

# Determining a Maximum-Tolerated Schedule of a Cytotoxic Agent

Thomas M. Braun\* and Zheng Yuan

Department of Biostatistics, School of Public Health, University of Michigan, Ann Arbor,  
Michigan 48109, U.S.A.

\* *email*: tombraun@umich.edu

and

Peter F. Thall

Department of Biostatistics and Applied Mathematics, Box 447, The University of Texas,  
M. D. Anderson Cancer Center, Houston, Texas 77030, U.S.A.

**SUMMARY.** Most phase I clinical trials are designed to determine a maximum-tolerated dose (MTD) for one initial administration or treatment course of a cytotoxic experimental agent. Toxicity usually is defined as the indicator of whether one or more particular adverse events occur within a short time period from the start of therapy. However, physicians often administer an agent to the patient repeatedly and monitor long-term toxicity due to cumulative effects. We propose a new method for such settings. It is based on the time to toxicity rather than a binary outcome, and the goal is to determine a maximum-tolerated schedule (MTS) rather than a conventional MTD. The model and method account for a patient's entire sequence of administrations, with the overall hazard of toxicity modeled as the sum of a sequence of hazards, each associated with one administration. Data monitoring and decision making are done continuously throughout the trial. We illustrate the method with an allogeneic bone marrow transplantation (BMT) trial to determine how long a recombinant human growth factor can be administered as prophylaxis for acute graft-versus-host disease (aGVHD), and we present a simulation study in the context of this trial.

**KEY WORDS:** Bone marrow transplantation; CRM; Dose escalation; KGF; Maximum-tolerated dose; Phase I trial.

## 1. Introduction

In an allogeneic bone marrow transplant (BMT), a subject (host) receives the bone marrow (graft) donated by another subject. The graft contains crucial T and natural killer cells that coordinate a positive immune response to any residual leukemia, called a graft-versus-leukemia (GVL) effect, yet also coordinate a negative immune response against host tissue, called graft-versus-host disease (GVHD). Acute GVHD (aGVHD) leads to substantial damage to the host's skin, liver, and/or gastrointestinal (GI) tract and remains a leading cause of morbidity and mortality in BMT patients.

Preclinical studies have shown that recombinant human keratinocyte growth factor (KGF) markedly reduces chemotherapy- or radiation-induced injury to the mucosal lining of the lower GI tract (Farrell et al., 1998). These results suggest that KGF can shield the GI tract of allogeneic BMT recipients from the detrimental effects of aGVHD while preserving the beneficial GVL effect (Panoskaltis-Mortari et al., 1998; Krijanovski et al., 1999). However, KGF carries the risk of several toxicities, including skin-related events such as rash, flushing, and edema, as well as increases in amylase and lipase, both of which are indicators of pancreas dysfunction.

A phase I study in colorectal cancer patients found that 40 mg/kg of KGF could be safely administered on each of 3 consecutive days (Meropol et al., 2003). In a phase I study in BMT, each patient received 60 mg/kg of KGF on each of the 2 days prior to BMT, and on the day of BMT. After 4 days of rest with no KGF, the patient received KGF for 3 more days. Thus, KGF was administered using the 10-day schedule (3-days-on/4-days-off/3-days-on), which we denote by (3+, 4-, 3+). Toxicity was monitored for 28 days, motivated by the assumption that any adverse effect due to a single administration of KGF is certain to occur within 18 days. Although one course of KGF using the (3+, 4-, 3+) schedule is considered safe, investigators believed that this may not be sufficient prophylaxis for aGVHD, which may take up to roughly 100 days after BMT to develop. Consequently, the investigators wished to study the safety of multiple courses of KGF, with 4 days of rest between consecutive courses, and focus on toxicity associated with the entire period of therapy. Thus, two courses would consist of the 24-day schedule (3+, 4-, 3+, 4-, 3+, 4-, 3+) and so on.

Conventional phase I studies determine the maximum-tolerated dose (MTD) for a single administration or course by characterizing a patient's outcome as a binary variable

indicating whether toxicity occurs within a relatively short time period from the start of therapy. Generally, the MTD is considered the highest dose that does not present a practical limitation to therapy (Storer, 1989; Goodman, Zahurak, and Piantadosi, 1995; Babb, Rogatko, and Zacks, 1998). This approach has seen widespread use largely because it facilitates adaptive dose-finding methods that use the doses and outcomes of previous patients to select doses for successive patients. However, such methods are inadequate for trials where the agent is administered repeatedly over time and evaluation of long-term effects is important, because they (1) base dose finding on one, initial administration or course of therapy, and (2) require that toxicity be evaluated quickly enough to allow for adaptive dose assignment. One exception is the time to event continual reassessment method (TiTE CRM) (Cheung and Chappell, 2000), which evaluates long-term toxicity without delaying accrual, although the TiTE CRM does not accommodate settings where treatment continues past the first course.

Our proposed method uses patient's time to toxicity as the outcome, with the hazard of toxicity modeled as the sum of a sequence of hazards, each associated with one administration. We determine the maximum-tolerated schedule (MTS) that the patient may receive based on the risk of toxicity occurring within a specified follow-up period that includes the maximum schedule being considered. Patient accrual, data monitoring, and outcome-adaptive decision making are done continuously throughout the trial under a Bayesian formulation. Each time a new patient is accrued, the most recent data are used to evaluate criteria that define the optimal schedule, which is assigned to the new patient.

Section 2 establishes notation and presents the probability model. Section 3 describes a method for prior elicitation, and Section 4 describes the trial design. Section 5 describes the application of the method to the KGF trial, and presents a computer simulation study to assess the design's performance. We conclude with a discussion in Section 6.

## 2. Probability Model for Toxicity

### 2.1 Preliminary Notation

Let  $t^*$  denote any given time from the start of the trial when one evaluates the data and makes a decision,  $n^*$  the number of patients enrolled up to  $t^*$ , and  $e_i$  the study time when patient  $i$  is enrolled, for  $i = 1, \dots, n^*$ . Let  $T_i$  be the actual, possibly unobserved, time after study entry when patient  $i$  experiences toxicity. At study time  $t^*$ , we denote the amount of time that patient  $i$  has been observed by

$$Y_i = \begin{cases} T_i & \text{if } e_i + T_i \leq t^*, \\ t^* - e_i & \text{if } e_i + T_i > t^*, \end{cases}$$

and we define  $\delta_i = 1$  if  $Y_i = T_i$  and  $\delta_i = 0$  if  $Y_i \neq T_i$ , that is,  $\delta_i$  is the indicator that patient  $i$  has had toxicity by  $t^*$ .

Let  $\mathbf{s}_i = \{s_{i,1}, \dots, s_{i,m_i}\}$  denote the successive *patient* times at which the  $i$ th patient receives the agent, where  $s_{i,1}$  coincides with study entry. Under this general notation, the agent can be administered whenever and as frequently as desired to each patient, and an arbitrary number of different treatment sequences can be studied. This allows some of a patient's actual administration times to deviate from his/

her scheduled times. We assume that the investigators wish to study  $k$  treatment sequences,  $\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(k)}$ , where  $\mathbf{s}^{(j)} = (s_1, s_2, \dots, s_{m^{(j)}})$  and that the  $j$ th sequence has a total of  $m^{(j)}$  administrations. Furthermore,  $\mathbf{s}^{(j)}$  is a subsequence of  $\mathbf{s}^{(j+1)}$  for each  $j = 1, \dots, k - 1$ , so that the duration of therapy increases with  $j$  and  $m^{(1)} < m^{(2)} \dots < m^{(k)}$ . In the KGF trial, one course of the (3+, 4-, 3+) schedule corresponds to  $\mathbf{s}^{(1)} = (1, 2, 3, 8, 9, 10)$ , two courses correspond to  $\mathbf{s}^{(2)} = (1, 2, 3, 8, 9, 10, 15, 16, 17, 22, 23, 24) = (\mathbf{s}^{(1)}, \mathbf{s}^{(1)} + 14)$ , and so on, with BMT at day 3 in any case.

Let  $m_i$  denote the index of the last administration received by patient  $i$  at interim study time  $t^*$ . Although  $m^{(j)}$  administrations are planned for a patient assigned to schedule  $\mathbf{s}^{(j)}$ , at  $t^*$  it may be the case that  $m_i < m^{(j)}$  either due to administrative censoring or because patient  $i$  had toxicity at study time  $e_i + s_{i,m_i}$  and thus received no further administrations. Let  $\tau$  denote the fixed maximum length of follow-up for each patient. This should be chosen by the physicians for clinical reasons, but must be large enough to accommodate the longest sequence,  $\mathbf{s}^{(k)}$ . In the KGF trial,  $\tau = 100$  days. A fixed target probability,  $p_\tau$ , is elicited from the physician and is defined as the targeted threshold probability of toxicity at any time from enrollment to  $\tau$ .

### 2.2 Toxicity Distribution

Temporarily suppress the patient index  $i$ . Consider a single patient, and assume that the same dose of the agent is given at each administration. Let  $h(u | \boldsymbol{\theta})$  be the hazard of toxicity attributed to a single administration, where  $\boldsymbol{\theta}$  is the model parameter vector. We define the total hazard of toxicity at time  $t^*$  for a patient treated with schedule  $\mathbf{s}$  to be

$$\lambda(Y | \boldsymbol{\theta}, \mathbf{s}) = \sum_{\ell=1}^m h(Y - s_\ell | \boldsymbol{\theta}), \tag{1}$$

with  $h(u | \boldsymbol{\theta}) = 0$  if  $u < 0$ . Under this definition, the patient's risk of toxicity at  $t^*$  depends on three quantities: (1)  $Y$ , the patient's time on study, (2)  $m$ , the number of administrations received up to  $t^*$ , and (3)  $\mathbf{s}$ , the times at which the treatment was administered. We assume that the form of  $h(\cdot)$  does not change with successive administrations, although this assumption can be relaxed.

The patient's cumulative hazard function (CHF) at  $t^*$  is

$$\begin{aligned} \Lambda(Y | \boldsymbol{\theta}, \mathbf{s}) &= \int_0^Y \sum_{\ell=1}^m h(u - s_\ell | \boldsymbol{\theta}) du, \\ &= \sum_{\ell=1}^m H(Y - s_\ell | \boldsymbol{\theta}), \end{aligned} \tag{2}$$

where

$$H(Y - s_\ell | \boldsymbol{\theta}) = \int_0^{Y-s_\ell} h(u - s_\ell | \boldsymbol{\theta}) du.$$

From equation (2), the cumulative distribution function (c.d.f.) and probability density function (p.d.f.) at  $t^*$  are

$$\begin{aligned} F(Y | \boldsymbol{\theta}, \mathbf{s}) &= 1 - \exp\{-\Lambda(Y | \boldsymbol{\theta}, \mathbf{s})\}, \\ f(Y | \boldsymbol{\theta}, \mathbf{s}) &= \lambda(Y | \boldsymbol{\theta}, \mathbf{s}) \exp\{-\Lambda(Y | \boldsymbol{\theta}, \mathbf{s})\}. \end{aligned} \tag{3}$$

### 2.3 Specifying the Single-Administration Hazard Function


The single-administration hazard  $h(u | \boldsymbol{\theta})$  can be quite general, provided that it reasonably reflects the risk of toxicity for the agent under study and is sufficiently tractable to facilitate the necessary computations. In general, because certain toxicities may be hard to identify before the trial, one can include in the definition of toxicity any adverse event sufficiently severe that it precludes further administration of the agent. We assume that the hazard of toxicity from a single administration has a finite duration and vanishes to zero within  $\theta_3$  days. In the KGF trial, based on the physicians' experience, the hazard vanishes after  $\theta_3 = 18$  days.

Because we assume that  $h(\cdot)$  has finite duration, we cannot model  $h(\cdot)$  as the hazard of a typical parametric lifetime distribution, such as the gamma or Weibull, unless  $h(\cdot)$  is truncated appropriately. As a simple, practical alternative, we assume that  $h$  increases linearly to a maximum and decreases linearly thereafter. Specifically, we define

$$h(u | \boldsymbol{\theta}) = \begin{cases} \frac{u}{\theta_1} \theta_2, & 0 \leq u \leq \theta_1, \\ \theta_2 \frac{\theta_3 - u}{\theta_3 - \theta_1}, & \theta_1 < u \leq \theta_3, \\ 0, & u > \theta_3 \text{ or } u < 0. \end{cases} \quad (4)$$

Thus,  $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$ , with  $\theta_1$  the time at which  $h(u | \boldsymbol{\theta})$  reaches its maximum,  $\theta_2$ , and  $\theta_3$  the time when  $h(\cdot)$  vanishes to zero. Figure 1 illustrates this function. Initially, we assumed that  $h(\cdot)$  had only the two parameters  $\theta_1$  and  $\theta_2$ , with  $\theta_3$  fixed and assumed known. However, we found that fixing  $\theta_3$  severely hindered the method's ability to locate the optimal schedule when the actual duration of  $h(\cdot)$  was much longer than the assumed value of  $\theta_3$ .

Other forms for  $h(u | \boldsymbol{\theta})$  are possible, depending on the particular application. For example, the Weibull hazard  $h_1(u | \boldsymbol{\theta}) = \theta_1 \theta_2 (\theta_1 u)^{\theta_2 - 1}$  allows the risk of toxicity to continue indefinitely, with the shape parameter  $\theta_2$  determining



**Figure 1.** Parametric hazard function for a single administration of an agent.

whether the risk increases, decreases, or remains constant over time. One also could vary the hazard of toxicity for each administration. For example, in our application, one may believe that each administration has a hazard that is proportional to its proximity to BMT. Because each patient has their BMT at time  $s_3$  and the last possible administration would occur 77 days later, an extended version of (1) is  $\lambda(Y | \boldsymbol{\theta}, \mathbf{s}) = \sum_{\ell} h_{\ell}(Y - s_{\ell} | \boldsymbol{\theta})$ , where  $h_{\ell}(Y - s_{\ell} | \boldsymbol{\theta}) = \alpha(s_{\ell}) h(Y - s_{\ell} | \boldsymbol{\theta})$ , and

$$\alpha(s_{\ell}) = 2 - \left| \frac{s_{\ell} - s_3}{77 - s_3} \right|$$

is an inflation factor that is largest at  $s_3$  and decreases linearly as the time of administration moves away from  $s_3$ .

### 2.4 Likelihood and Posterior

The most recent data at study time  $t^*$  collected on patient  $i$ , for  $i = 1, 2, \dots, n^*$ , are  $\mathcal{D}_i = (\mathbf{s}_i, Y_i, \delta_i)$ . The optimal treatment sequence assigned to patient  $n^* + 1$  who enters the trial at  $t^*$  is based on the posterior of  $\boldsymbol{\theta}$  given the data available at  $t^*$ , which we denote by  $\mathcal{D}^* = (t^*, \mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_{n^*})$ . The likelihood at  $t^*$  is

$$\mathcal{L}(\mathcal{D}^* | \boldsymbol{\theta}) = \prod_{i=1}^{n^*} \{f(Y_i | \boldsymbol{\theta}, \mathbf{s}_i)\}^{\delta_i} \{1 - F(Y_i | \boldsymbol{\theta}, \mathbf{s}_i)\}^{1 - \delta_i}, \quad (5)$$

and, denoting the prior by  $p(\boldsymbol{\theta})$ , the posterior of  $\boldsymbol{\theta}$  is

$$g(\boldsymbol{\theta} | \mathcal{D}^*) = \frac{\mathcal{L}(\mathcal{D}^* | \boldsymbol{\theta}) p(\boldsymbol{\theta})}{\int \mathcal{L}(\mathcal{D}^* | \boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}}$$

Because the above integral cannot be obtained analytically under our assumed model, we compute posterior quantities via Markov chain Monte Carlo (MCMC) methods (Robert and Casella, 1999). Our computer code (in the C programming language) for the MCMC computations is available by request.

## 3. Establishing Prior Distribution for $\boldsymbol{\theta}$

### 3.1 Eliciting Prior Parameters from Investigators

The prior  $p(\boldsymbol{\theta})$  must be sufficiently uninformative so that it is dominated by the data in order to allow the schedule-finding algorithm to provide a safe and reliable design. We specify the prior distribution of  $\theta_1$  conditional on  $\theta_3$  because  $\theta_1 \leq \theta_3$ . Because no such restriction is necessary for  $\theta_2$ , we assume a priori that  $\theta_2$  is independent of both  $\theta_1$  and  $\theta_3$ . Thus,  $p(\boldsymbol{\theta}) = p_3(\theta_3) p_1(\theta_1 | \theta_3) p_2(\theta_2)$ .

To develop  $p(\boldsymbol{\theta})$ , we first ask the investigator to specify a range  $[T_{\ell}, T_u]$  of plausible values for the duration of  $h(\cdot)$ , as well as  $\mu_{\theta_3}$ , the anticipated average duration. We assume that  $\theta_3$  follows the generalized beta distribution

$$p_3(u) = \frac{(u - T_{\ell})^{a_3 - 1} (T_u - u)^{b_3 - 1}}{(T_u - T_{\ell})^{a_3 + b_3 - 1} B(a_3, b_3)}, \quad T_{\ell} \leq u \leq T_u,$$

where  $B(a, b) = \int_0^1 x^{a-1} (1-x)^{b-1} dx$ ,  $a_3 = k_3(\mu_{\theta_3} - T_{\ell})$ , and  $b_3 = k_3(T_u - \mu_{\theta_3})$ . The tuning constant  $k_3$  scales the values of  $a_3$  and  $b_3$  and modulates the variability of  $p_3(u)$ , which will increase as  $k_3$  decreases. For example, in the KGF trial, if the investigators believed that the duration of the hazard could range from 4 to 100 days, but was 18 days on average,  $a_3 \propto 14$  and  $b_3 \propto 82$ . Through preliminary simulations,

an appropriate value for  $k_3$  may be determined so that  $p_3(u)$  is sufficiently sensitive to the data collected on the first few patients.

Next, we assume that  $\theta_1 | \theta_3$  follows the generalized beta distribution

$$p_1(u | \theta_3) = \frac{u^{a_1-1}(\theta_3 - u)^{b_1-1}}{\theta_3^{a_1+b_1-1}B(a_1, b_1)}, \quad 0 \leq u \leq \theta_3,$$

where the parameters  $a_1$  and  $b_1$  are determined similarly to  $a_3$  and  $b_3$  above. However, a different approach for determining  $a_1$  and  $b_1$  is appropriate when the investigator summarizes the range of plausible values as the interval  $\mu_{\theta_1} \pm d$ . Viewing this as a 95% credible interval and assuming approximate symmetry about  $E(\theta_1 | \theta_3) = \mu_{\theta_1}$ , we obtain the equations  $\text{var}(\theta_1 | \theta_3) = \theta_3^2 a_1 b_1 / \{(a_1 + b_1)^2 (a_1 + b_1 + 1)\} = d^2/4$  and  $\mu_{\theta_1} = \theta_3^3 a_1 / (a_1 + b_1)$ , which yield the prior parameters

$$a_1 = \frac{\mu_{\theta_1}}{\theta_3} \left[ \frac{4\mu_{\theta_1}(\theta_3 - \mu_{\theta_1})}{d^2} - 1 \right],$$

$$b_1 = \frac{\theta_3 - \mu_{\theta_1}}{\theta_3} \left[ \frac{4\mu_{\theta_1}(\theta_3 - \mu_{\theta_1})}{d^2} - 1 \right].$$

Conventional methods for developing priors that are alternatives to the two approaches given above are given by Robert (2001, Chapter 3).

To determine  $p_2(u)$ , the investigator must select the schedule,  $\mathbf{s}^*$ , from the set  $\{\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(k)}\}$  that (s)he believes a priori to be the MTS. Recall that each patient is followed for up to  $\tau$  days, and  $p_\tau$  is the targeted probability of toxicity by  $\tau$ . From the a priori hypothesis of the optimal schedule, we determine  $\mu_{\theta_2} = E(\theta_2)$  as the value for which  $\mathbf{s}^*$  satisfies  $F(\tau | \theta_2 = \mu_{\theta_2}, \theta_1, \theta_3 = \theta_3^*, \mathbf{s}^*) = p_\tau$ , where  $\theta_3^*$  is a fixed value of  $\theta_3$ . A fixed value of  $\theta_1$  is unnecessary because the triangle in Figure 1 has area  $\theta_2 \theta_3 / 2$ , making the cumulative hazard (and cumulative probability) of toxicity independent of  $\theta_1$ . For example, in the KGF trial, if the investigators select  $\mathbf{s}^* = \mathbf{s}^{(2)}$  as optimal with  $\theta_3^* = 18$ , then each administration has cumulative hazard  $9\mu_{\theta_2}$ . Because there are 12 administrations, the entire schedule has a cumulative hazard of  $108\mu_{\theta_2}$  and, as a result,  $p_\tau = 1 - \exp(-108\mu_{\theta_2})$ , yielding  $\mu_{\theta_2} = -\log(1 - p_\tau)/108$ .

Because  $\theta_2$  is the height of the single-administration hazard, it is qualitatively different from  $\theta_1$  and  $\theta_3$ , which are on the time domain. Therefore, we do not restrict the upper bound of  $p_2(u)$ , and assume that  $\theta_2$  has a Gamma prior with parameters  $a_2 = k_2$  and  $b_2 = k_2/\mu_{\theta_2}$ , where  $k_2$  is a tuning constant used to modulate  $\text{var}(\theta_2) = \mu_{\theta_2}^2/k_2$ .

### 3.2 Calibrating the Prior Distribution of $\theta$

The ability of the data to dominate the prior distribution of  $\theta$  is heavily influenced by the variances of  $p_3(\theta_3)$ ,  $p_1(\theta_1 | \theta_3)$ , and  $p_2(\theta_2)$ , as quantified by the tuning constants  $k_2$  and  $k_3$ , and  $d$ , the width of the credible interval for  $\theta_1$ . Thus, it is essential to carefully evaluate the design's sensitivity to these parameters. By simulating the toxicity times of a small number of patients, one can compare the prior means for  $\theta$  to their respective posterior values and evaluate the effects of a small amount of data on  $p(\theta)$ . For example, if  $p(\theta_3)$  reflects the belief that

toxicity is unlikely beyond 25 days after administration, by simulating a few patients to have toxicities that occur far beyond 25 days, one can determine whether the prior allows the posterior mean of  $\theta_3$  to shift beyond its prior mean and reflect the data appropriately. If not,  $k_2$ ,  $k_3$ , and  $d$  may be calibrated and the exercise repeated until the desired effect is achieved.

The prior variances cannot be made arbitrarily large, as is usually done with Bayesian analyses of large data sets. In any small-scale clinical trial using adaptive methods, very few data are available, especially early in the trial. If there is substantial prior probability mass over too broad a range, this often cannot be overcome by a small amount of data, depending on the particular model, data structure, and decision-making algorithm. In the present setting, unduly large prior variances would severely hinder the algorithm's ability to assign optimal schedules during the trial and select an optimal MTS at the end.

It also is important to examine the prior on  $F(\tau | \theta, \mathbf{s}^{(j)})$  for each  $\mathbf{s}^{(j)}$  to determine whether  $p(\theta)$  may be producing pathological behavior by causing too much of the probability mass of a given  $F(\tau | \theta, \mathbf{s}^{(j)})$  to be placed near 0 or 1. That is, it is not  $p(\theta)$  or  $p(\theta | \text{data})$  per se that matter, but rather the consequent distributions of the  $F(\tau | \theta, \mathbf{s}^{(j)})$ 's. For example, suppose that one assumes a priori that  $4 \leq \theta_3 \leq 100$  days, with a mean  $E(\theta_3) = 18$  days, while  $0 \leq \theta_1 \leq 4$ . If schedule 2 is optimal a priori for the targeted threshold of  $p_\tau = 0.20$  for  $\tau = 100$  days, this motivates the assumption that  $\theta_2$  has a Gamma distribution with mean 0.0021. In this example, we also define  $k_2 = k_3 = 1$  based upon a preliminary calibration, the methods of which were discussed above. Figure 2 displays the resulting priors of  $F(100 | \theta, \mathbf{s}^{(j)}, 100)$  for  $j = 1, \dots, 6$ . As desired,  $E\{F(100 | \theta, \mathbf{s}^{(j)}, 100)\}$  is closest to 0.20 for  $j = 2$ ; also, the dispersion of  $F(100 | \theta, \mathbf{s}^{(j)}, 100)$  increases with  $j$  due to the cumulative nature of the schedules.

## 4. Trial Conduct

A maximum of  $N$  patients are enrolled in the trial, with each patient assigned a treatment administration sequence upon arrival. The first patient is assigned the shortest sequence,  $\mathbf{s}^{(1)}$ . Each patient is followed for up to  $\tau$  days, with treatment terminated if toxicity is observed. Given a desired threshold  $p_\tau$  for  $F(\tau | \theta, \mathbf{s}^{(j)})$ , we will consider two alternative criteria for choosing each patient's sequence.

*Criterion 1.* At time  $t^*$ , for each  $j = 1, \dots, k$ , compute

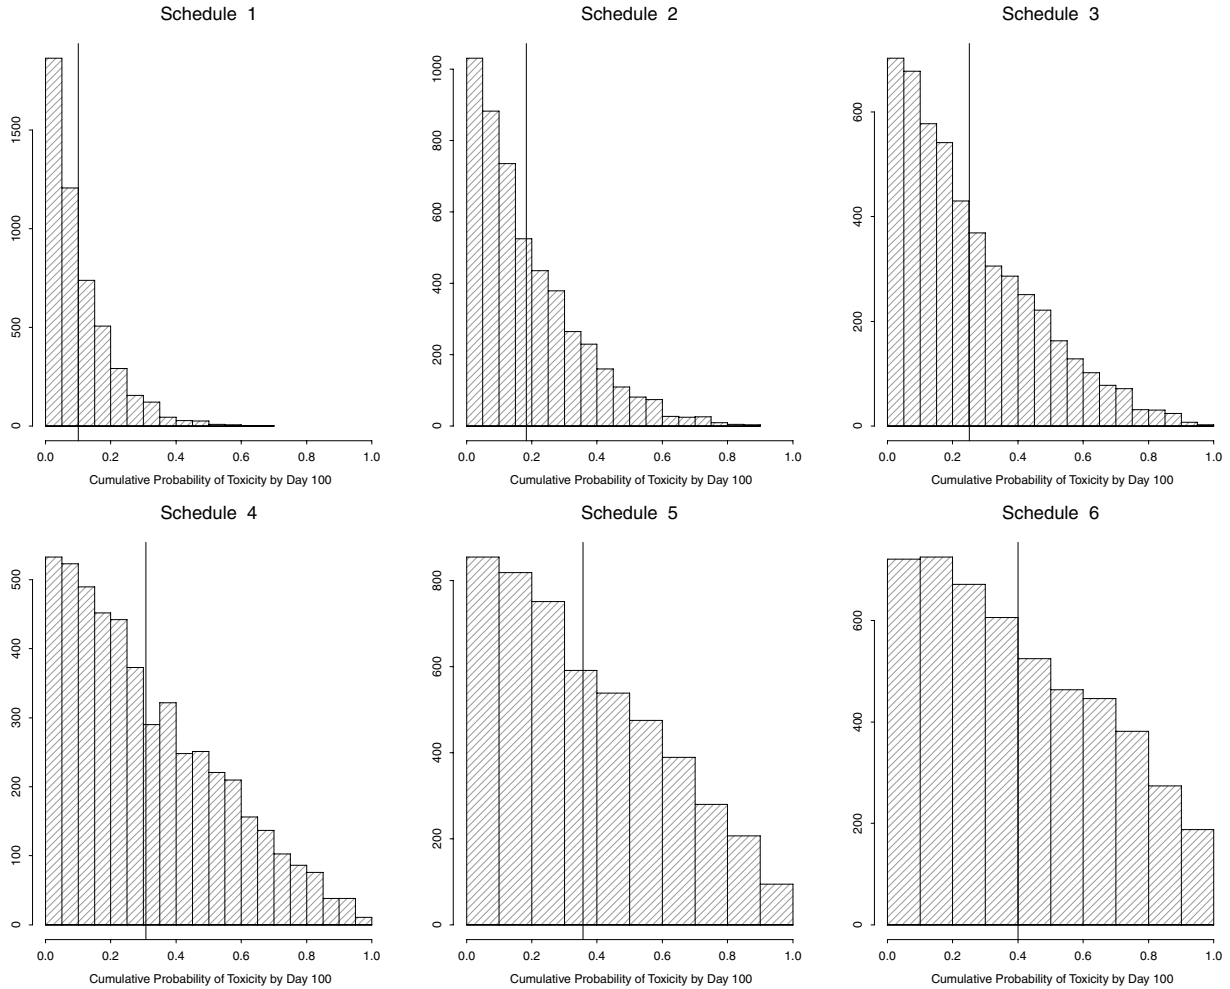
$$F_j^*(\tau) = E\{F(\tau | \theta, \mathbf{s}^{(j)}) | \mathcal{D}^*\}. \tag{6}$$

The best sequence is defined as that having  $F_j^*(\tau)$  closest to  $p_\tau$ , that is, minimizing  $|F_j^*(\tau) - p_\tau|$ . This criterion, as a function of treatment sequence, is analogous to the CRM criterion (O'Quigley, Pepe, and Fisher, 1990) based upon the posterior mean probability of the more usual binary toxicity, as a function of dose.

*Criterion 2.* At time  $t^*$ , for each  $j = 1, \dots, k$ , compute

$$\phi_j(\tau) = \Pr\{F(\tau | \theta, \mathbf{s}^{(j)}) > p_\tau | \mathcal{D}^*\}. \tag{7}$$

Because  $F(\tau | \theta, \mathbf{s}^{(j)})$  is monotone increasing in  $j$ , it follows that  $\phi_1(\tau) \leq \phi_2(\tau) \leq \dots \leq \phi_k(\tau)$ . Given a fixed upper limit,  $\bar{p}$ ,



**Figure 2.** Empirical prior distributions for cumulative probability of toxicity by day 100 for each schedule. Each histogram is based upon 5000 observations. The solid vertical line represents the mean of the distribution.

the best sequence is defined as the longest sequence for which  $\phi_j(\tau) < \bar{p}$ .

The second criterion is similar to the acceptability criteria used by Thall and Russell (1998) for dose finding based on both efficacy and toxicity, and it also is similar to the criterion for overdose control proposed by Babb, Rogatko, and Zacks (1998). Under either Criterion 1 or Criterion 2, the best sequence is assigned to patient  $n^* + 1$ . At the end of the trial, the MTS is defined as the best sequence based on the final data.

To protect patient safety, we impose the additional constraint that only incremental schedule escalation is permitted. Formally, if  $\mathbf{s}^{(j)}$  is determined to be the best schedule, the next patient is assigned  $\mathbf{s}^{(j)}$  only if each of the schedules  $\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(j-1)}$  previously has been assigned to at least  $M$  patients, for a predetermined  $M \geq 1$ . Otherwise, the next patient enrolled is assigned to the longest schedule that meets this criterion. In our application, we have set  $M = 1$ , although larger values of  $M$  are possible. However, an important practical consideration is that larger values of  $M$  will slow the speed of schedule escalation, so that the trial will take longer

to accrue patients on longer schedules. Schedule deescalation is permitted per the criterion without any constraint.

**5. Application**

5.1 *Simulation Study Design*

In the KGF trial, the investigators wished to study  $k = 6$  schedules corresponding to 2, 4, 6, 8, 10, or 12 weeks of therapy. Because aGVHD is defined to occur during the first 100 days after transplant, the maximum period to monitor toxicity was specified to be  $\tau = 100$  days. Per the adaptive design, a schedule is specified for each patient, and KGF is discontinued if the patient experiences toxicity before 100 days. The goal is to determine how long a BMT patient can receive KGF as prophylaxis for aGVHD while controlling the probability of toxicity within 100 days to be 20%.

The investigators believed a priori that the hazard of toxicity for a single administration vanishes after an average of 18 days, with a range of 4–100 days. They also believed that  $\theta_1$ , the time when the hazard of toxicity from one administration of KGF is largest, should be at most 4 days. Using the methods described in Section 3, we derived parameters  $a_3$  and

**Table 1**  
*Probability of toxicity within 100 days for each schedule under each of the simulation scenarios*

Scenario	Duration of hazard (days)	10000 $\theta_2$	Schedule (weeks)					
			1 (2)	2 (4)	3 (6)	4 (8)	5 (10)	6 (12)
1	18	4.13	0.20	0.36	0.49	0.59	0.67	0.74
	50	1.49	0.20	0.36	0.49	0.59	0.67	0.72
2	18	2.07	0.11	0.20	0.28	0.36	0.43	0.49
	50	0.74	0.11	0.20	0.28	0.36	0.42	0.47
3	18	1.38	0.07	0.14	0.20	0.26	0.31	0.36
	50	0.50	0.07	0.14	0.20	0.26	0.31	0.35
4	18	1.03	0.05	0.11	0.15	0.20	0.24	0.28
	50	0.37	0.05	0.11	0.15	0.20	0.24	0.27
5	18	0.83	0.04	0.09	0.13	0.16	0.20	0.24
	50	0.30	0.04	0.09	0.13	0.17	0.20	0.23
6	18	0.69	0.04	0.07	0.11	0.14	0.17	0.20
	50	0.26	0.04	0.08	0.11	0.15	0.18	0.20
7	18	1.60	0.08	0.16	0.23	0.29	0.35	0.40
8	n/a	n/a	0.07	0.14	0.20	0.26	0.31	0.36

$b_3$  for  $p(\theta_3)$  and  $a_1$  and  $b_1$  for  $p(\theta_1 | \theta_3)$ . Based upon the results of an earlier trial, the rate of toxicity on the lowest schedule of KGF was very low, and the investigators believed that even 12 weeks of KGF would not cause appreciably more toxicity, leading us to assume that  $p(\theta_2)$  was a Gamma distribution

with mean 0.0007. Through a detailed sensitivity analysis, we determined that  $k_3 = 0.1$  and  $k_2 = 0.2$  allowed the data to have adequate influence on the posterior of  $\theta$ .

After studying Criterion 2 with  $0.20 \leq \bar{p} \leq 0.80$ , we found that  $\bar{p} = 0.60$  worked best. Values of  $\bar{p} > 0.60$  tended to make

**Table 2**  
*Performance of the design with 30 patients, assuming the single-administration hazard duration is 18 days. Each entry is the schedule's selection percentage, with number of patients assigned to that schedule given in parentheses. Values within one schedule of the true MTS are given in boldface. The final column is the mean absolute deviation (in percentage points) across all 1000 simulations of the cumulative probability of toxicity of the selected MTS and the threshold  $p_\tau$ .*

Scenario		Schedule (no. of weeks)						$\Delta_{p_\tau}$
		1 (2)	2 (4)	3 (6)	4 (8)	5 (10)	6 (12)	
Criterion 1	1	<b>86.6</b> ( <b>20.9</b> )	<b>13.3</b> ( <b>6.4</b> )	0.1 (1.4)	0.0 (0.6)	0.0 (0.3)	0.0 (0.3)	2.0
	2	<b>27.8</b> ( <b>8.9</b> )	<b>53.5</b> ( <b>11.8</b> )	<b>16.4</b> ( <b>4.8</b> )	2.0 (1.9)	0.2 (1.5)	0.1 (1.1)	0.8
	3	7.9 (4.3)	<b>33.0</b> ( <b>7.0</b> )	<b>37.1</b> ( <b>9.4</b> )	<b>14.4</b> ( <b>3.8</b> )	4.9 (3.4)	2.7 (2.1)	1.3
	4	1.7 (2.7)	21.0 (4.7)	<b>23.2</b> ( <b>6.6</b> )	<b>32.6</b> ( <b>6.9</b> )	<b>10.9</b> ( <b>6.2</b> )	10.6 (3.0)	1.9
	5	0.8 (2.1)	10.8 (4.8)	23.7 (4.7)	<b>24.2</b> ( <b>5.9</b> )	<b>26.3</b> ( <b>9.1</b> )	<b>14.2</b> ( <b>3.4</b> )	3.5
	6	0.1 (1.7)	5.7 (3.7)	16.9 (4.5)	18.6 (3.8)	<b>20.1</b> ( <b>5.0</b> )	<b>38.6</b> ( <b>11.3</b> )	4.1
Criterion 2 ( $\bar{p} = 0.60$ )	1	<b>93.3</b> ( <b>21.9</b> )	<b>6.4</b> ( <b>4.6</b> )	0.3 (1.7)	0.0 (0.8)	0.0 (0.5)	0.0 (0.5)	1.1
	2	<b>40.1</b> ( <b>8.9</b> )	<b>41.9</b> ( <b>10.4</b> )	<b>14.0</b> ( <b>4.8</b> )	3.0 (2.3)	0.8 (2.2)	0.2 (1.4)	1.9
	3	14.2 (5.0)	<b>28.6</b> ( <b>6.5</b> )	<b>31.4</b> ( <b>7.5</b> )	<b>14.6</b> ( <b>5.6</b> )	5.6 (3.0)	5.6 (2.4)	1.2
	4	4.7 (2.8)	16.6 (4.7)	<b>20.1</b> ( <b>5.2</b> )	<b>24.8</b> ( <b>7.4</b> )	<b>20.2</b> ( <b>6.0</b> )	13.6 (3.9)	1.1
	5	1.9 (2.3)	9.8 (3.2)	14.0 (3.6)	<b>18.4</b> ( <b>5.0</b> )	<b>39.1</b> ( <b>10.8</b> )	<b>16.8</b> ( <b>5.1</b> )	2.6
	6	0.4 (1.6)	3.9 (2.6)	9.9 (3.4)	13.4 (3.3)	<b>19.0</b> ( <b>5.3</b> )	<b>53.4</b> ( <b>13.8</b> )	3.1

**Table 3**

Performance of the design with 30 patients, assuming the single administration hazard duration is 50 days. Each entry is the schedule's selection percentage, with number of patients assigned to that schedule given in parentheses. Values within one schedule of the true MTS are given in boldface. The final column is the mean absolute deviation (in percentage points) across all 1000 simulations of the cumulative probability of toxicity of the selected MTS and the threshold  $p_\tau$ .

		Schedule (no. of weeks)						
	Scenario	1 (2)	2 (4)	3 (6)	4 (8)	5 (10)	6 (12)	$\Delta_{p_\tau}$
Criterion 1	1	<b>83.3</b> <b>(20.4)</b>	<b>11.4</b> <b>(6.0)</b>	0.3 (1.8)	0.0 (0.8)	0.0 (0.5)	0.0 (0.5)	0.9
	2	<b>25.8</b> <b>(7.6)</b>	<b>50.9</b> <b>(10.9)</b>	<b>18.1</b> <b>(5.5)</b>	4.2 (2.5)	0.9 (2.2)	0.1 (1.3)	0.0
	3	6.3 (3.5)	<b>31.9</b> <b>(6.9)</b>	<b>34.0</b> <b>(8.1)</b>	<b>17.9</b> <b>(4.3)</b>	6.1 (4.7)	3.8 (2.5)	0.5
	4	1.7 (2.4)	13.4 (4.9)	<b>26.7</b> <b>(5.3)</b>	<b>30.4</b> <b>(6.5)</b>	<b>15.6</b> <b>(7.6)</b>	12.2 (3.3)	1.0
	5	0.9 (1.9)	7.9 (3.9)	20.1 (4.6)	<b>21.0</b> <b>(5.1)</b>	<b>32.2</b> <b>(11.0)</b>	<b>17.9</b> <b>(3.5)</b>	2.7
	6	0.3 (1.6)	4.9 (3.4)	14.6 (4.4)	15.6 (3.7)	<b>19.9</b> <b>(4.4)</b>	<b>44.7</b> <b>(12.5)</b>	3.6
Criterion 2 ( $\bar{p} = 0.60$ )	1	<b>96.4</b> <b>(21.7)</b>	<b>3.4</b> <b>(4.1)</b>	0.2 (1.9)	0.0 (1.0)	0.0 (0.7)	0.0 (0.6)	0.6
	2	<b>39.2</b> <b>(8.2)</b>	<b>44.0</b> <b>(10.1)</b>	<b>11.8</b> <b>(4.5)</b>	3.1 (3.0)	1.4 (2.6)	0.5 (1.6)	1.7
	3	13.0 (4.6)	<b>29.1</b> <b>(5.9)</b>	<b>30.7</b> <b>(7.7)</b>	<b>13.5</b> <b>(5.3)</b>	6.7 (3.9)	7.0 (2.6)	0.8
	4	5.5 (3.1)	18.7 (4.5)	<b>21.6</b> <b>(5.8)</b>	<b>23.2</b> <b>(8.3)</b>	<b>18.8</b> <b>(4.8)</b>	12.2 (3.5)	1.7
	5	2.6 (2.1)	6.4 (3.1)	15.7 (3.8)	<b>22.6</b> <b>(5.8)</b>	<b>37.9</b> <b>(11.0)</b>	<b>14.8</b> <b>(4.2)</b>	2.6
	6	0.9 (1.8)	5.2 (2.6)	11.8 (3.1)	12.8 (3.3)	<b>15.1</b> <b>(4.4)</b>	<b>54.2</b> <b>(14.7)</b>	3.2

the algorithm unacceptably conservative, frequently selecting an MTS at schedules lower than the true MTS, while  $\bar{p} < 0.60$  caused the algorithm to select overly toxic schedules too often. We studied the design with a maximum sample size of 30 patients, which is feasible but sufficient to determine the MTS with reasonable accuracy. Patient interarrival times were assumed to be uniformly distributed from 12 to 16 days. In each simulation, posterior values were based upon 2000 MCMC samples from the posterior of  $\theta$  preceded by a burn-in of 2000 samples.

We first examined the design's performance using both Criterion 1 and Criterion 2 (with  $\bar{p} = 0.60$ ) under each of six scenarios, with schedule  $s^{(j)}$  optimal under the  $j$ th scenario. The times to toxicity were simulated assuming that  $\theta_1$  occurs 2 days after administration and that  $\theta_3$  is either 18 or 50 days. The results are summarized in Tables 2 and 3. We also examined the design's performance, using Criterion 1, under two additional scenarios. In scenario 7, the duration of the hazard was 18 days, but the actual MTS was located between schedules 2 and 3. In scenario 8, schedule 3 was the MTS, but the form of  $F(\cdot)$  deviated from that used in our model, with toxicity occurring uniformly over the interval  $[10 + 14(j - 1), 10 + 14j]$  under  $s^{(j)}$ . This leads to relatively late-onset toxicities, and is very different from our assumed additive triangular hazard model. Results for scenarios 7 and 8 are given in Table 4. In the first seven scenarios, the value of  $\theta_2$  is varied to reflect which schedule is optimal; those

**Table 4**

Performance of the design with 30 patients under two additional settings. Each entry is the schedule's selection percentage, with number of patients assigned to that schedule given in parentheses. In scenario 7, the actual MTS is in-between schedules 2 and 3. In scenario 8, schedule 3 is the MTS, but times to toxicity deviate from the assumed triangular hazard.

		Schedule (no. of weeks)					
	Scenario	1 (2)	2 (4)	3 (6)	4 (8)	5 (10)	6 (12)
7		11.2 (5.2)	48.0 (10.7)	28.3 (6.5)	8.1 (3.2)	3.0 (1.7)	1.4 (2.8)
		2.0 (2.4)	24.3 (5.7)	37.1 (6.1)	20.5 (4.5)	9.1 (3.1)	7.0 (8.2)

values are shown in Table 1, which also contains the actual day 100 probabilities of toxicity for each schedule under all eight scenarios.

5.2 Simulation Results

Tables 2 and 3 display how frequently each schedule was selected as the MTS, and the mean number of patients assigned to each schedule. In interpreting the selection percentages given in these tables, it is important to bear in mind that

the difference in toxicity probabilities between neighboring schedules generally decreases across the six scenarios, making location of the MTS most difficult in scenario 6. To quantify this, the final columns in Tables 2 and 3 display the average absolute deviation between the toxicity probability of the selected MTS and the desired threshold  $p_r = 0.20$ . Larger values of this statistic indicate greater difficulty in locating the MTS.

When  $\theta_3 = 18$  days, Table 2 shows that the algorithm using Criterion 1 selected the MTS within one schedule of the optimum MTS 99.9%, 97.7%, 84.5%, 66.7%, 64.7%, and 58.7% of the time for scenarios 1–6, respectively. The algorithm tends to err conservatively, in that it is more likely to select schedules shorter than the MTS over schedules longer than the MTS. In addition, in all six scenarios more patients are assigned to the MTS than to any of the other schedules. Similar conclusions are reached when using Criterion 2 with  $\bar{p} = 0.60$ , where the algorithm selected the MTS within one schedule of the optimum MTS in 99.7%, 96.0%, 74.6%, 65.1%, 74.3%, and 72.4% of the time for scenarios 1–6.

Criterion 2 was more likely than Criterion 1 to correctly identify the MTS when the true MTS was one of the shortest or longest schedules, and Criterion 1 outperformed Criterion 2 when the true MTS was in the middle of all the schedules. Note that  $\theta_1$  has no impact upon each course's estimated cumulative probability of toxicity. Instead,  $\theta_1$  influences the time at which each patient experiences toxicity and the process of escalation and deescalation during the study. This was supported by additional simulations (not shown) in which we found the prior distribution for  $\theta_1$  had negligible impact on the results shown in Table 2.

The prior belief that the longest schedule,  $\mathbf{s}^{(6)}$ , was optimal was reflected in the prior mean of  $\theta_2$ . Given that the algorithm was conservative despite the fact that the longest schedule was the a priori optimum MTS, the first five scenarios supply evidence that the method reliably shifts the posterior distribution of  $\theta_2$  away from a misspecified prior mean, even with the limited sample size of 30 patients.

Table 3 contains results analogous to those in Table 2 when the true duration of the toxicity hazard for a single administration is extended to  $\theta_3 = 50$  days. Because  $p(\theta_3)$  has a mean of 18 days, here toxicities can occur much later than the prior distribution anticipates. If  $p(\theta_3)$  is not sufficiently flexible, dose escalation will move too rapidly, essentially because the algorithm places insufficient weight on any "late-onset" toxicities when computing the posterior of  $\boldsymbol{\theta}$ . As a result, one might expect the algorithm to be liberal, i.e., likely to identify the MTS at a schedule longer than the optimum MTS, as well as assigning too many patients to overly toxic schedules. However, Table 3 shows that the algorithm is relatively insensitive to the fixed value of  $\theta_3$  in all six scenarios. We also found that, with either Criterion 1 or Criterion 2, the algorithm is slightly more likely to identify the MTS at higher schedules when the true  $\theta_3$  is larger than expected (simulations not shown). This finding is most apparent for scenarios 5 and 6, where the algorithm correctly identifies the MTS more often when  $\theta_3 = 50$  than when  $\theta_3 = 18$ .

The first two lines of Table 4 summarize the algorithm's performance when the true MTS lies between schedules 2 and 3. The algorithm performs well in this setting, with ei-

ther schedule 2 or schedule 3 selected 76.3% of the time. The last two lines of Table 4 summarize the algorithm's performance under scenario 8, where schedule 3 is the MTS but the actual times to toxicity have a uniformly late onset and do not follow the assumed additive triangular hazard model. The algorithm's performance is similar to that under scenario 3, for which the true toxicity distribution matches that of the model. In scenario 8, more subjects are assigned to the longest schedule than they are in scenario 3. This slight skewness toward longer schedules is likely from influence of the prior and indicates that late-onset toxicities affect schedule assignment during the study without substantively changing the final decision of the algorithm.

## 6. Concluding Remarks

While we have focused on a fixed number of nested schedules, the threshold probability  $p_r$  could be attained by schedules that are not included in  $\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(6)}$ . Regardless, investigators often select schedules that are nested in this way and easy to apply in clinical practice. For example, it is unlikely that the schedule (1, 3, 4, 5, 10, 12) would be preferred over (1, 2, 3, 8, 9, 10), because the second schedule is simpler and easier to implement. One may envision a statistical approach that optimizes one of our criteria over all possible schedules. While this may be interesting from a statistical viewpoint, it may be difficult to implement. In clinical application,  $\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(k)}$  are likely to be determined by practical considerations, rather than from a desire to obtain "best" statistical properties.

An alternate design for this study was recently published by Braun, Levine, and Ferrara (2003), who treated each schedule as a "dose" and determined the MTS using the TiTE CRM with study-specific modifications. However, by considering each schedule to be a dose, patients who received an incomplete schedule were only evaluated up to the point of their last fully completed schedule. Furthermore, the doses overlapped, leading to some ambiguity as to which dose contributed to a late-onset toxicity. Our method avoids both of these limitations by accounting for each administration separately in the model. These limitations also underscore the fact that the TiTE CRM is well suited to identify a single optimal dose, but it is not intended to identify an optimal sequence of doses. Choosing a schedule rather than a dose may be regarded as a paradigm shift, and our model and method are motivated to more closely reflect administration of the agent, whose cumulative effects could lead to toxicity, possibly of late onset.

Our method may be extended to accommodate a wider range of applications. One important generalization is to model the single administration hazard so that the dose administered on each day can vary within each patient, and between patients. One also may allow the hazard to vary with patient characteristics, so that the schedules may be chosen more specifically for each patient. Our model parameterization relies on a strong homogeneity assumption about the hazard of toxicity over the sequence of treatment times. For example, under our model, one could observe nothing other than patients who receive the shortest sequence,  $\mathbf{s}^{(1)}$ , learn about  $\boldsymbol{\theta}$ , and make predictions about the risk of toxicity for a patient who is scheduled to receive any  $\mathbf{s}^{(j)}$ ,  $j \geq 1$ . However, this is a feature of any parametric regression model that



is assumed to hold over a given domain of covariates, which in this case are the  $s^{(j)}$ 's. If this assumption is considered invalid in a particular setting, then a more general model accounting for inhomogeneity across schedules would be needed. Currently, we are investigating these generalizations.

## ACKNOWLEDGEMENTS

Tom Braun and Zheng Yuan's research was partially supported by the University of Michigan's NIH Cancer Center support grant 5-P30-CA-46592. Peter Thall's research was partially supported by NCI grant RO1-CA-83932.

## REFERENCES

- Babb, J., Rogatko, A., and Zacks, S. (1998). Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine* **17**, 1103–1120.
- Braun, T. M., Levine, J. E., and Ferrara, J. L. M. (2003). Determining a maximum tolerated cumulative dose: Dose reassignment within the TITE-CRM. *Controlled Clinical Trials* **24**, 669–681.
- Cheung, Y. and Chappell, R. (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56**, 1177–1182.
- Farrell, C., Bready, J., Rex, K., et al. (1998). Keratinocyte growth factor protects mice from chemotherapy and radiation-induced gastrointestinal injury and mortality. *Cancer Research* **58**, 933–939.
- Goodman, S., Zahurak, M., and Piantadosi, S. (1995). Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine* **14**, 1149–1161.
- Krijanovski, O. I., Hill, G. R., Cooke, K. R., Teshima, T., Crawford, J. M., Brinson, Y. S., and Ferrara, J. L. M. (1999). Keratinocyte growth factor separates graft-versus-leukemia effects from graft-versus-host disease. *Blood* **94**, 825–831.
- Meropol, N., Somer, R., Gutheil, J., Pelley, R. J., Modiano, M. R., Rowinsky, E. K., Rothenberg, M. L., Redding, S. W., Serdar, C. M., Yao, B., Heard, R., and Rosen, L. S. (2003). Randomized phase I trial of recombinant human keratinocyte growth factor plus chemotherapy: Potential role as mucosal protectant. *Journal of Clinical Oncology* **21**, 1452–1458.
- 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.
- Panoskaltis-Mortari, A., Lacey, D., Valleria, D., and Blazar, B. (1998). Keratinocyte growth factor administered before conditioning ameliorates graft-versus-host disease after allogeneic bone marrow transplantation in mice. *Blood* **92**, 3960–3967.
- Robert, C. P. (2001). *The Bayesian Choice*, 2nd edition. New York: Springer.
- Robert, C. P. and Casella, G. (1999). *Monte Carlo Statistical Methods*. New York: Springer.
- Storer, B. E. (1989). Design and analysis of phase I clinical trials. *Biometrics* **45**, 925–937.
- 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.

Received April 2004. Revised August 2004.

Accepted September 2004.