the $x \in \mathcal{X}_{\pi^*}(\mathcal{D}_n)$ closest to this average is assigned to the cohort. At the end of the trial, any $x \in \mathcal{X}_{\pi^*}(\mathcal{D}_n)$ is a solution. Thus, for example, one may choose three final dose pairs on $\mathcal{X}_{\pi^*}(\mathcal{D}_N)$, one on L_1 , one to the left of L_1 and one to the right of L_1 , and randomize patients among these three x pairs in a subsequent phase II trial.

In their illustrative application, TMML use a cohort size of 2 with 60 patients divided into $N_1 = 20$ (10 cohorts) in stage 1 and $N_2 = 40$ (20 cohorts) in stage 2. Simulations show that using (N_1, N_2) either (30, 30) or (40, 20) gives a design with inferior properties compared to (20, 40). Cohort-by-cohort computations show that the target isocontour $\mathcal{X}_{\pi^*}(\mathcal{D}_n)$ varies substantially with each new cohort’s data even after $n = 30$ or 40 patients, but stabilizes by $n = 50$ or 60. This is the case essentially because a binary outcome is a very small amount of information per patient. For total sample size $N = N_1 + N_2 = 60$, however, the method is quite reliable in terms of choosing dose pairs that have true toxicity probability close to π^* .

3. USING BOTH EFFICACY AND TOXICITY

From a clinical perspective, the primary purpose of treatment is to fight disease, and safety is never a secondary concern in any medical setting. Thus, both scientifically and medically, both efficacy and toxicity matter at all stages of clinical investigation. In many dose-finding trials, the entry criteria specify patients with such poor prognosis that response is very unlikely, between 4% and 10%. In such settings, targeting even a low response rate $\pi_{R^*} = 0.10$ or 0.20 and using a phase II type rule to stop accrual if the observed response rate is likely to be below π_{R^*} at any acceptable dose may be impractical, since few or no responses are expected. This is the most common rationale for conducting phase I based on toxicity alone, while recording data on biological effects and possible clinical anti-disease effects. However, actual response rates vary widely between phase I trials, and many have complete or partial response rates well over 20% (cf. Horstmann et al., 2005). Moreover, patients enroll in a phase I trial motivated by the hope that the new treatment will achieve an anti-disease effect, not simply the desire that no toxicity will occur.

3.1 Outcomes and Models

These considerations lead to the idea that, when it is realistic to target a response rate of 10% or larger, dose-finding should be done using a phase I/II design based on both $E =$ efficacy and $T =$ toxicity. Many phase I/II designs have been proposed (Gooley et al., 1994; O’Quigley, Fenton and Hughes, 2001; Braun, 2002; Ivanova, 2003). The following phase I/II methodology, “EffTox,” is based on the developments given by Thall and Russell (1998) and Thall and Cook (2004). Illustrations are given by Thall, Cook and Estey (2006) and Whelan et al. (2008). Patient outcome may be either a three-category or bivariate binary variable. The former case applies when E and T are defined in such a way that they are disjoint but $E \neq T^c$, so that Y takes on values in $\{E, T, N\}$ where $N = (E \cup T)^c = \{\text{no response and no toxicity}\}$. This is appropriate if, for example, toxicity is irreversible organ damage or regimen-related death. When it is possible for both E and T to occur, the outcome is bivariate binary, $Y = (Y_E, Y_T)$, where Y_k indicates the outcome $k = E, T$. For either case, denote the outcome probabilities for a patient given dose x by $\pi(x, \theta) = (\pi_E(x, \theta), \pi_T(x, \theta))$.

For the trinary outcome case, the three-parameter model used by Thall and Russell (1998) is motivated by the idea that the three outcomes are ordered in the sense that $N < E < T$, with the idea

that higher x is more likely to push the patient’s outcome upward along this scale. The model is given by $\pi_T(x, \theta) = g^{-1}(\beta_0 + \beta_2x)$ and $\pi_E(x, \theta) = g^{-1}(\beta_0 + \beta_1 + \beta_2x) - \pi_T(x, \theta)$, where g is a link function, $\theta = (\beta_0, \beta_1, \beta_2)$ and $\beta_1, \beta_2 > 0$. This model forces $\pi_E(x, \theta)$ to be very nonmonotone in x . A more flexible four-parameter model (Thall and Cook, 2004, Section 3) is given by $\Pr(E|T^c, x, \theta) = g^{-1}(\beta_{E,0} + \beta_{E,1}x)$ and $\pi_T(x, \theta) = g^{-1}(\beta_{T,0} + \beta_{T,1}x)$, where $\theta = (\beta_{E,0}, \beta_{E,1}, \beta_{T,0}, \beta_{T,1})$ and $\beta_{E,1}, \beta_{T,1} > 0$. Using this model, $\pi_E(x, \theta) = \Pr(E|T^c, x, \theta)\{1 - \pi_T(x, \theta)\}$. For either model, $f(Y|x, \theta) = \prod_{y=E,N,T} \{\pi_y(x, \theta)\}^{I(Y=y)}$.

For the bivariate binary case, the model must specify the four elementary outcome probabilities $\pi_{a,b}(x, \theta) = \Pr(Y_E = a, Y_T = b|x, \theta)$ for $a, b \in \{0, 1\}$. The p.m.f. of a patient treated at dose x is $f(Y|x, \theta) = \prod_{a=0}^1 \prod_{b=0}^1 \{\pi_{a,b}(x, \theta)\}^{I(Y_E=a, Y_T=b)}$. The general approach used by Thall and Cook (2004) and Thall, Nguyen and Estey (2008) is to first specify the two marginal dose-outcome distributions $\pi_k(x, \theta) = g^{-1}\{\eta_k(x, \theta)\}$, in terms of link function g and linear terms $\eta_k(x, \theta)$ for $k = E$ and T , and then define the joint distribution in terms of the marginals. Temporarily suppressing x and θ , $\pi_{a,b}$ is determined by $(a, b, \pi_E, \pi_T, \psi)$, where ψ is an association parameter. This may be done tractably using a Gumbel distribution,

$$\begin{aligned} \pi_{a,b} &= \pi_E^a(1 - \pi_E)^{1-a} \pi_T^b(1 - \pi_T)^{1-b} \\ &+ (-1)^{a+b} \pi_E(1 - \pi_E) \\ &\cdot \pi_T(1 - \pi_T) \left(\frac{e^\psi - 1}{e^\psi + 1} \right), \end{aligned} \tag{6}$$

with ψ real-valued, or a Gaussian copula, $C(u, v) = \Phi_\psi(\Phi^{-1}(u), \Phi^{-1}(v))$, for $0 \leq u, v \leq 1$, where Φ_ψ is the bivariate standard normal c.d.f. with correlation ψ and Φ is the univariate $N(0, 1)$ c.d.f. Under this copula, $\pi_{0,0} = \Phi_\psi(\Phi^{-1}(1 - \pi_E), \Phi^{-1}(1 - \pi_T))$ with $\pi_{1,0} = 1 - \pi_T - \pi_{0,0}$, and $\pi_{1,1} = \pi_E + \pi_T + \pi_{0,0} - 1$. If g is the probit link, $\pi_k = \Phi(\eta_k)$ and $\pi_{0,0} = \Phi_\psi(-\eta_E, -\eta_T)$.

A major practical issue is that the $\eta_k(x, \theta)$ ’s should be realistic but the model must be numerically tractable, to facilitate the processes of fitting historical data if available, prior elicitation, and computing posterior decision criteria thousands of times while simulating the trial during the design process. It often is important to allow $\pi_E(x, \theta)$ to be nonmonotone in x , which may be appropriate for biological agents, such as viral vectors expressing cytokines aimed at triggering an immune response to kill tumor cells. This may be done very effectively by assuming a simple

quadratic $\eta_E(x, \theta) = \beta_{E,0} + x\beta_{E,1} + x^2\beta_{E,2}$, although other functions may be used. While $\eta_T(x, \theta) = \beta_{T,0} + x\beta_{T,1}$ with $\Pr(\beta_{T,1} > 0) = 1$ is appropriate for cytotoxic agents, in other settings a quadratic also may be used for $\eta_T(x, \theta)$. For example, if “toxicity” includes infection and an anti-cancer agent also kills bacterial or fungal infections, then $\pi_T(x, \theta)$ may be nonmonotone and actually decrease with higher x .

Thall and Cook (2004) provide a penalized least squares method for establishing the prior $p(\theta|\xi)$ based on elicited means $\mu_{y,j}^{(e)}$ and standard deviations (s.d.’s) $\sigma_{y,j}^{(e)}$ of $\pi_y(x_j, \theta)$ for $y = E, T$ and several doses x_1, \dots, x_m . Since each prior mean $\mu_{y,j}(\xi)$ and s.d. $\sigma_{y,j}(\xi)$ of $\pi_y(x_j, \theta)$ is a function of the fixed hyperparameters ξ characterizing $p(\theta|\xi)$, nonlinear least squares may be used to solve for ξ by minimizing the objective function

$$\begin{aligned} \text{SS}(\xi) &= \sum_{y=E,T} \sum_{j=1}^m [\{\mu_{y,j}^{(e)} - \mu_{y,j}(\xi)\}^2 \\ &+ \{\sigma_{y,j}^{(e)} - \sigma_{y,j}(\xi)\}^2] \\ &+ c \sum_{1 \leq j < k \leq m} \{\tilde{\sigma}_j - \tilde{\sigma}_k\}^2, \end{aligned} \tag{7}$$

where each $\tilde{\sigma}_k$ is a prior standard deviation in ξ . The second sum in (7) is included to limit the variability among the prior s.d.’s, using a small penalty constant $c > 0$.

3.2 Dose Admissibility and Efficacy-Toxicity Trade-offs

The dose-finding algorithm relies on two different types of posterior decision criteria to choose x from a finite set of possibilities, $\mathcal{X} = \{x^{(1)}, \dots, x^{(k)}\}$. The first criterion determines which doses are acceptable, and the second chooses the best acceptable dose. Let $\underline{\pi}_E$ be a fixed lower limit on $\pi_E(x, \theta)$ and $\bar{\pi}_T$ a fixed upper limit on $\pi_T(x, \theta)$. The fixed limits are specified by the physician. Let p_E^* and p_T^* both be fixed upper probability cut-offs, usually selected from the range 0.80 to 0.95. A dose x is *unacceptable* if it is likely to have either unacceptably low efficacy or unacceptably high toxicity, formally if

$$\begin{aligned} \Pr\{\pi_E(x, \theta) < \underline{\pi}_E | \mathcal{D}_n\} &> p_E^* \quad \text{or} \\ \Pr\{\pi_T(x, \theta) > \bar{\pi}_T | \mathcal{D}_n\} &> p_T^*. \end{aligned} \tag{8}$$

A dose x is *acceptable* if neither inequality in (8) holds. The set of acceptable doses in \mathcal{X} based on \mathcal{D}_n is denoted by \mathcal{A}_n . These criteria are essentially those used

by Thall, Simon and Estey (1995) as stopping rules in phase II trials, and the second criterion in (8) is used routinely for deciding whether to stop a phase I trial, for example, when using the CRM, if the lowest dose is too toxic. For example, when using EffTox, if $p_E^* = p_T^* = 0.90$, the rules in (8) are equivalent to saying that x is *acceptable* if $\Pr\{\pi_E(x, \theta) > \underline{\pi}_E | \mathcal{D}_n\} > 0.10$ and $\Pr\{\pi_T(x, \theta) < \bar{\pi}_T | \mathcal{D}_n\} > 0.10$. While, intuitively, these may seem like rather weak requirements, if one replaces 0.10 by a large cut-off such as 0.80 by setting $p_E^* = p_T^* = 0.20$ in (8), then the rules are nearly certain to stop any trial very quickly after a very small number of patients, due to the large variability of the posterior probabilities used in (8). This gets at the important distinction between determining the acceptability of x for the purpose of dose-finding with small to moderate sample sizes, and the confirmatory statement “ x is safe and effective” formalized by inequalities such as $\Pr\{\pi_E(x, \theta) > \underline{\pi}_E | \mathcal{D}_n\} > 0.95$ and $\Pr\{\pi_T(x, \theta) < \bar{\pi}_T | \mathcal{D}_n\} > 0.95$. Such confirmatory conclusions are inappropriate based on early phase trial results since they can only be established convincingly by a large sample size, regardless of what the posterior probabilities may be.

To describe the second decision criterion, for simplicity, I will focus on the bivariate binary case, where $\pi(x, \theta) \in [0, 1]^2$. To compare two acceptable doses, say, $x^{(1)}$ and $x^{(2)}$, based on the posteriors of $\pi(x^{(k)}, \theta) = (\pi_E(x^{(k)}, \theta), \pi_T(x^{(k)}, \theta))$ for $k = 1, 2$, some method for reducing each pair $\pi(x^{(k)}, \theta)$ to a one-dimensional criterion is required, as inevitably is the case when a statistic of dimension ≥ 2 is used for comparison. The EffTox method does this by formalizing the idea that a higher risk of toxicity is a reasonable trade-off for a higher probability of achieving anti-disease effect. The method first computes the posterior means $\mu^{(n)}(x) = (\mu_E^{(n)}(x), \mu_T^{(n)}(x)) = (E\{\pi_E(x, \theta) | \mathcal{D}_n\}, E\{\pi_T(x, \theta) | \mathcal{D}_n\})$ for each $x \in \mathcal{X}$. Each $\mu^{(n)}(x)$ is then reduced to a one-dimensional criterion by the following geometric construction, which begins by eliciting several pairs of fixed probabilities, $p^{(j)} = (p_E^{(j)}, p_T^{(j)})$, $j = 1, \dots, m$, that the physician considers equally desirable. A target curve, \mathcal{C} , is fit to the elicited pairs, treating p_T as a monotone increasing function of p_E , or, equivalently, reversing the roles of p_T and p_E . This should be done using a graphical representation of \mathcal{C} to provide a means for the physician to adaptively modify his/her target pairs. Given $p \in [0, 1]^2$, let p_C denote the point where the straight line segment in $[0, 1]^2$ passing through p and the ideal

point $(1, 0)$ intersects \mathcal{C} . The *desirability* of p may be defined as

$$(9) \quad \delta(p) = \exp\left\{-\frac{\|p - (1, 0)\|}{\|p_C - (1, 0)\|}\right\},$$

where $\|\cdot\|$ denotes Euclidean distance. This has maximum $\delta(1, 0) = 1$, with $\delta(p)$ decreasing as p moves away from $(1, 0)$ along any straight line in $[0, 1]^2$. Several other definitions of $\delta(p)$ may be used, although (9) is reasonable and tractable. The contour of all p having desirability u is $\mathcal{C}_u = \{p \in [0, 1]^2 : \delta(p) = u\}$, so that $\mathcal{C} = \mathcal{C}_{e^{-1}}$. Denote the set of real numbers u such that $\mathcal{C}_u \neq \emptyset$ by R_C . Since $u \neq v \implies \mathcal{C}_u \cap \mathcal{C}_v = \emptyset$, the family $\{\mathcal{C}_u, u \in R_C\}$ partitions $[0, 1]^2$. This construction is used to quantify the desirability of a dose x by evaluating (9) at $p = \mu^{(n)}(x)$. To compare doses $x^{(1)}$ and $x^{(2)}$, we compute $\mu^{(n)}(x^{(1)})$ and $\mu^{(n)}(x^{(2)})$, illustrated by the two round dots in Figure 2a, and then compute their desirabilities $\delta(\mu^{(n)}(x^{(1)}))$ and $\delta(\mu^{(n)}(x^{(2)}))$, as shown in Figure 2b. The elicited pairs are represented by the symbol “ \times ” in Figure 2b, which also shows \mathcal{C} and several \mathcal{C}_u along with their numerical u values. During the trial, if no dose is acceptable, formally if $\mathcal{A}_n = \emptyset$, then accrual is stopped with no dose selected; otherwise each cohort is given the dose x maximizing $\delta(\mu^{(n)}(x))$ among all $x \in \mathcal{A}_n$. This methodology has been used for dose-finding trials in acute stroke, treatment for GVHD in SCT, chemotherapy of acute leukemia, and anergized cells given post-transplant to accelerate immune reconstitution following allotx.

It is important to consider the consequences of how one sets goals in the bivariate binary outcome case. Recall that $\pi_{a,b} = \Pr(Y_E = a, Y_T = b)$ with $\pi_E = \pi_{1,1} + \pi_{1,0}$ and $\pi_T = \pi_{1,1} + \pi_{0,1}$. Several authors have proposed choosing x to maximize $\pi_{1,0}$, the probability of the best possible outcome, efficacy and no toxicity, or the conditional probability $\pi_{E|T^c} = \pi_{1,0}/(1 - \pi_T)$. Unfortunately, most new treatments simply don't work that way. A new therapy that is either more aggressive, for example, a higher dose, or highly active biologically is likely to decrease $\pi_{0,0}$ and increase some combination of $\pi_{1,0}$, $\pi_{0,1}$ and $\pi_{1,1}$. Treating $\pi = (\pi_{0,0}, \pi_{1,0}, \pi_{0,1}, \pi_{1,1})$ as fixed for simplicity, suppose that standard treatment gives outcome probability vector $\pi^{(0)} = (0.50, 0.10, 0.30, 0.10)$, which has marginals $(\pi_E, \pi_T) = (0.20, 0.40)$. Suppose that experimental treatment $x^{(1)}$ has $\pi^{(1)} = \pi(x^{(1)}) = (0.30, 0.20, 0.30, 0.20)$, which has marginals $(\pi_E^{(1)}, \pi_T^{(1)}) = (0.40, 0.50)$, a doubling of $\pi_E^{(0)}$ from 0.20 to 0.40 and a 25% increase in $\pi_T^{(0)}$ from 0.40 to 0.50. Suppose that a competing experimental treatment $x^{(2)}$ has

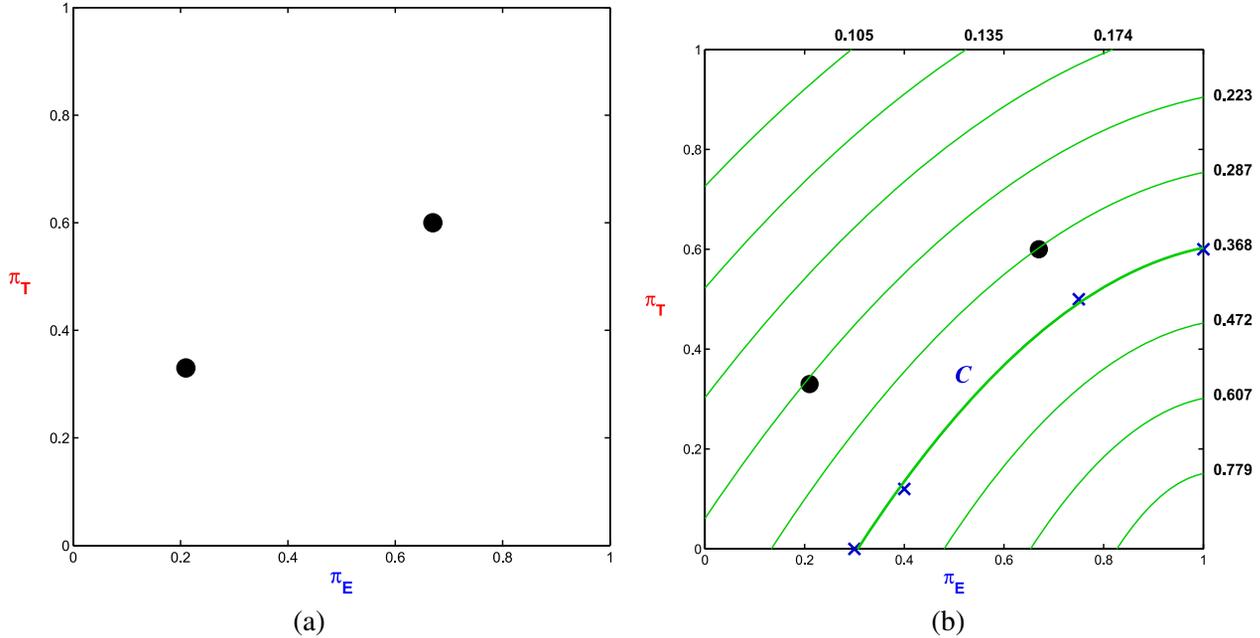


FIG. 2. Example of posterior means $\mu^{(n)}(x^{(j)}) = (\mathbb{E}\{\pi_E(x^{(j)}, \theta)|data_n\}, \mathbb{E}\{\pi_T(x^{(j)}, \theta)|data_n\})$ for two dose pairs $x^{(1)}$ and $x^{(2)}$ (denoted by round dots), given alone in the left-hand graph (Figure 2a), and with the addition of the target contour \mathcal{C} constructed from elicited target points (denoted by \times) and several resulting desirability contours (Figure 2b).

$\pi^{(2)} = \pi(x^{(2)}) = (0.30, 0.20, 0.45, 0.05)$, which has marginals $(\pi_E^{(2)}, \pi_T^{(2)}) = (0.25, 0.50)$, a slight increase in $\pi_E^{(0)}$ from 0.20 to 0.25 with the same increase in $\pi_T^{(0)}$ as given by $x^{(1)}$. Since $\pi_{1,0}(x^{(1)}) = \pi_{1,0}(x^{(2)}) = 0.20$ and $\pi_{E|T^c}(x^{(1)}) = \pi_{E|T^c}(x^{(2)}) = 0.40$, a method based on either $\pi_{1,0}$ or $\pi_{E|T^c}$ would consider $x^{(1)}$ and $x^{(2)}$ to be equivalent. In contrast, the trade-off based method would consider $x^{(1)}$ superior to $x^{(2)}$.

4. FINDING PATIENT-SPECIFIC DOSES

Thall, Nguyen and Estey (2008) generalized EffTox to account for patient heterogeneity by using the patient’s vector $Z = (Z_1, \dots, Z_q)$ of covariates observed at enrollment. The method requires historical data, \mathcal{H} , to obtain an informative distribution on covariate effect parameters for use in trial design and conduct. The model and method account for dose effects, covariate effects and possible dose-covariate interactive effects on π_E and π_T . The design assigns each patient a dose that is individualized based on the patient’s Z vector. This is very different from conventional early phase trial designs, since (i) patients with different covariates may receive different doses at the same point in the trial, (ii) the entry criteria may change adaptively, with the possibility that enrollment may be shut down for some patients but continued for others, and (iii) at

the end of the trial a computer-based rule is provided for assigning each future patient’s x based on his/her Z vector, rather than choosing a single dose for all patients.

For designs with individualized treatment assignment rules utilizing Z (cf. Ratain et al., 1996; Babb and Rogatko, 2001), the i th patient’s data are $(x^{(i)}, Z^{(i)}, Y^{(i)})$, the probability model is elaborated by defining $f(y|Z, x, \theta)$ for a patient with covariates Z who receives treatment x , and α_n is a function of (Z, \mathcal{D}_n) . To accommodate Z and historical data in the design described here, let τ denote either a dose x in the trial or historical treatment from the set $\{\tau_1, \dots, \tau_m\}$. The probability model given earlier is extended by defining the marginal probabilities, $\pi_k(\tau, Z, \theta) = g^{-1}\{\eta_k(\tau, Z, \theta)\}$, $k = E, T$, for a patient with covariates Z given dose x , assuming linear terms of the general form

$$\eta_k(\tau, Z, \theta) = \beta_k Z + \sum_{j=1}^m (\mu_{k,j} + \xi_{k,j} Z) I(\tau = \tau_j) + \{\omega_k(x, \alpha_k) + \gamma_k Z\} I(\tau = x)$$

(10)

for $k = E, T$, where $\beta_k Z = \beta_{k,1} Z_1 + \dots + \beta_{k,q} Z_q$ account for covariate main effects, $\gamma_k Z = \gamma_{k,1} Z_1 + \dots + \gamma_{k,q} Z_q$ account for dose-covariate interactions, $\mu_k = (\mu_{k,1}, \dots, \mu_{k,m})$ are historical main treatment effects, $\xi_{k,j} Z = \xi_{k,j,1} Z_1 + \dots + \xi_{k,j,q} Z_q$ account for covariate interactions with the j th historical treatment, and

$\omega_E(x, \alpha_E)$ and $\omega_T(x, \alpha_T)$ are the usual dose-outcome functions characterizing main dose effects, and ω_E and ω_T may be quadratic or linear functions of x , as given earlier. For fitting the historical data, (10) takes the form

$$(11) \quad \eta_k(\tau_j, Z, \theta) = \mu_{k,j} + \beta_k Z + \xi_{k,j} Z$$

for $j = 1, \dots, m$ and $k = E, T$.

For fitting the data obtained during the trial, (10) is

$$(12) \quad \eta_k(x, Z, \theta) = \omega_k(x, \alpha_k) + \beta_k Z + x\gamma_k Z$$

for $k = E, T$.

A much more parsimonious model that accounts for dose-covariate interactions is obtained by replacing $x\gamma_k Z$ in (12) with either $\gamma_k\{\omega_k(x, \alpha_k) \times \beta_k Z\}$ or $\gamma_k\{x \times \beta_k Z\}$ where each γ_k is now a single parameter rather than a q -dimensional vector. This model requires only 2 dose-covariate interaction parameters instead of $2q$. This is motivated by the idea that $\gamma_k\{\omega_k(x, \alpha_k) \times \beta_k Z\}$ is similar to the one-degree of freedom interaction term in the model for a two-way layout with one observation per cell given by Tukey (1949). Unfortunately, in practice, this parsimonious model is a complete disaster since, using either $\omega_k(x, \alpha_k)$ or x , it gives a very poor fit to the trial data when dose-covariate interactions of any complexity are present. So this more parsimonious model is a cute idea that simply doesn't work.

Generalizing the EffTox design to accommodate Z requires much more than writing down a model. The set $\mathcal{A}_n(Z)$ of acceptable doses for a patient with covariates Z is defined to be all $x \in \mathcal{X}$ satisfying the constraints

$$(13) \quad \Pr\{\pi_E(x, Z, \theta) < \underline{\pi}_E(Z) | \mathcal{D}_n \cup \mathcal{H}\} < p_E^* \quad \text{and} \\ \Pr\{\pi_T(x, Z, \theta) > \bar{\pi}_T(Z) | \mathcal{D}_n \cup \mathcal{H}\} < p_T^*,$$

where $\underline{\pi}_E(Z)$ and $\bar{\pi}_T(Z)$ are *acceptability bounding functions*, constructed as follows. First, a representative set of covariate vectors, $\{Z^{(1)}, \dots, Z^{(K)}\}$, is determined. For each $Z^{(j)}$, the physician specifies the smallest probability of efficacy, $\underline{\pi}_E^{(j)}$, and the largest probability of toxicity, $\bar{\pi}_T^{(j)}$, that are acceptable for a patient having those covariates. For $k = E, T$, denote $\zeta_k(Z) = E(\beta_k Z | \mathcal{H})$, the historical posterior mean of the covariate main effect linear combination. To construct the bounding function $\underline{\pi}_E(Z)$ for $\pi_E(x, Z, \theta)$, the K pairs $(\zeta_E(Z^{(1)}), \underline{\pi}_E^{(1)})$, \dots , $(\zeta_E(Z^{(K)}), \underline{\pi}_E^{(K)})$ of estimated linear terms and elicited lower bounds on

π_E are used as regression data to fit a simple linear or quadratic curve by least squares, using $\zeta_E(Z^{(j)})$ as the predictor and $\underline{\pi}_E^{(j)}$ as the outcome variable. Denoting the estimated outcome under the fitted regression model by $\hat{\underline{\pi}}_E(\zeta_E(Z))$, the *efficacy lower bounding function* is $\underline{\pi}_E(Z) = \hat{\underline{\pi}}_E \circ \zeta_E(Z)$. The *toxicity upper bounding function* $\bar{\pi}_T(Z) = \hat{\bar{\pi}}_T \circ \zeta_T(Z)$ is computed similarly from $(\zeta_T(Z^{(1)}), \bar{\pi}_T^{(1)})$, \dots , $(\zeta_T(Z^{(K)}), \bar{\pi}_T^{(K)})$. When constructing these functions, it is important to plot the scattergrams of the constructed regression data sets along with the fitted curves, which the physician may use to guide adjustment of some $\underline{\pi}_E^{(j)}$ or $\bar{\pi}_T^{(j)}$ values, if desired, to obtain acceptability bounding functions $\underline{\pi}_E(Z)$ and $\bar{\pi}_T(Z)$ that make sense clinically. These constructions map each patient's Z vector into the probability bounds used in (13) to determine whether each $x \in \mathcal{X}$ is acceptable for that patient. To define a covariate-specific dose desirability index, we evaluate $\delta(p)$ given by (9) at $p = \mu^{(n)}(x, Z) = (E\{\pi_E(x, Z, \theta) | \mathcal{D}_n\}, E\{\pi_T(x, Z, \theta) | \mathcal{D}_n\})$, and denote this by $\delta_n(x, Z)$. For two patients with different covariates $Z_1 \neq Z_2$, it may be the case that $\mathcal{A}_n(Z_1) \neq \mathcal{A}_n(Z_2)$, including the possibility that $\mathcal{A}_n(Z) = \emptyset$ for one patient but not the other. Even if $\mathcal{A}_n(Z_1) = \mathcal{A}_n(Z_2)$, the x that maximizes $\delta_n(x, Z_1)$ may not be the same as that maximizing $\delta_n(x, Z_2)$.

Figure 3 illustrates how the dose-efficacy and dose-toxicity probability curves in x also may change with Z . The curves are given for a particular fixed θ^{true} in which the interactive effects of x and Z are substantial, taken from the acute leukemia application discussed by Thall, Nguyen and Estey (2008) where $Z = (\text{AGE}, \text{cytogenetic abnormality})$, with the second covariate coded as a three-category variable having possible values {Good, Intermediate, Poor} defined in terms of prognostic level. The rows in Figure 3 correspond to three different Z values. In the left column, the probabilities $\pi_E(x, \theta)$ are represented by circles and $\pi_T(x, \theta)$ by triangles, with an open (filled) circle or triangle representing an unacceptable (acceptable) dose. The corresponding desirabilities are given in the right column, obtained by evaluating $\delta(\pi_E(x, Z, \theta^{true}), \pi_T(x, Z, \theta^{true}))$. The figure shows that the dose-outcome functions $\pi_E(x, \theta)$ and $\pi_T(x, \theta)$ may change dramatically with Z , that the effect of prognosis may be as large as or larger than that of dose, and that interactions between x and Z may be quite important. Figure 3 also illustrates how the

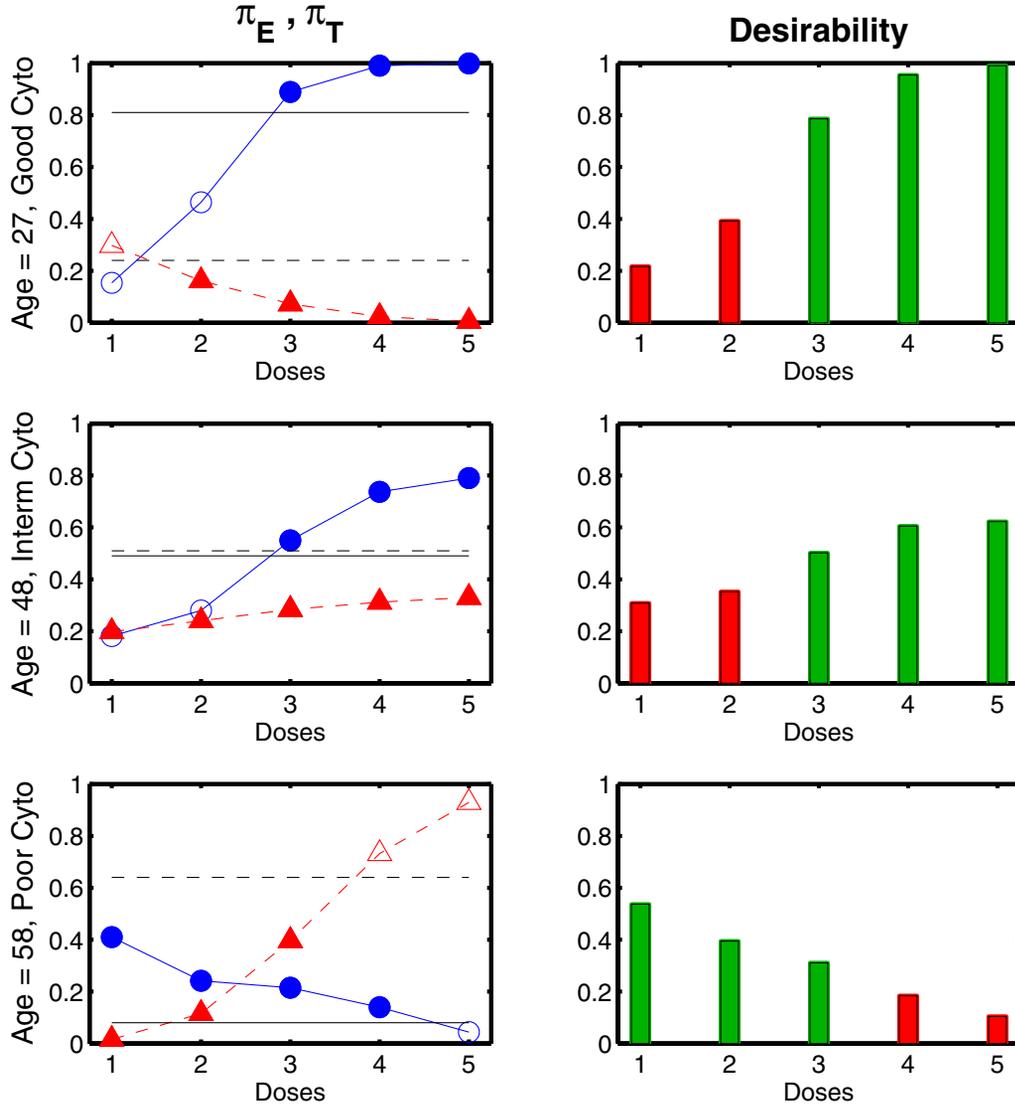


FIG. 3. Marginal outcome probabilities $\pi_E(x, \theta)$, given as circles, and $\pi_T(x, \theta)$, given as triangles (left column) and corresponding desirabilities (right column) for each of three different patient prognostic vectors, under a fixed θ . Solid (open) points correspond to acceptable (unacceptable) doses.

desirability function δ reduces each two-dimensional (p_E, p_T) to a one-dimensional value that may be used to compare doses for each Z .

To apply this methodology, the first step is to analyze \mathcal{H} under several models, choose the model providing the best fit, compute $p(\beta, \psi | \mathcal{H})$, and determine noninformative priors on α and γ . During the trial, when a patient with covariates Z is enrolled, $\mathcal{A}_n(Z)$ is computed. If $\mathcal{A}_n(\mathbf{Z}) = \phi$, the patient is not treated on protocol. If $\mathcal{A}_n(\mathbf{Z}) \neq \phi$, the patient is treated with the dose x maximizing $\delta_n(x, Z)$. If $\mathcal{A}_n(Z^{(j)}) = \phi$ for all representative covariates, then the trial is stopped. After the trial, given final data \mathcal{D}_N , the decision rules

based on $p(\theta | \mathcal{H} \cup \mathcal{D}_N)$ are used to select doses for future patients.

Our computer simulation studies of this new methodology have produced some disquieting messages. The first is that ignoring established prognostic covariates may lead to either very unsafe or very ineffective dose assignments for many patients both during and after a phase I or phase I/II trial. The second message is that, if dose-covariate interactions are present, ignoring them by using an additive model for the effects of x and Z also may lead to very poor dose assignments. That is, the common practice of ignoring known patient heterogeneity in early phase trials may lead to bad science and bad clinical practice.

5. ACCOUNTING FOR MULTIPLE TOXICITIES

5.1 Outcomes and Model

Bekele and Thall (BT, 2004) proposed a dose-finding method based on a vector $Y = (Y_1, \dots, Y_J)$ of several qualitatively different types of toxicity, with Y_j an ordinal variable recording the j th toxicity's severity. The method was motivated by a phase I trial to choose a dose of gemcitabine, in mg/m², from {100, ..., 1000} when combined with a fixed dose of 50 cGy external beam radiation, both given prior to surgery, for patients with soft tissue sarcoma. The design was developed working with a team of three oncologists who had extensive experience treating sarcomas. The point of departure from conventional methods is that the design distinguishes between different types of toxicity, and it also accounts for the severity levels of each.

Denote the $m_j + 1$ severity levels of Y_j by $\{y_{j,0}, y_{j,1}, \dots, y_{j,m_j}\}$. For example, in the sarcoma trial the 4 levels of liver toxicity were $y_{j,0} = \{\text{grade 0 or 1}\}$, $y_{j,1} = \{\text{grade 2}\}$, $y_{j,2} = \{\text{grade 3}\}$ and $y_{j,3} = \{\text{grade 4}\}$. Binary Y_j corresponds to $m_j = 1$. Using standardized doses $x = \log\{(\text{raw dose})/1000\}$, so that $\mathcal{X} = \{-2.30, -1.61, \dots, 0\}$, the distribution of $Y|x$ was modeled using the method of Albert and Chib (1993), in terms of the J -vector of Gaussian latent variables $\zeta = (\zeta_1, \dots, \zeta_J)$ with $E(\zeta_j|x) = \beta_{j,0} + x\beta_{j,1}$, $\text{var}(\zeta_j) = 1$ and correlation matrix Ω , by defining $Y_j = y_{j,k}$ if $\gamma_{j,k} \leq \zeta_j < \gamma_{j,k+1}$ for $k = 0, 1, \dots, m_j$ and $j = 1, \dots, J$ for cut-off parameters $\gamma_j = (\gamma_{j,1}, \dots, \gamma_{j,m_j})$ satisfying $-\infty = \gamma_{j,0} < \gamma_{j,1} < \dots < \gamma_{j,m_j} < \gamma_{j,m_j+1} = +\infty$, with $\gamma_{j,1} \equiv 0$ to ensure identifiability. This formulation greatly facilitates MCMC computations used to obtain posterior quantities. Denoting the $2J$ -vector of regression parameters $\beta = (\beta_{1,0}, \beta_{1,1}, \dots, \beta_{J,0}, \beta_{J,1})$, the vector $\gamma = (\gamma_1, \dots, \gamma_J)$ having $m_+ = m_1 + \dots + m_J$ entries, and the $J(J-1)/2$ off-diagonal elements of Ω by $\rho = (\rho_{1,2}, \rho_{1,3}, \dots, \rho_{J-1,J})$, the model parameter vector is $\theta = (\beta, \gamma, \rho)$. The marginal distribution of $Y_j|x$ is given by

$$\begin{aligned} \pi_{j,k}(x, \theta) &= \Pr(Y_j = y_{j,k}|x, \theta) \\ (14) \quad &= \Phi(\gamma_{j,k+1} - \beta_{j,0} - \beta_{j,1}x) \\ &\quad - \Phi(\gamma_{j,k} - \beta_{j,0} - \beta_{j,1}x). \end{aligned}$$

To obtain an expression for the joint distribution, denote the p.d.f. of a multivariate normal random vector W with mean vector μ and variance-covariance matrix Σ by $\phi_W(\cdot|\mu, \Sigma)$. In matrix notation, $E(\zeta|x) = X\beta'$, where X is the $J \times 2J$ block diagonal matrix with J identical blocks $(1 \ x)$. Denote the intervals $G_{j,k} =$

$(\gamma_{j,k}, \gamma_{j,k+1}]$. For observed vector $k = (k_1, \dots, k_J)$ of toxicity severity levels, the outcome is $Y = y(k) = (y_{1,k_1}, \dots, y_{J,k_J})$, which corresponds to latent ζ values in the J -dimensional set $G(k, \gamma) = G_{1,k_1} \times \dots \times G_{J,k_J}$. A single patient's likelihood contribution is

$$\begin{aligned} \mathcal{L}(Y|x, \theta) \\ (15) \quad &= \prod_{k_1=0}^{m_1} \dots \prod_{k_J=0}^{m_J} \left\{ \int_{G(k,\gamma)} \phi_Z(z|X\beta', \Omega) dz \right\}^{I[Y=y(k)]}. \end{aligned}$$

For priors, BT assume $\beta \sim N(\mu, \Sigma)$, subject to $\Pr(\beta_{j,1} > 0) = 1$ for all $j = 1, \dots, J$, so that β is $2J$ -variate normal with all slope coefficients truncated below at 0, but μ and Σ correspond to the untruncated $2J$ -variate normal. This ensures that $\Pr(Y_j > y_{j,k}|x, \beta) = 1 - \Phi\{\gamma_{j,k} - \beta_{j,0} - \beta_{j,1}x\}$ increases with x for each j and $k > 1$. For each j with $m_j \geq 2$ (3 or more levels), $\{\gamma_{j,2}, \dots, \gamma_{j,m_j}\}$ follow independent, uninformative priors on the domain $[0, 10]$, with each $p(\gamma_{j,k}) \propto 1$, subject to the constraint $0 < \gamma_{j,2} < \gamma_{j,3} < \dots < \gamma_{j,C_j}$, where the upper limit 10 on the support of each $p(\gamma_{j,k})$ was chosen for numerical convenience. The $\rho_{j,k}$'s are assumed to be i.i.d. $N(0, 1000)$, truncated to have support $[-1, +1]$, with Ω positive definite.

5.2 Total Toxicity Burden and Trial Conduct

The dose-finding method is based on *toxicity severity weights*, elicited as follows. The oncologists are first asked to specify the J toxicities to be monitored, including the severity levels of each. They are then asked to specify a numerical severity weight for each level of each toxicity within a positive-valued numerical range with which they are comfortable, such as 0 to 10, or 0 to 100. The severity weights are denoted by $w = (w_1, \dots, w_J)$, where $w_j = (w_{j,0}, w_{j,1}, \dots, w_{j,m_j})$ are the severity weights of the possible values $(y_{j,0}, y_{j,1}, \dots, y_{j,m_j})$ of Y_j , with the obvious requirement $w_{j,0} < w_{j,1} < \dots < w_{j,m_j}$; otherwise, if $w_{j,k} = w_{j,k+1}$, then levels k and $k + 1$ of Y_j should be combined. The elicited severity weights used in the sarcoma trial are illustrated in Figure 4. An interesting practical point arose while assigning severity weights to myelosuppression, which is defined in terms of low blood cell counts and is caused by effects of chemotherapy on the bone marrow. At first, no distinction was made between myelosuppression occur-

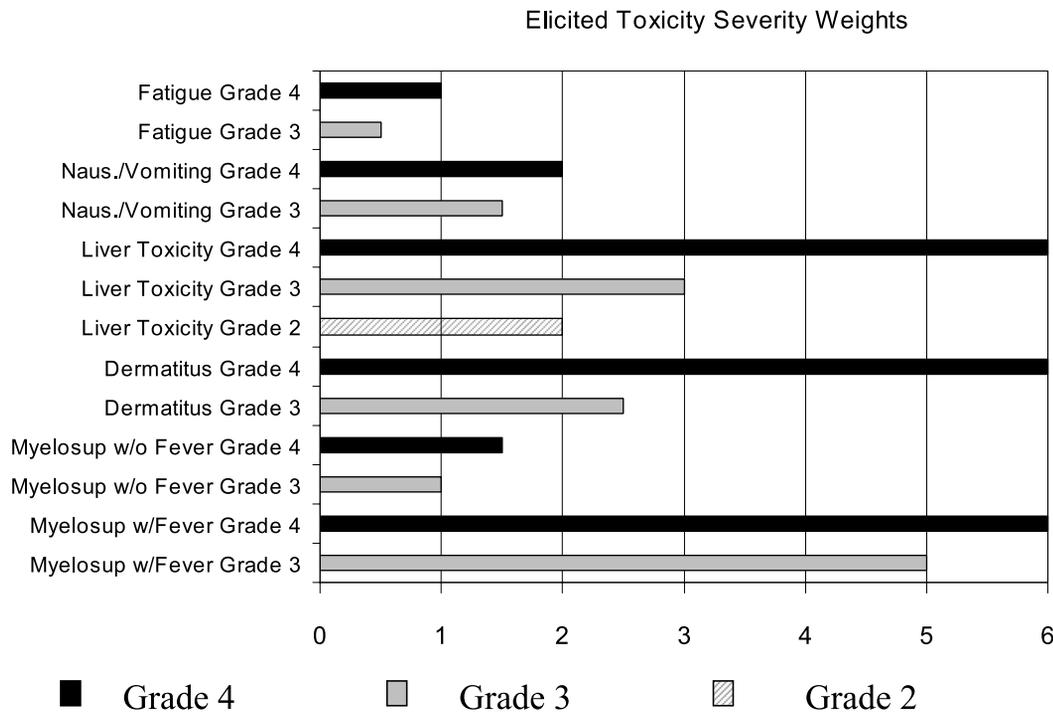


FIG. 4. The elicited toxicity severity weights used in the soft tissue sarcoma trial.

ring either with or without fever. During the process of establishing w , however, the oncologists explained that myelosuppression is much more severe when it occurs with fever, since it is then life-threatening and may be an impediment to further chemotherapy. This led us to redefine myelosuppression occurring without or with fever as two different types of toxicity. Figure 4 shows the severity weights 1 and 1.5 for grade 3 and 4 myelosuppression without fever, compared to weights 5 and 6 for grade 3 and 4 myelosuppression with fever. Thus, in general, Y and w are elicited together, and this process is not unlikely to involve iteration.

A patient's total toxicity burden (TTB) is defined to be

$$(16) \quad TTB = \sum_{j=1}^J \sum_{k=1}^{m_j} w_{j,k} I(Y_j = y_{j,k}).$$

For example, from Figure 4, a patient with grade 3 fatigue, grade 3 nausea/vomiting and grade 4 myelosuppression without fever would have $TTB = 0.5 + 1.5 + 1.5 = 3.5$, whereas a patient with grade 4 myelosuppression with fever would have $TTB = 6.0$. Using the conventional approach of defining a single binary outcome Y indicating at least one grade 3 or 4 toxicity, these two patients would be scored identically, with both having $Y = 1$.

The posterior expected TTB of dose x is

$$(17) \quad \begin{aligned} \tau(x, \mathcal{D}_n) &= E(TTB|x, \mathcal{D}_n) \\ &= \sum_{j=1}^J \sum_{k=1}^{m_j} w_{j,k} E\{\pi_{j,k}(x, \theta) | \mathcal{D}_n\}. \end{aligned}$$

The trial is conducted by establishing a targeted total toxicity burden, TTB^* , and choosing each cohort's x to minimize $|\tau(x, \mathcal{D}_n) - TTB^*|$. This is analogous to choosing a dose, based on a binary Y with $\pi(x, \theta) = \Pr(Y = 1|x, \theta)$, using the CRM criterion to minimize $|E\{\pi(x, \theta) | \mathcal{D}_n\} - \pi^*|$ for given target probability π^* . It is easy to show that, since $w_{j,k-1} < w_{j,k}$ and $\beta_{j,1} > 0$ for all j and k , $\tau(x, \mathcal{D}_n)$ is increasing in x , so x may be determined by a monotone search. The process proposed by BT for establishing the target TTB^* is straightforward, albeit somewhat elaborate. The physicians are first asked to specify a set of hypothetical patient cohorts and toxicity outcomes for each patient in each cohort, with the cohorts defined so that the toxicity severities vary substantially between cohorts. BT provide a detailed description of this process, and in the sarcoma trial there were 16 hypothetical cohorts of 4 patients each with the mean TTB of each cohort varying from $\overline{TTB} = 1.25$ to 5.62. For each hypothetical cohort, the oncologists are asked whether observing its toxicity outcomes would lead them to

escalate, repeat the current dose, or de-escalate for the next cohort. The target TTB^* is then defined as the mean of the \overline{TTB} values for which the decision would be to repeat the current dose. For the sarcoma trial, this yielded $TTB^* = 3.04$. Computer simulations of this methodology provided by BT show that it has remarkably attractive OCs and makes decisions very differently from conventional phase I designs. A cohort of four patients all with myelosuppression grade 4 without fever, patients #1, #2 and #3 with grade 3 fatigue, and patient #4 with grade 3 nausea/vomiting would have $\overline{TTB} = \{(1.5 + 0.5) + (1.5 + 0.5) + (1.5 + 0.5) + (1.5 + 1.5)\}/4 = 2.25$. The three oncologists all agreed that the appropriate decision based on these outcomes would be to escalate, whereas any conventional method based on one binary toxicity indicator would score this as 4 “toxicities” in 4 patients and certainly would de-escalate.

6. OPTIMIZING DOSE AND SCHEDULE

6.1 A New Paradigm for Phase I Trials

Braun et al. (BTND, 2007) proposed a new paradigm for phase I trials that jointly optimizes schedule of administration and per-administration dose (PAD) based on time-to-toxicity. This extends Braun, Yuan and Thall (2005), who optimized schedule while assuming a fixed PAD. Although the model used by BTND is very different from that underlying the TITE CRM (Cheung and Chappell, 2000) for dose-finding based on time-to-toxicity, the BTND method is a practical extension in that it allows schedule as well dose to be varied. The treatment regime is $x = (s, d_s)$, where $s = (s_1, \dots, s_k)$ are successive administration times and $d_s = (d(s_1), \dots, d(s_k))$ are the doses given at those times. BTND address the problem of evaluating a $K \times J$ matrix of K nested schedules, $s^{(1)} \subset s^{(2)} \subset \dots \subset s^{(K)}$, where the k th schedule is $s^{(k)} = (s_1, s_2, \dots, s_{m^{(k)}})$, so that $m^{(1)} < m^{(2)} < \dots < m^{(K)}$, and J PADs, $d^{(1)} < d^{(2)} < \dots < d^{(J)}$. The treatment set evaluated by the design is $\mathcal{X} = \{(s^{(k)}, d^{(j)}) : k = 1, \dots, K, j = 1, \dots, J\}$, and the total amount of the agent given to the patient increases with both dose and schedule. For example, a patient assigned PAD $d^{(3)}$ under schedule $s^{(2)} = (s_1, s_2, \dots, s_{m^{(2)}})$ receives total dose $d^{(3)}m^{(2)}$ of the agent in $m^{(2)}$ successive administrations of $d^{(3)}$ each, unless therapy is terminated early due to toxicity, so the planned $d_{s^{(2)}}$ in $x = (s^{(2)}, d_{s^{(2)}})$ is the $m^{(2)}$ -vector with all entries $d^{(3)}$.

For this regime, it is helpful to distinguish between two time scales, *study time* and *patient time*. Starting at study time 0 when the trial begins, let e be a

given patient’s entry time, so that the patient’s assigned schedule s is administered at study times $e + s = (e + s_1, \dots, e + s_k)$. Denote a patient’s time from entry at e to toxicity by T , so that at study time t the patient’s observed time to toxicity or last follow-up is $T^o(t) = T$ if $e + T \leq t$ and $T^o(t) = t - e$ if $e + T > t$. Defining $\delta(t) = I(e + T \leq t)$, the patient’s outcome data at study time t are $Y(t) = (T^o(t), \delta(t))$. The probability model is constructed from the patient’s hazard of toxicity, $h(u|d, \theta)$, associated with a single administration of dose d given u days previously, and we denote $H(x|d, \theta) = \int_0^x h(u|d, \theta) du$. Under the assumption that effects of successive administrations of the agent are additive, the overall hazard of toxicity at study time t for a patient entering at e and treated with $x = (s, d_s)$ is

$$(18) \quad \lambda(t|e, (s, d_s), \theta) = \sum_{j=1}^k h(t - e - s_j|d(s_j), \theta),$$

where $h(u|d, \theta) = 0$ for all $u < 0$. The patient’s cumulative hazard of toxicity at study time t is thus

$$(19) \quad \Lambda(t|e, (s, d_s), \theta) = \sum_{j=1}^k H(t - e - s_j|d(s_j), \theta),$$

and the probability that the patient has not had toxicity by study time t is $\Pr(e + T > t|e, (s, d_s), \theta) = \exp\{-\Lambda(t|e, (s, d_s), \theta)\}$. Thus, h and H are expressed in terms of patient time, whereas λ and Λ are expressed in terms of study time. The probability distribution of T is determined by the particular form of the single administration hazard function h .

The model allows each patient’s actual $x = (s, d_s)$ received to fall outside the set of KJ treatment configurations in \mathcal{X} , provided that each dose in d_s is an element of $\{d^{(1)}, \dots, d^{(K)}\}$. In particular, the elements of d_s need not be identical. This accommodates the possibility that a patient’s treatment does not go as planned, for example, due to interim dose reductions following moderate toxicity or deviations from the planned schedule. It also allows the possibility that the patient’s planned x may be changed before the schedule is completed, based on other patients’ data observed during the patient’s therapy.

At study time t , let $x_i(t)$ denote the portion of the i th patient’s treatment regime x_i that has been administered by that time and let $\mathcal{D}_t = \{(T_i^o(t), \delta_i(t), e_i, x_i(t)), i = 1, \dots, n(t)\}$ denote the current data. The likelihood at study time t is

$$(20) \quad \mathcal{L}(\mathcal{D}_t|\theta) = \prod_{i=1}^{n(t)} \{\lambda(T_i^o(t)|e_i, x_i(t), \theta)\}^{\delta_i(t)} \cdot \exp\{-\Lambda(T_i^o(t)|e_i, x_i(t), \theta)\}.$$

BTND assume that, for each $j = 1, \dots, J$, the single-administration hazard function associated with dose $d^{(j)}$ is a triangle, formally

$$(21) \quad h(u|d^{(j)}, \theta_j) = \frac{2\alpha_j u}{(\beta_j + \gamma_j)\beta_j} I(0 \leq u \leq \beta_j) + \frac{2\alpha_j(\beta_j + \gamma_j - u)}{(\beta_j + \gamma_j)\gamma_j} \cdot I(\beta_j < u \leq \beta_j + \gamma_j),$$

where $\theta_j = (\alpha_j, \beta_j, \gamma_j)$, so that $\theta = (\theta_1, \dots, \theta_J)$ has $3J$ elements. The j th triangle has base of length $\beta_j + \gamma_j$, area α_j , and maximum height $h(\beta_j|d^{(j)}, \theta_j) = 2\alpha_j/(\beta_j + \gamma_j)$. Thus, for $u \geq \beta_j + \gamma_j$, the cumulative single-administration hazard is $H_j(u|d^{(j)}, \theta_j) = \alpha_j$. Under this model, since any schedule (s_1, \dots, s_k) is finite, given θ and k -vector of PADs $d_s = (d^{(1)}, \dots, d^{(k)})$, the cumulative hazard $\Lambda(t|(s, d_s), \theta)$ has the finite maximum value $k\alpha_j$. Consequently, given θ , the probability that the patient never experiences toxicity is $\bar{F}(t|(s, d_s), \theta) = \exp(-k\alpha_j)$ for all $t > s_k + \beta_j + \gamma_j$, with the obvious elaboration of the upper limit on $\bar{F}(t|(s, d_s), \theta)$ if the elements of d_s are not identical.

The triangular form of h may seem to be an oversimplification of a complex phenomenon. In application, however, it is quite flexible and yields a very robust trial design. Figure 5 shows the cumulative hazard of toxicity for a patient treated with a fixed PAD according to the 4-administration schedule $s = (0, 3, 10, 13)$. The shaded area represents $H(12)$, the cumulative hazard of toxicity by day 12. The smoothness of the curve $H(u)$ for $0 \leq u \leq 40$ that results from summing four

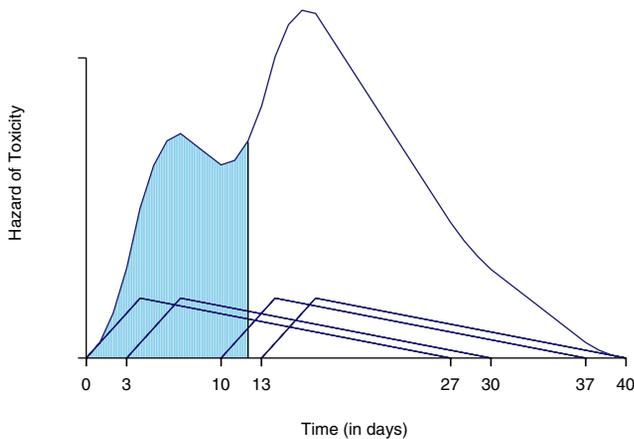


FIG. 5. Illustration of triangular component hazards for administrations on days 0, 3, 10 and 13, and the resulting cumulative hazard function. The shaded region is the cumulative hazard, $H(12)$, of toxicity by day 12.

triangles and the fact that the parameters $(\alpha_j, \beta_j, \gamma_j)$ characterizing the j th triangle corresponding to $d^{(j)}$ are estimated from the accumulating data together provide an intuitive motivation for the model's flexibility and robustness. This was borne out by the extensive simulation studies reported by BTND. In the setting where only schedule is varied with PAD fixed, Liu and Braun (2009) have studied the use of a smooth component hazard function, a 2-parameter Weibull, $h(u|\alpha, \beta) = e^\beta \alpha u^{\alpha-1} \exp(-u^\alpha e^\beta)$ for $u > 0$, where $\alpha > 0$ and β is real-valued. This allows h to be non-monotone if $\alpha \geq 2$ or decreasing if $0 \leq \alpha < 2$.

6.2 Trial Conduct

Given an interval $[0, t^*]$ large enough to reliably evaluate T under the longest schedule, the physician specifies a target $\pi^* = \Pr(T \leq t^*)$. For brevity, I will temporarily index a patient's assigned treatment $(s^{(k)}, d^{(j)})$ by (k, j) , and denote the c.d.f. of T associated with $x = (k, j)$ by $F_{k,j}(\theta) = \Pr(T \leq t^*|(k, j), \theta)$. Since $\mathcal{L}(\mathcal{D}_t|\theta)$ and hence the posterior $p(\theta|\mathcal{D}_t)$ change continuously with t during the trial, necessarily, $x \in \mathcal{X}$ is chosen for each newly enrolled patient, that is, $c = 1$. A patient accrued at study time t is assigned the pair $(k, j) \in \mathcal{X}$ minimizing the objective function $|\mathbb{E}\{F_{k,j}(\theta)|\mathcal{D}_t\} - \pi^*|$, similar to the CRM. Assignment of x using this criterion is subject to the following two safety rules. Given a maximum toxicity probability, F_{max} , specified by the physician, the schedule-dose pair (k, j) is *acceptable* if $\Pr(F_{k,j}(\theta) > F_{max}|\mathcal{D}_t) < p^*$, where p^* is a fixed upper cut-off such as 0.80 or 0.90. If no pair in \mathcal{X} is acceptable, the trial is stopped. This is similar to the toxicity portion of the acceptability criteria (8) of the EffTox methodology. The second safety rule is that escalation from (k, j) is restricted in that no untried dose-schedule combination may be skipped, specifically the next patient may be treated at $x = (k+1, j)$, $(k, j+1)$ or $(k+1, j+1)$, but at no higher pair. There is no such constraint on de-escalation. While developing this methodology, we initially tried using the more restrictive constraint that does not allow diagonal escalation, from (k, j) to untried $(k+1, j+1)$, but this yielded a design with very poor properties. This is the case essentially because this constraint makes exploration of the 2-dimensional set of KJ schedule-PAD pairs unfeasible, and, in fact, it provides no additional measure of safety. While BTND recommended that the first patient be treated at the safest pair $(k, j) = (1, 1)$, in practice, the physician might wish to start at $(1, 2)$, $(2, 1)$ or $(2, 2)$.

It may seem self-evident that this method is greatly superior to any comparable method that fixes schedule and only varies dose, since an optimal combination $(s^{(k)}, d^{(j)})$ is simply ignored if the fixed schedule is not $s^{(k)}$. The simulations reported by BTND clearly illustrate this point. Currently, however, it still is standard practice in phase I trials to guess what schedule might be best, possibly based on animal data, and proceed in humans by varying only dose. As described by BTND, this new methodology has been used to conduct an allotx trial of the post-transplant agent 5-azacitidine, which is thought to kill leukemia cells by reactivating tumor suppressor genes while also enhancing graft-versus-leukemia effect.

7. DISCUSSION

Each of the designs reviewed here includes one or more more aspects of treatment or outcome in an early phase trial that are ignored by standard designs. The price of accommodating such additional complexity is a much more structured model and method. This often requires substantially more work for trial design and conduct, including analysis of historical data, elicitation of priors and design parameters, development of computer software, carrying out simulations to calibrate design parameters and establish operating characteristics, and the difficult process of real time data monitoring during trial conduct. In each case, however, the design provides advantages over standard methods so large that the comparisons may seem unfair. Evidently, accounting for both anti-disease effect and toxicity is a good idea, ignoring covariates is a bad idea, quantifying the clinical importance of different types and grades of toxicities is a good idea, and ignoring schedule effects is a bad idea. For example, the optimal treatment pair determined by the design in the 5-azacitidine trial was (40 mg/m² per administration, 3 cycles), which simply could not have been found using a conventional dose-finding method that fixes schedule at 1 cycle. While each of the designs relies on a model with a nontrivial number of parameters, which in turn often requires elaborate prior specification and sophisticated numerical methods, the amount of information per patient also is much greater. The final questions are whether such designs have good properties, which the computer simulations show they do, and whether they can be implemented in practice, which has been the case for all of the designs discussed here.

For most early phase trials, conventional methods for determining sample size based on hypothesis testing or

estimation may be of little use. To determine a planned maximum sample size for a phase I or I/II trial, I ask the physicians the anticipated accrual rate, which often is a range of values, the desired maximum trial duration, and cost or other resource limitations, such as the amount of a specialized agent that feasibly can be produced in the laboratory. For each of several feasible maximum sample sizes, I simulate the trial and also compute posterior estimates of important parameters based on illustrative data sets. I then show these results to the physician and ask him/her to choose a maximum sample size on that basis. If the largest feasible sample size does not yield a reasonably reliable design, I recommend that the trial not be conducted. Using this practical approach, in my experience planned phase I or I/II sample sizes usually range from 24 to 60.

The most severe difficulties in achieving widespread implementation of outcome-adaptive methods in early phase clinical trials are computational and sociological. The first practical requirement is portable, high quality computer software for implementation, including statistical programs that perform the necessary computations for specific methods and, ideally, graphical user interfaces that communicate with established databases and statistical programs to facilitate real-time data entry and computation of adaptive decision criteria during the trial. An “elephant in the living room” of outcome-adaptive methods for clinical trials is that constructing and implementing such information systems in medical environments is often much more difficult, expensive and time-consuming than developing a particular statistical method. Moreover, once such a system is in place, the process of entering the patient data required by an outcome-adaptive method is time-consuming and potentially error prone.

A natural question is whether one can construct practical designs that address the problems that arose in the SCT trial described in the [Introduction](#). Such designs would optimize multiple schedule-dose combinations of several agents used in combination based on a vector of appropriately chosen efficacy and toxicity outcomes, possibly accounting for patient covariates and schedule-dose-covariate interactive effects on the outcomes, while adaptively choosing patient-specific schedule-dose combinations in real time. This also would require a decision criterion based on multiple outcomes, possibly using either efficacy-toxicity trade-offs or numerical utilities (Houede et al., 2010). Currently, we are working to develop new designs that include various combinations of these extensions. In my experience, however, early phase trials are so complex

that most trials cannot be optimally designed until after they already have been carried out, and a “one size fits all” design simply does not exist.

Clinical trials are viewed very differently by the pharmaceutical companies who produce and supply new agents, by regulatory agencies, by institutional review boards, by administrators who provide infrastructure and resources for trial conduct, by the physicians and nurses who actually treat the patients in a trial, and by the patients themselves. Individuals with decision-making authority in all of these different groups must agree on a trial design before a trial may be conducted. Many of these people regard the structure and properties of a particular statistical design as technicalities too complicated to understand and at most marginally relevant. Most early phase trials are conducted using very simple conventional methods that do not require computers and are easy to implement. Accounting more fully for the complexities of both the actual treatment regimes and the patients’ clinical outcomes is a double edged sword, since the greater safety and reliability that such methods provide is obtained only by working much harder in both design formulation and trial conduct. Physicians who understand the advantages of properly constructed outcome-adaptive designs and want to use such methods are a minority, although they often provide the initial motivation for developing new statistical designs. Their desire to use outcome-adaptive methods, and the recent shift in the pharmaceutical community, at least among statisticians, to embrace all things “adaptive” in clinical trials seem to be harbingers of a different future. How this may actually translate into practical reality in the coming years remains to be seen.

The website <http://biostatistics.mdanderson.org/SoftwareDownload> contains computer programs for implementing the methods described in Sections 2 (ToxFinder), 3.1 and 3.2 (EffTox) and 5 (Dose Schedule Finder).

ACKNOWLEDGMENT

This research was partially supported by NCI Grant 2RO1 CA083932.

REFERENCES

ALBERT, J. H. and CHIB, S. (1993). Bayesian analysis of binary and polytomous response data. *J. Amer. Statist. Assoc.* **88** 669–679. [MR1224394](#)

- ANDERSSON, B. S., THALL, P. F., MADDEN, T., COURIEL, D., WANG, X., TRAN, H. T., ANDERLINI, P., DE LIMA, M., GAJEWSKI, J. and CHAMPLIN, R. E. (2002). Busulfan systemic exposure relative to regimen-related toxicity and acute graft vs. host disease; defining a therapeutic window for IVBuCy2 in chronic myelogenous leukemia. *Biology of Blood and Marrow Transplantation* **8** 477–485.
- BABB, J. S. and ROGATKO, A. (2001). Patient specific dosing in a phase I cancer trial. *Stat. Med.* **20** 2079–2090.
- BEKELE, B. N. and THALL, P. F. (2004). Dose-finding based on multiple toxicities in a soft tissue sarcoma trial. *J. Amer. Statist. Assoc.* **99** 26–35. [MR2061885](#)
- BEKELE, B. N., JI, Y., SHEN, Y. and THALL, P. F. (2008). Monitoring late onset toxicities in phase I trials using predicted risks. *Biostatistics* **9** 442–457.
- 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.
- BRAUN, T. M., YUAN, Z. and THALL, P. F. (2005). Determining a maximum tolerated schedule of a cytotoxic agent. *Biometrics* **61** 335–343. [MR2140904](#)
- BRAUN, T. M., THALL, P. F., NGUYEN, H. and DE LIMA, M. (2007). Simultaneously optimizing dose and schedule of a new cytotoxic agent. *Clinical Trials* **4** 113–124.
- CHEUNG, Y. K. (2005). Coherence principles in dose-finding studies. *Biometrika* **92** 863–873. [MR2234191](#)
- CHEUNG, Y. and CHAPPELL, R. (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56** 1177–1182. [MR1815616](#)
- DETTE, H., BRETZ, F., PEPELYSHEV, A. and PINHIRO, J. (2008). Optimal designs for dose-finding studies. *J. Amer. Statist. Assoc.* **103** 1225–1237. [MR2462895](#)
- GOODMAN, S. G., ZAHURAK, M. L. and PIANTADOSI, S. (1995). Some practical improvements in the continual reassessment method. *Stat. Med.* **14** 1149–1161.
- 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.
- HAINES, L. M., PEREVOZSKAYA, I. and ROSENBERGER, W. F. (2003). Bayesian optimal designs for phase I clinical trials. *Biometrics* **59** 591–600. [MR2004264](#)
- HORSTMANN, E., MCCABE, M. S., GROCHOW, L., YAMAMOTO, S., RUBINSTEIN, L., BUDD, T., SHOEMAKER, D., EMANUEL, E. J. and GRADY, C. (2005). Risks and benefits of phase I oncology trials, 1991 through 2002. *New England J. Medicine* **352** 895–904.
- HOUEDE, N., THALL, P. F., NGUYEN, H., PAOLETTI, X. and KRAMAR, A. (2010). Utility-based optimization of combination therapy using ordinal toxicity and efficacy in phase I/II trials. *Biometrics*. In press.
- IVANOVA, A. (2003). A new dose-finding design for bivariate outcomes. *Biometrics* **59** 1001–1007. [MR2025124](#)
- LIU, C. A. and BRAUN, T. M. (2009). Parametric non-mixture cure models for schedule-finding of therapeutic agents. *Appl. Statist.* **58** 225–236.
- MORITA, S., THALL, P. F. and MÜLLER, P. (2008). Determining the effective sample size of a parametric prior. *Biometrics* **64** 595–602. [MR2432433](#)

- O'QUIGLEY, J. (1990). Sequential design and analysis of dose-finding studies in patients with life threatening disease. *Fund. Clin. Pharmacology* **4** (Suppl. 2) 81s–91s.
- O'QUIGLEY, J., HUGHES, M. D. and FENTON, T. (2001). Dose-finding designs for HIV studies. *Biometrics* **57** 1018–1029. [MR1973811](#)
- 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. [MR1059105](#)
- PALMER, C. R. (2002). Ethics, data-dependent designs, and the strategy of clinical trials: Time to start learning-as-we-go? *Stat. Methods Med. Res.* **5** 381–402.
- RATAIN, M. J., MICK, R., JANISCH, L., BEREZIN, F., SCHILSKY, R. L., VOGELZANG, N. J. and KUT, M. (1996). Individualized dosing of amonafide based on a pharmacodynamic model incorporating acetylator phenotype and gender. *Pharmacogenetics* **6** 93–101.
- THALL, P. F. and COOK, J. D. (2004). Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* **60** 684–693. [MR2089444](#)
- 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., COOK, J. D. and ESTEY, E. H. (2006). Adaptive dose selection using efficacy-toxicity trade-offs: Illustrations and practical considerations. *J. Biopharm. Statist.* **16** 623–638. [MR2252311](#)
- THALL, P. F., NGUYEN, H. and ESTEY, E. H. (2008). Patient-specific dose-finding based on bivariate outcomes and covariates. *Biometrics* **64** 1126–1136.
- THALL, P. F., LEE, J. J., TSENG, C.-H. and ESTEY, E. H. (1999). Accrual strategies for phase I trials with delayed patient outcome. *Stat. Med.* **18** 1155–1169.
- THALL, P. F., MILLIKAN, R. E., MÜLLER, P. and LEE, S.-J. (2003). Dose-finding with two agents in phase I oncology trials. *Biometrics* **59** 487–496. [MR2004253](#)
- THALL, P. F., SIMON, R. and ESTEY, E. H. (1995). Bayesian sequential monitoring designs for single-arm clinical trials with multiple outcomes. *Stat. Med.* **14** 357–379.
- TUKEY, J. W. (1949). One degree of freedom for non-additivity. *Biometrics* **5** 232–242.
- WHELAN, H. T., COOK, J. D., AMLIE-LEFOND, C. M., HOVINGA, C. A., CHAN, A. K., ICHORD, R. N., DEVEBER, G. A. and THALL, P. F. (2008). Practical model-based dose-finding in early phase clinical trials: Optimizing tissue plasminogen activator dose for treatment of ischemic stroke in children. *Stroke* **39** 2627–2636.