Biometrics 59, 487–496
September 2003

Dose-Finding with Two Agents in Phase I Oncology Trials

Peter F. Thall,[1,*] Randall E. Millikan,[2] Peter Mueller,[1] and Sang-Joon Lee[1]

[1]Department of Biostatistics, Box 447, University of Texas, M. D. Anderson Cancer Center, 1515 Holcombe Blvd., Houston, Texas 77030, U.S.A.

[2]Department of Genitourinary Medical Oncology, Box 422, University of Texas, M. D. Anderson Cancer Center, 1515 Holcombe Blvd., Houston, Texas 77030, U.S.A.

[*]*email:* rex@mdanderson.org

Summary. We propose an adaptive two-stage Bayesian design for finding one or more acceptable dose combinations of two cytotoxic agents used together in a Phase I clinical trial. The method requires that each of the two agents has been studied previously as a single agent, which is almost invariably the case in practice. A parametric model is assumed for the probability of toxicity as a function of the two doses. Informative priors for parameters characterizing the single-agent toxicity probability curves are either elicited from the physician(s) planning the trial or obtained from historical data, and vague priors are assumed for parameters characterizing two-agent interactions. A method for eliciting the single-agent parameter priors is described. The design is applied to a trial of gemcitabine and cyclophosphamide, and a simulation study is presented.

Key words: Adaptive design; Bayesian design; Dose-finding; Phase I clinical trial.

1. Introduction

The goal of a Phase I clinical trial of a new agent is to find a dose having acceptable toxicity, where dose-limiting toxicity is typically defined as side effects sufficiently morbid that they constitute a practical limitation to the delivery of treatment. The particular effects, and the severity at which the limitation to treatment is imposed, are specific to each clinical setting. For ethical reasons, most Phase I trials are conducted adaptively (Storer, 1989), with the dose for each successive patient cohort chosen using the dose-toxicity data from patients treated previously in the trial. The challenges in Phase I trial design and conduct are that little is known *a priori* about the dose-toxicity probability curve, and decisions must be based on very small sample sizes.

In a Phase I trial of two agents used in combination for the first time in a particular patient disease group, the goal is to determine an acceptable dose pair, $\mathbf{x} = (x_1, x_2)$. The problem of constructing an algorithm for choosing doses sequentially in this case is more difficult than the analogous single-agent problem. Although prior knowledge of each agent's individual dose-toxicity curve typically is available from previous Phase I trials, such information is of limited use for predicting the probability of toxicity as a function of $\mathbf{x}$ when the two agents are used together. This is because the biochemical and biological effects of the combination may be quite complex, and the dose-toxicity probability surface may depend largely on unknown interactions between the two agents. Simon and Korn (1990) give a detailed discussion of this problem. In oncology, this phenomenon is commonly seen, for example, with combination chemotherapies, selective immunotoxins combined with T-cell infusion, and graft-versus-host disease prophylaxis following ablative chemotherapy in allogeneic bone marrow transplantation. An additional difficulty is that even a carefully defined bounded planar region of potential $\mathbf{x}$ values is a much larger set than the line segment within which dose-finding is usually done in a single-agent Phase I trial. An exhaustive search on even a coarse grid of (x_1, x_2) pairs is simply not feasible in Phase I, due to both the limited sample size and ethical constraints.

Despite these difficulties, practical methods may be constructed for dose-finding with two agents in Phase I. Simon and Korn (1990, 1991) propose a method based on the idea of the "total equivalent dose" of the combination. Another approach is to limit the search to a sequence of combinations such that each level is obtained from the previous one by increasing one or both of the single-agent doses. Since this ensures that the probability of toxicity is monotone in the levels, a single-agent dose-finding algorithm requiring only monotonicity, such as the continual reassessment method (CRM) (O'Quigley, Pepe and Fisher, 1990) may be applied. This type of approach severely restricts the domain of possible dose pairs that are considered, however.

In this article, we propose a two-stage Bayesian design for finding one or more acceptable dose pairs in the two-agent setting. We assume a parametric model for the toxicity probability, $\pi(\mathbf{x}, \boldsymbol{\theta})$, as a function of $\mathbf{x}$ and a parameter vector, $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3)$. The subvectors $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ parameterize the two single-agent toxicity probabilities, $\pi_1(x_1, \boldsymbol{\theta}_1)$ and $\pi_2(x_2, \boldsymbol{\theta}_2)$, while $\boldsymbol{\theta}_3$ accounts for interaction between the two agents. Mantel (1974) defined therapeutic synergism as a clinical

effect of two agents used together that is not achievable with either agent used alone, regardless of dose. Our approach exploits the facts that $\pi(\mathbf{x}, \boldsymbol{\theta})$ is increasing in each of x_1 and x_2, and that the curves $\{\pi(\mathbf{x}, \boldsymbol{\theta}) : x_2 = 0, x_1 \geq 0\}$ and $\{\pi(\mathbf{x}, \boldsymbol{\theta}) : x_1 = 0, x_2 \geq 0\}$ on the edges of the toxicity surface $\{\pi(\mathbf{x}, \boldsymbol{\theta}) : x_1 \geq 0, x_2 \geq 0\}$ must coincide with the respective single-agent dose-toxicity curves, $\{\pi_1(x_1, \boldsymbol{\theta}_1) : x_1 \geq 0\}$ and $\{\pi_2(x_2, \boldsymbol{\theta}_2) : x_2 \geq 0\}$. Thus, while prior knowledge about the two single-agent curves cannot predict interactions, it provides useful information about the edges of the toxicity surface.

Implementation of our proposed method requires close collaboration with the physician(s) planning the trial, who must provide a definition of dose-limiting toxicity, hereafter "toxicity," a fixed target toxicity probability, π^*, and information characterizing the dose-toxicity curve of each component as a single agent. This information is used to construct informative priors on $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$, and a vague prior is assumed for $\boldsymbol{\theta}_3$. We define an *acceptable dose* (AD) to be any combination $\mathbf{x}$ having posterior mean toxicity probability

$$\mathrm{E}\{\pi(\mathbf{x}, \boldsymbol{\theta}) \mid \mathrm{data}\} = \pi^*. \tag{1}$$

In contrast with the single-agent setting, when $\mathbf{x}$ is two-dimensional, equation (1) does not have a unique solution. Rather, the set of dose pairs satisfying (1) is the random contour

$$L_2(\pi^*, \mathrm{data}) = \{\mathbf{x} : E\{\pi(\mathbf{x}, \boldsymbol{\theta}) \mid \mathrm{data}\} = \pi^*\} \tag{2}$$

in the two-dimensional domain of $\mathbf{x}$ (Figure 1(a)). We propose a two-stage, outcome-adaptive Bayesian design, with doses for successive patient cohorts chosen from the plane of two-dose combinations based on the most recently updated posteriors. In stage 1, dose-finding is restricted to a fixed line segment, L_1. The search in stage 2 chooses doses from L_2 (π^*, *data*) subject to additional optimality criteria. At the end of the trial, several ADs may be selected for study in a subsequent trial.

Section 2 describes the motivating trial that we use to illustrate the method. Probability models and a method for eliciting priors on the single-agent dose-toxicity parameters are presented in Section 3. The dose-finding algorithm is presented in Section 4. Computational methods are described in Section 5. Section 6 describes an application of the method, including a simulation study to evaluate the design's operating characteristics. Robustness is discussed in Section 7, and we close with a discussion in Section 8.

2. A Trial of Gemcitabine + Cyclophosphamide

Our illustrative application, which originally motivated this research, is a trial of gemcitabine (Gem) and cyclophosphamide (CTX), which may have several clinical applications. Gem is a chemotherapeutic agent with broad spectrum anticancer activity. This drug has a complex pharmacology and its biologic effects show remarkable schedule dependence. For example, when given as a single dose administered over 30 minutes, repeated every three weeks, doses in excess of 4000 mg/m^2 are tolerated. When given over 96 hours, however, only about 90 mg/m^2 can be tolerated. It is well established that Gem shows substantial synergy in combination

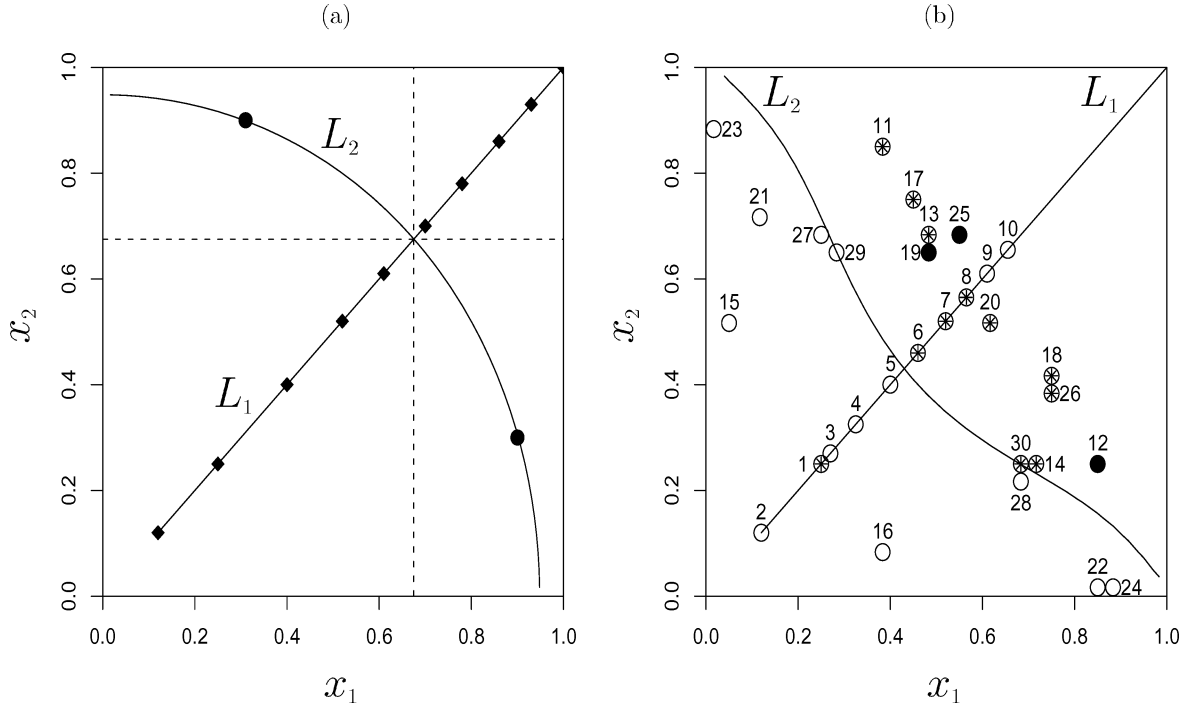

Figure 1. (**a**) The fixed line L_1 and random contour L_2 of the dose-finding algorithm. (**b**) Illustration of the successive dose pairs chosen in a typical case. Each cohort of two patients is represented by a circle numbered by the cohort's order in the trial, with 0, 1, or 2 patients experiencing toxicity in the two-patient cohort represented, respectively, by an empty circle, a circle enclosing a star, and a shaded circle.

with other agents, and this is especially true of combinations with DNA-damaging agents, such as a classical group of drugs known as alkylators. Because the archetype of this class is CTX, there is considerable interest in exploring the Gem/CTX combination. The goal of this trial is to determine three acceptable dose combinations for study in a subsequent Phase II trial: one that is mostly Gem, one that is mostly CTX, and one that is a relatively even mix of the two agents. The clinical protocol for the Gem/CTX trial, using the methodology described here, is under review at the University of Texas M. D. Anderson Cancer Center, and we anticipate that the trial soon will be activated.

3. Dose-Toxicity Model

3.1 *Toxicity Probabilities*

Given a dose d_i and single-agent AD d_i^* for agent i, we formulate the dose-toxicity probability model in terms of the standardized doses $x_i = d_i/d_i^*$, for $i = 1, 2$. The standardized doses of the two agents will have similar domains, with most values between 0 and 1, so that model parameters pertaining to the two doses will have similar numerical scales which, in turn, stabilizes numerical computations. The ADs d_1^* and d_2^* may be elicited from the physician(s) along with the priors on $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ while planning the trial. We provide a method for eliciting the single-agent priors and ADs in Section 3.2, below.

For any $\boldsymbol{\theta}$, the toxicity probabilities π, π_1, and π_2 must satisfy the following admissibility conditions:

(a) $\pi(x_1, 0, \boldsymbol{\theta}) = \pi_1(x_1, \boldsymbol{\theta}_1)$ for all x_1, and $\pi(0, x_2, \boldsymbol{\theta}) = \pi_2(x_2, \boldsymbol{\theta}_2)$ for all x_2;
(b) $\pi(x_1, x_2, \boldsymbol{\theta})$ is increasing in both x_1 and x_2;
(c) $\pi_i(0, \boldsymbol{\theta}_i) = 0$ for $i = 1, 2$.

Properties (a) and (b) together imply that each $\pi_i(x_i, \boldsymbol{\theta}_i)$ is increasing in x_i, and consequently that $\pi(x_1, x_2, \boldsymbol{\theta}) > \max\{\pi_1(x_1, \boldsymbol{\theta}_1), \pi_2(x_2, \boldsymbol{\theta}_2)\}$ for all $x_1 > 0$ and $x_2 > 0$. This says that adding any amount of agent 2 to a given amount of agent 1 must increase the probability of toxicity. Properties (a) and (c) together imply that $\pi(0, 0, \boldsymbol{\theta}) = 0$.

There are many models that satisfy these conditions. Because the sample size typically is very small in Phase I, especially early in the trial, the model must be parsimonious. At the same time, the model must be sufficiently flexible to allow a wide variety of possible shapes that the dose-toxicity probability surface may assume. We will use the following model, which provides a reasonable balance between these conflicting practical requirements. Define

$$\pi(\mathbf{x}, \boldsymbol{\theta}) = \frac{\alpha_1 x_1^{\beta_1} + \alpha_2 x_2^{\beta_2} + \alpha_3 \left(x_1^{\beta_1} x_2^{\beta_2}\right)^{\beta_3}}{1 + \alpha_1 x_1^{\beta_1} + \alpha_2 x_2^{\beta_2} + \alpha_3 \left(x_1^{\beta_1} x_2^{\beta_2}\right)^{\beta_3}}, \tag{3}$$

with $\boldsymbol{\theta}_i = (\alpha_i, \beta_i)$ for $i = 1, 2, 3$ and each entry of the six-dimensional parameter vector $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3) = (\alpha_1, \beta_1, \alpha_2, \beta_2, \alpha_3, \beta_3)$ required to be positive real-valued. Denoting $\eta_i(x_i, \boldsymbol{\theta}_i) = \log(\alpha_i) + \beta_i \log(x_i)$, for $i = 1, 2$, it follows that each $\pi_i(x_i, \boldsymbol{\theta}_i) = \text{logit}^{-1}\{\eta_i(x_i, \boldsymbol{\theta}_i)\}$, and it is easy to verify that this model satisfies the admissibility conditions (a)–(c).

Denoting the dose combinations and toxicity indicators of the first n patients in the trial by $\mathbf{Z}_n = \{(\mathbf{x}_k, Y_k), k = 1, \ldots, n\}$, the likelihood is

$$f(\mathbf{Z}_n \mid \boldsymbol{\theta}) = \prod_{k=1}^{n} \pi(\mathbf{x}_k, \boldsymbol{\theta})^{Y_k} \{1 - \pi(\mathbf{x}_k, \boldsymbol{\theta})\}^{1-Y_k}. \tag{4}$$

Denoting the prior on $\boldsymbol{\theta}$ at the start of the trial by $f(\boldsymbol{\theta})$, by Bayes' theorem the posterior given $\mathbf{Z}_n$ is $f(\boldsymbol{\theta} \mid \mathbf{Z}_n) \propto f(\mathbf{Z}_n \mid \boldsymbol{\theta}) \times f(\boldsymbol{\theta})$. In the present setting, posterior integrals are analytically intractable and numerical integration is required. This is described in Section 5, below.

Our dose-finding algorithm begins with independent informative priors on $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ and a vague prior $f_3^0(\boldsymbol{\theta}_3)$ on $\boldsymbol{\theta}_3$. The informative priors may be obtained either based on historical data from previous single-agent studies or by elicitation from the physician(s). In either case, we will assume that each parameter follows a gamma prior, and we denote the gamma distribution with mean ab and variance ab^2 by $G(a, b)$.

3.2 *Priors Based on Historical Data*

If dose-toxicity data are available from previous single-agent studies, they may be used to obtain priors on $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ for use in the planned trial. For $i = 1, 2$, denote the toxicity indicator and dose of the kth patient in the historical trial of agent i by $(\mathcal{Y}_{i,k}, x_{i,k})$ and $\mathcal{Z}_i = \{(\mathcal{Y}_{i,k}, x_{i,k}), k = 1, \ldots, n_i)\}$, where n_i is the number of patients. The likelihood of $\mathcal{Z}_i$ is

$$f_i(\mathcal{Z}_i \mid \boldsymbol{\theta}_i) = \prod_{k=1}^{n_i} \{\pi_i(x_{i,k}, \boldsymbol{\theta}_i)\}^{\mathcal{Y}_{i,k}} \{1 - \pi_i(x_{i,k}, \boldsymbol{\theta}_i)\}^{1-\mathcal{Y}_{i,k}}. \tag{5}$$

Denoting a vague prior on $\boldsymbol{\theta}_i$ that may be assumed before the historical data are observed by f_i^0 for $i = 1, 2$, the posterior of $\boldsymbol{\theta}_i$ given the historical data is then $f_i(\boldsymbol{\theta}_i \mid \mathcal{Z}_i) \propto f_i(\mathcal{Z}_i \mid \boldsymbol{\theta}_i) f_i^0(\boldsymbol{\theta}_i)$, and the prior used at the start of the two-agent trial is

$$f(\boldsymbol{\theta} \mid \mathcal{Z}_1, \mathcal{Z}_2) = f_1(\boldsymbol{\theta}_1 \mid \mathcal{Z}_1) f_2(\boldsymbol{\theta}_2 \mid \mathcal{Z}_2) f_3^0(\boldsymbol{\theta}_3). \tag{6}$$

3.3 *Elicited Priors*

When individual patient data from trials of the single agents are not available, informative priors $f_1(\boldsymbol{\theta}_1)$ and $f_2(\boldsymbol{\theta}_2)$ must be elicited from the physician(s). This may be done in various ways, with the particular elicitation method tailored to the clinical setting and physicians' level of technical expertise. We employed the following method in the Gem/CTX trial. Temporarily restrict attention to one agent and suppress subscripts, so that the two parameters of interest are α and β, and $\alpha x^\beta = \pi(\mathbf{x}, \boldsymbol{\theta})/\{1 - \pi(\mathbf{x}, \boldsymbol{\theta})\}$. We assume that $\alpha \sim G(a_1, a_2)$ and $\beta \sim G(b_1, b_2)$, so that each single-agent prior has four hyperparameters. The following four questions pertain to the toxicity probability curve of a single agent. Before asking these questions, the definition of "toxicity" and the target probability π^* must be established. The technical objective is to obtain four equations in the parameters $\{a_1, a_2, b_1, b_2\}$ having a solution that provides a corresponding prior on α, β and hence on $\pi(x, \boldsymbol{\theta})$.

1. What is the highest dose having negligible toxicity? That is, what is the highest dose that is almost certain to be toxic in less than 5% of patients?
2. What is the targeted dose that will have on average $100\pi^*$ toxicities?

3. What dose above the AD has a prohibitively high toxicity rate, say 60%?
4. What is the smallest dose above the AD that you are almost certain has toxicity rate above the targeted $100\pi^*$?

In practice, the physician(s) may wish to use specific numerical values different from those given above. For example, they may replace the 5% in question 1 with some other small percentage or, similarly, replace the 60% in question 3 with a different large percentage. This is not only acceptable, but desirable, and it should lead to more-accurate priors. The doses $d^{(1)}, d^{(2)}, d^{(3)} = d^*$, and $d^{(4)}$ given as answers to these questions may be used to obtain the following four probability statements. Denote $z_j = d^{(j)}/d^*$ and $g(\eta) = \eta/(1+\eta)$.

$$\Pr\left\{g\left(\alpha z_1^{\beta}\right) < 0.05\right\} = 0.99 \tag{7}$$

$$\mathrm{E}(\alpha) = a_1 a_2 = \pi^*/(1-\pi^*) \tag{8}$$

$$\mathrm{E}\left(\alpha z_3^{\beta}\right) = a_1 a_2 \mathrm{E}\left(z_3^{\beta}\right) = 0.60/0.40 \tag{9}$$

$$\Pr\left\{g\left(\alpha z_4^{\beta}\right) > \pi^*\right\} = 0.99 \tag{10}$$

Assuming that, *a priori*, $\alpha \sim G(a_1, a_2)$ and $\beta \sim G(b_1, b_2)$, equations (7)–(10) may be solved numerically for a_1, a_2, b_1, and b_2. This process is carried out twice, once for each single agent. Table 1 summarizes the elicited values for Gem and CTX as single agents.

Vague priors on α_3 and β_3 are appropriate because nothing is known *a priori* about interactions between the two agents. A preliminary sensitivity analysis of the priors on α_3 and β_3 showed that either large values of $\mathrm{E}(\beta_3)$ or values of $\mathrm{E}(\alpha_3)$ differing substantially from 1 give unrealistically skewed priors on $\pi(\mathbf{x}, \boldsymbol{\theta})$ that, in turn, cause the dose-finding algorithm to behave pathologically. For example, $\mathrm{E}(\beta_3) \geq 2$ puts so much prior probability mass on large doses that the final selected doses are higher than desired. The effects of the variances are less pronounced for moderately large values ($\geq$3), and the range 3 to 10 gives very reasonable behavior. We thus used $\mathrm{E}(\beta_3) = 0.05$, $\mathrm{E}(\alpha_3) = 1$, and $\mathrm{var}(\alpha_3) = \mathrm{var}(\beta_3) = 3$, and we recommend these or similar values for general application of the method.

Table 1
Elicited doses and priors for gemcitabine and cyclophosphamide as single agents

	Gemcitabine		Cyclophosphamide	
	Elicited dose	Median (π)	Elicited dose	Median (π)
d_1	600	.0018	350	.0055
$d_2 = d^*$	1200	.2578	600	.2687
d_3	1400	.5420	700	.5463
d_4	2000	.9361	800	.7677
	Mean	Variance	Mean	Variance
α	.4286	.1054	.4286	.0791
β	7.6494	5.7145	7.8019	3.9933

4. Dose-Finding Algorithm

4.1 *Structure of L_1 and L_2*

Dose-finding in stage 1 is done on the fixed line segment, L_1, illustrated in Figure 1(a). To determine L_1, the physician(s) first must choose a combination, $\mathbf{x}^{(1)}$ that will be the lowest, hence least toxic, dose pair considered in stage 1. The main criteria for choosing $\mathbf{x}^{(1)}$ are that its prior mean toxicity, $\mathrm{E}\{\pi(\mathbf{x}^{(1)}, \boldsymbol{\theta})\}$, must be low relative to π^*, but $\mathbf{x}^{(1)}$ must not be so low that it is very unlikely to be therapeutically effective. If $d_{i,1}$ is the dose of agent i elicited by question 1, hence thought to have negligible toxicity, then a reasonable choice is $\mathbf{x}^{(1)} = (d_{1,1}/d_1^*, d_{2,1}/d_2^*)$. At the other extreme, given that $x_{0,1}^*$ and $x_{0,2}^*$ are the prior single-agent ADs, the combination $\mathbf{x}_0^* = (x_{0,1}^*, x_{0,2}^*)$ is likely to be unacceptably toxic. Thus, a safety requirement is that $x_1^{(1)} < x_{0,1}^*$ and $x_2^{(1)} < x_{0,2}^*$, and in practice $x_i^{(1)}$ should be well below $x_{0,i}^*$ for each $i = 1, 2$. We define L_1 to be the straight line segment from $\mathbf{x}^{(1)}$ to $\mathbf{x}_0^*$. The physician(s) then must specify a set $D_1 = \{\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(k)}\}$ of dose combinations along L_1 where dose-finding in stage 1 is done initially.

Let $\mathbf{Z}_n$ denote the dose and toxicity data from the first n patients in the trial, for $n \leq N$. Due to the monotonicity of $\pi(\mathbf{x}, \boldsymbol{\theta})$ in $\mathbf{x}$, the posterior expected toxicity probability $\bar{\pi}_n(\mathbf{x}) = \mathrm{E}\{\pi(\mathbf{x}, \boldsymbol{\theta}) \mid \mathbf{Z}_n\}$ is monotone increasing in x_1 and x_2, hence is monotone along L_1. While $\mathbf{x}^{(1)}$ is the lowest dose in stage 1, in practice the physician(s) may decide to treat the first cohort at $\mathbf{x}^{(2)}$ or $\mathbf{x}^{(3)}$, with $\mathbf{x}^{(1)}$ included as a fallback option if $\mathbf{x}^{(2)}$ turns out to be unexpectedly toxic. In practice, the elements of D_1 may be equally spaced with $\mathbf{x}^{(k)} = (x_{0,1}^*, x_{0,2}^*)$, or possibly with $\mathbf{x}^{(k)}$ a pair of smaller values. Denoting $L_{2,n} = L_2(\pi^*, \mathbf{Z}_n)$, let $\mathbf{x}_n^* = L_1 \cap L_{2,n}$ be the unique point on L_1 having mean posterior toxicity probability π^*. Unlike the fixed line segment L_1, $L_{2,n}$ changes randomly as the data from each new cohort are added to $\mathbf{Z}_n$. Since $\pi(\mathbf{x}, \boldsymbol{\theta})$ and hence $\bar{\pi}_n(\mathbf{x})$ is increasing in both x_1 and x_2, it follows that $L_{2,n}$ lies entirely within the upper left and lower right quadrants of the Cartesian plane having origin $\mathbf{x}_n^*$. Denote, respectively, the portions of $L_{2,n}$ in these quadrants by $L_{2,n}^{\uparrow \mathrm{left}}$ and the $L_{2,n}^{\downarrow \mathrm{right}}$.

4.2 *A Two-Stage Algorithm*

Let c = cohort size. Once informative priors on $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$, the design parameters $\{n_1, n_2, c, \pi^*\}$, and L_1 and D_1 all have been specified, the trial may be conducted. The following two-stage algorithm aims to choose one or more acceptable dose combinations for later study.

Stage 1: Treat the first cohort at the lowest dose, $\mathbf{x}^{(1)}$. Thereafter, treat each cohort at the dose $\mathbf{x}^{(j)} \in D_1$ that minimizes $|\bar{\pi}_n(\mathbf{x}^{(j)}) - \pi^*|$, subject to the constraint that no untried dose level in D_1 may be skipped when escalating. Once the first toxicity is observed, say at $\mathbf{x}^{(r)}$, expand the set of allowable doses by adding doses midway between the consecutive pairs above $\mathbf{x}^{(r)}$, and also all doses below $\mathbf{x}^{(r)}$. Denoting $D_1^* = D_1 \cup \{\frac{1}{2}(\mathbf{x}^{(r)} + \mathbf{x}^{(r+1)}), \ldots, \frac{1}{2}(\mathbf{x}^{(k-1)} + \mathbf{x}^{(k)})\}$, and the line segment from $\mathbf{x}^{(1)}$ to $\mathbf{x}^{(r)}$ by L_1^*, expand the set of allowable doses from D_1 to $D_1^* \cup L_1^*$. As before, no untried dose combination in D_1^* may be skipped when escalating. When n_1 patients have been treated, proceed to stage 2.

Stage 2: Treat successive cohorts at dose combinations selected alternately from $L_{2,n}^{\uparrow\text{left}}$ and $L_{2,n}^{\downarrow\text{right}}$. Stop when a total of $N = n_1 + n_2$ patients have been treated.

Since L_1 is one-dimensional, stage 1 is similar to a conventional single-agent Phase I trial. In stage 1, untried dose combinations may not be skipped in D_1, or in the expanded set D_1^* once a toxicity has been observed, due to the ethical desire to avoid overdosing early in the trial, when uncertainty is greatest. The stage 1 sample size should be large enough to establish with a reasonable degree of certainty how high the contour $L_{2,n}$ is likely to be on the $\pi(\mathbf{x}, \boldsymbol{\theta})$ surface. This motivated our choice of $n_1 = 20$ in the Gem/CTX trial, since this is a typical Phase I sample size. The risk of overdosing is less of a concern in stage 2, where all patients are treated at doses on $L_{2,n}$. In contrast with conventional Phase I designs, where dose-finding is restricted to a finite number of dose levels, in the latter part of stage 1 and throughout stage 2, we choose dose pairs from a continuum. While conventional Phase I practice is convenient, and is safer than choosing doses from a continuum, there is no reason to assume that the final optimal dose must be one of the arbitrarily chosen dose levels. Consequently, we preserve safety by escalating in discrete steps initially, until the first toxicity is encountered, but then allowing dose pairs to be chosen from the line segment L_1^* in the latter part of stage 1 and from the contour $L_{2,n}$ in stage 2.

4.3 *Criteria for Stage 2*

In stage 2, successive dose combinations may be chosen from $L_{2,n}^{\uparrow\text{left}}$ and $L_{2,n}^{\downarrow\text{right}}$ in several ways. This is because, based on the toxicity criterion (1) alone, all dose pairs $\mathbf{x} \in L_{2,n}$ are equally acceptable. The following algorithm for stage 2 is based on (1) and two additional criteria. First, it is clinically desirable to choose each successive dose combination to maximize the potential to kill cancer cells, while maintaining the posterior mean toxicity rate at π^*. At the same time, it is also desirable to maximize the amount of information about $\pi(\mathbf{x}, \boldsymbol{\theta})$ obtained from each new cohort. To formalize the first goal, we temporarily assume for simplicity that, in terms of cancer-killing effect, one unit change in x_1 is equivalent to one unit change in x_2. This says that the amount of potential increase in cancer-killing effect obtained by moving from $\mathbf{x}_n^*$ to $\mathbf{x}$ on $L_{2,n}$ is proportional to the total increase in dose, $\mathcal{K}(\mathbf{x}, \mathbf{x}_n^*) = (x_1 - x_{n,1}^*) + (x_2 - x_{n,2}^*)$. Any dose combination $\mathbf{x} \in L_{2,n}^{\uparrow\text{left}} - \{\mathbf{x}_n^*\}$ must satisfy the inequalities $x_1 - x_{n,1}^* < 0 < x_2 - x_{n,2}^*$ and, similarly, any $\mathbf{x} \in L_{2,n}^{\downarrow\text{right}} - \{\mathbf{x}_n^*\}$ must satisfy $x_1 - x_{n,1}^* > 0 > x_2 - x_{n,2}^*$. That is, the two summands of $\mathcal{K}(\mathbf{x}, \mathbf{x}_n^*)$ must have opposite signs. It follows that choosing a dose combination $\mathbf{x}$ from either $L_{2,n}^{\uparrow\text{left}}$ or $L_{2,n}^{\downarrow\text{right}}$ to maximize cancer-killing potential compared to $\mathbf{x}_n^*$ amounts to choosing $\mathbf{x}$ so that the negative summand of $\mathcal{K}(\mathbf{x}, \mathbf{x}_n^*)$ is small relative to its positive summand. The assumption that the two agents have the same cancer-killing effect per standard dose unit may be relaxed if prior data quantifying the relative effects of the two agents are available, say, from an animal experiment or a study of *in vitro* effects on cancer cell cultures. If λ is the cancer-killing effect of one standard dose unit of agent 1 relative to one standard dose unit of agent 2, then the total gain in cancer-killing effect obtained by moving from $\mathbf{x}_n^*$ to $\mathbf{x}$ may be defined as

$$\mathcal{K}_\lambda\big(\mathbf{x}, \mathbf{x}_n^*\big) = \lambda\big(x_1 - x_{n,1}^*\big) + \big(x_2 - x_{n,2}^*\big). \tag{11}$$

This is similar to the idea of total equivalent dose used by Simon and Korn (1990).

A different criterion for selecting a dose combination $\mathbf{x} \in L_{2,n}$ is the amount of information provided by the data obtained from treating the next cohort at $\mathbf{x}$. Since the likelihood for a patient treated at $\mathbf{x}$ is $\pi(\mathbf{x}, \boldsymbol{\theta})^Y\{1 - \pi(\mathbf{x}, \boldsymbol{\theta})\}^{1-Y}$, denoting $\pi(\mathbf{x}, \boldsymbol{\theta})^{(j)} = \partial\pi(\mathbf{x}, \boldsymbol{\theta})/\partial\theta_j$ where θ_j is the jth entry of $\boldsymbol{\theta}$, the Fisher information matrix associated with treating the patient at dose $\mathbf{x}$ is

$$\begin{aligned} I(\mathbf{x}, \boldsymbol{\theta})^{(6\times 6)} &= \big\{I^{(j,k)}(\mathbf{x}, \boldsymbol{\theta})\big\} \\ &= \left[\frac{\pi(\mathbf{x}, \boldsymbol{\theta})^{(j)}\pi(\mathbf{x}, \boldsymbol{\theta})^{(k)}}{\pi(\mathbf{x}, \boldsymbol{\theta})\{1 - \pi(\mathbf{x}, \boldsymbol{\theta})\}}\right]. \end{aligned} \tag{12}$$

The dose on $L_{2,n}$ may be chosen for the next cohort to maximize the posterior expectation of the log determinant of the Fisher information matrix given the current data,

$$\mathcal{I}_n(\mathbf{x}) = \mathrm{E}\left[\log\{\det I(\mathbf{x}, \boldsymbol{\theta})\} \mid \mathbf{Z}_n\right]. \tag{13}$$

To compute $\mathcal{I}_n(\mathbf{x})$, denote the numerator of $\pi(\mathbf{x}, \boldsymbol{\theta})$ by $\gamma(\mathbf{x}, \boldsymbol{\theta})$. Since each partial derivative $\pi^{(j)}(\mathbf{x}, \boldsymbol{\theta})$ is of the form $f^{(j)}/\{1 + \gamma(\mathbf{x}, \boldsymbol{\theta})\}^2$, it follows that $I^{(j,k)}(\mathbf{x}, \boldsymbol{\theta}) = f^{(j)}f^{(k)}/[\gamma(\mathbf{x}, \boldsymbol{\theta})\{1 + \gamma(\mathbf{x}, \boldsymbol{\theta})\}^2]$ for $j, k = 1, \ldots, 6$. Writing $f^{(\alpha_j)}$ for the numerator of $\partial\pi/\partial\alpha_j$, and so on, the six partial derivatives are $f^{(\alpha_j)} = x_j^{\beta_j}$ and $f^{(\beta_j)} = \log(x_j)(\alpha_j x_j^{\beta_j} + \alpha_3\beta_3 x_1^{\beta_1\beta_3}x_2^{\beta_2\beta_3})$ for $j = 1, 2$, $f^{(\alpha_3)} = x_1^{\beta_1\beta_3}x_2^{\beta_2\beta_3}$, and $f^{(\beta_3)} = \alpha_3\log(x_1^{\beta_1}x_2^{\beta_2}) \times (x_1^{\beta_1}x_2^{\beta_2})^{\beta_3}$. The criterion (13) is similar to that used in Bayesian D-Optimal design (Chaloner and Larntz, 1989; Zocchi and Atkinson, 1999), although our application is not formally an example of this method.

Because $\mathcal{K}_\lambda(\mathbf{x}, \mathbf{x}_n^*)$ and $\mathcal{I}_n(\mathbf{x})$ take on values on different scales, combining them is not straightforward. We address this problem by first choosing the optimal dose pair under each criterion and then averaging these two doses. If the next cohort in stage 2 is to be treated at a dose chosen from $L_{2,n}^{\uparrow\text{left}}$, then, restricting attention to $L_{2,n}^{\uparrow\text{left}}$, let $\mathbf{x}(\mathcal{K}_\lambda, L_{2,n}^{\uparrow\text{left}})$ be the dose pair that maximizes the cancer-killing criterion function $\mathcal{K}_\lambda(\mathbf{x}, \mathbf{x}_n^*)$, and let $\mathbf{x}(\mathcal{I}_n, L_{2,n}^{\uparrow\text{left}})$ be the pair maximizing the information $\mathcal{I}_n(\mathbf{x})$. The next cohort is treated at the dose, $\mathbf{x}(\mathcal{K}_\lambda, \mathcal{I}_n, L_{2,n}^{\uparrow\text{left}})$, that is midway between $\mathbf{x}(\mathcal{K}_\lambda, L_{2,n}^{\uparrow\text{left}})$ and $\mathbf{x}(\mathcal{I}_n, L_{2,n}^{\uparrow\text{left}})$ on $L_{2,n}^{\uparrow\text{left}}$. If the midway dose is not on $L_{2,n}$, then we define $\mathbf{x}(\mathcal{K}_\lambda, \mathcal{I}_n, L_{2,n}^{\uparrow\text{left}})$ as the projection of the midway dose onto $L_{2,n}$. The dose $\mathbf{x}(\mathcal{K}_\lambda, \mathcal{I}_n, L_{2,n}^{\downarrow\text{right}})$ chosen from $L_{2,n}^{\downarrow\text{right}}$ is defined similarly.

Based on the final data, $\mathbf{Z}_N$, we select three dose combinations: $\mathbf{x}_N^{\downarrow\text{right}} = \mathbf{x}(\mathcal{K}_\lambda, \mathcal{I}_N, L_{2,N}^{\downarrow\text{right}})$, $\mathbf{x}_N^{\uparrow\text{left}} = \mathbf{x}(\mathcal{K}_\lambda, \mathcal{I}_N, L_{2,N}^{\uparrow\text{left}})$, and the unique combination $\mathbf{x}_N^{\text{middle}} = L_1 \cap L_{2,N}$. Roughly speaking, $\mathbf{x}_N^{\text{middle}}$ contains substantial quantities of both agents, $\mathbf{x}_N^{\downarrow\text{right}}$ contains more of agent 1 and less agent 2, and $\mathbf{x}_N^{\uparrow\text{left}}$ contains more of agent 2 and less agent 1. These three combinations may be studied in a subsequent randomized trial in which tumor response or possibly survival time or disease-free survival time are the primary outcomes.

4.4 *Steps for Application of the Method*

Given the complexity of the model and method, it is worthwhile to list the necessary steps for application. Initially, the trial entry criteria, the patient accrual rate, the two agents and their dose ranges, and a target π^* must be established. Then priors should be determined, as described in Section 3. The design parameters $\{n_1, n_2, c\}$ may be chosen by using simulation to evaluate the design over a range of possible values. We illustrate how this may be done below. Once this structure is in place, the trial may be conducted. To facilitate broad application, we have made computer programs for simulation and trial conduct freely available at `biostatistics.mdanderson.org`.

5. Computational Methods

Implementation of the proposed method is computationally intensive in that it involves repeated calculation of the posterior mean toxicity surface, $\{\bar{\pi}_n(\mathbf{x}) : 0 < x_1 < 1, 0 < x_2 < 1\}$, after each new cohort has been evaluated. Consequently, the use of computationally efficient algorithms is critical for a successful implementation. We use Markov chain Monte Carlo (MCMC) simulation (Gelfand and Smith, 1990; Gilks et al., 1993), accelerated by the use of the maximum *a posteriori* (MAP) estimate as an initial starting point, and an asymptotic posterior approximation to define transition probabilities in the Markov chain. We first find the MAP by maximization of the unstandardized posterior $f(\boldsymbol{\theta} \mid \mathbf{Z}_n)$, given in (4). For later use, in the construction of a proposal distribution, we compute the negative inverse Hessian $\hat{\Sigma}$ of the unstandardized log posterior distribution, evaluated at $\hat{\boldsymbol{\theta}}$. Starting with initial state $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$, we then initiate an MCMC simulation. We use transition probabilities defined by alternate use of a random walk proposal and an independence chain. The independence chain uses a proposal based on an $N(\hat{\boldsymbol{\theta}}, \hat{\Sigma})$ approximation to the joint posterior distribution. After an initial transient of 50 iterations, we evaluate the toxicity surface after each 20th iteration, with a total of 2000 iterations. The ergodic average of these surfaces provides an estimate of the desired posterior expected toxicity probability $\bar{\pi}_n(\mathbf{x})$. To obtain these values over the entire domain of $\mathbf{x}$, we evaluate $\bar{\pi}_n(\mathbf{x})$ on a grid of doses. During stage 1, we need only evaluate toxicities for the doses on the grid D_1 in L_1. An additional complication arises when the first toxicity is observed and we switch to the expanded set $D_1^* \cup L_1^*$, since L_1^* is a continuous interval. To deal with this, we still constrain MCMC evaluation of $\bar{\pi}_n(\mathbf{x})$ to the doses on the finite grid D_1^*, and we use interpolation for any combination falling between doses in D_1^*. A problem in the calculation of $L_2(\pi^*, \mathbf{Z}_n)$ is that this contour changes randomly each time the posterior is updated. To deal with this, we first run the MCMC and save the simulated (approximate) posterior Monte Carlo sample of $\boldsymbol{\theta}$ values, and then search for $L_2(\pi^*, \mathbf{Z}_n)$ on a two-dimensional grid over (x_1, x_2). In the course of this search, each time we consider a grid point $\mathbf{x} = (x_1, x_2)$ we compute $\bar{\pi}_n(\mathbf{x})$ as the appropriate ergodic average over the saved Monte Carlo sample, evaluating $\pi(\boldsymbol{\theta}, \mathbf{x})$ for that dose pair $\mathbf{x}$ only. To further reduce the computational burden, we update $\bar{\pi}_n(\mathbf{x})$ only after every other cohort, i.e., after assigning one dose from $\mathrm{L}_{2,n}^{\uparrow\text{left}}$ and one dose from $\mathrm{L}_{2,n}^{\downarrow\text{right}}$.

6. Application to the Gem/CTX Trial

In the Gem/CTX trial, patients are treated in cohorts of size $c = 2$ with $n_1 = 20$ patients in stage 1 and $n_2 = 40$ in stage 2. The line L_1 was defined to be the segment connecting the standardized doses $\mathbf{x}^{(1)} = (0.12, 0.12)$ and $(x_{0,1}^*, x_{0,2}^*) = (1, 1)$, which correspond to the (Gem, CTX) dose pairs (144, 72) and (1200, 600) mg/m^2. Ten dose pairs were used for stage 1, with the eight intermediate doses on L_1 chosen conservatively by the physician, R. Millikan, so that anticipated toxicity probabilities were dominated by the corresponding prior mean values. Specifically, $D_1 = \{(x, x) : x = 0.12, 0.25, 0.40, 0.52, 0.61, 0.70, 0.78, 0.86, 0.93, 1\}$. The first cohort was treated at (0.25, 0.25), with no untried dose skipped while escalating. Although we did not do so in the Gem/CTX trial, one also may impose an additional safety rule to stop the trial if the lowest dose is too toxic, for example, if $\Pr\{\pi(\mathbf{x}^{(1)}, \boldsymbol{\theta}) > \pi^* \mid \text{data}\}$ exceeds an upper probability bound such as 0.90 or 0.95.

Figure 1(b) gives the 30 successive dose pairs chosen by this design in a hypothetical trial, including the final $L_{2,N}$ based on the data from all 60 patients. Each two-patient cohort is numbered by the order in which its patients were treated in the trial. The sequence of locations of the dose pairs 11 through 30, chosen in stage 2 of the trial, as well as their variability, illustrate the manner in which L_2 changes as the data from successive cohorts are obtained and the posterior is updated. In this context, it is important to bear in mind that each selected $\mathbf{x}$ in stage 2 is the average of the doses that optimize the cancer-killing criterion and Fisher information, and that the posterior is updated after each pair of "upper left and lower right" cohorts as a computational convenience.

To assess average behavior, we simulated the trial under each of five dose-toxicity scenarios, illustrated in terms of their contours of constant toxicity probability in Figure 2. Table 2 summarizes the selected dose pairs $\mathbf{x}_N^{\text{middle}}$, $\mathbf{x}_N^{\downarrow\text{right}}$, and $\mathbf{x}_N^{\uparrow\text{left}}$ and corresponding values of the true $\pi(\mathbf{x})$ and the posterior mean, $\bar{\pi}_n(\mathbf{x})$, at each selected $\mathbf{x}$. The values $\mathbf{x}_{n_1}^{\text{middle}}$ and $\mathbf{x}_N^{\text{middle}}$ together show how much the selected dose on L_1 changes by carrying out the second stage. The desired contour under scenario 1 is the diagonal line in the middle of the $\mathbf{x}$ domain, running roughly from $(1, 0)$ to $(0, 1)$, and the toxicity probability surface is not very steep. Under scenario 2, the desired dose pairs lie on a contour with one or both doses very close to their maximum values; scenario 3 is the opposite case, with all acceptable dose pairs having one or both doses very close to their minimum values. Scenario 4 is the most complex, with S-shaped contours for π at or near the target 0.30.

Under scenarios 1, 3, and 4, the algorithm has very accurate average behavior, choosing dose pairs with both $\pi(\mathbf{x})$ and posterior mean $\bar{\pi}_n(\mathbf{x})$ on average within 0.04 of the target 0.30. The largest deviation is the mean value 0.22 of $\pi(\mathbf{x}^{\uparrow\text{left}})$ under scenario 2, which is quite reasonable given that the toxicity surface rises very rapidly with x_1 and x_2, and thus the contours where $0.20 \le \pi(\mathbf{x}) \le 0.40$ are very close together. The fact that the average toxicity probability is below, rather than above, the target in this dangerous case is reassuring, in terms of the procedure's safety. The large drop in $\mathrm{Var}\{\pi(\mathbf{x}) \mid \mathbf{Z}_n\}$, seen in all scenarios, as n increases from $n_1 = 20$ to $N = 60$ shows the gain in reliability from conducting

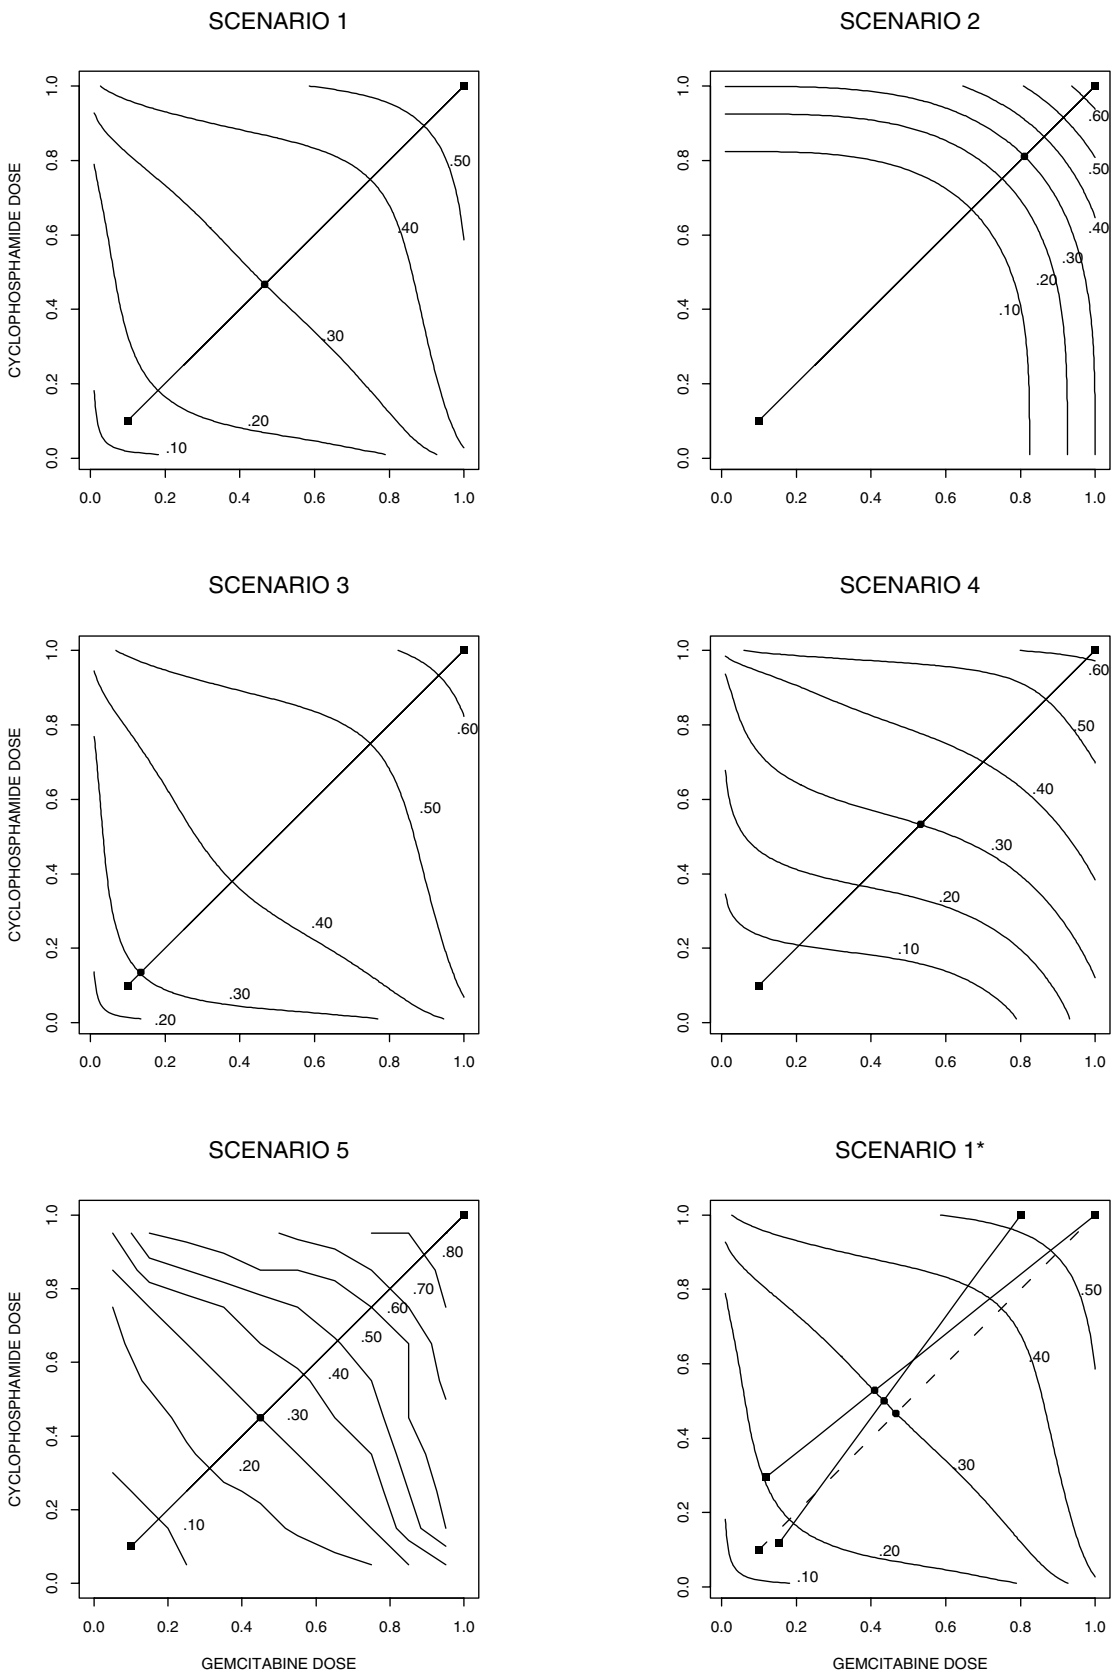

Figure 2. Scenarios for the probability of toxicity as a function of the cyclophosphamide and gemcitabine doses in the prostate cancer trial. For each agent, the doses are given in standard units on the domain from 0 to 1.

Table 2
Dose combinations selected for Gem/CTX. Standard deviations are given as subscripts.

	$\mathbf{x}_{n_1}^{\text{middle}}$	$\mathbf{x}_N^{\text{middle}}$	$\mathbf{x}_N^{\downarrow\text{right}}$	$\mathbf{x}_N^{\uparrow\text{left}}$
		Scenario 1		
Selected dose $\mathbf{x}$	$(0.47_{0.16},0.47_{0.16})$	$(0.50_{0.10},0.50_{0.10})$	$(0.66_{0.11},0.24_{0.19})$	$(0.34_{0.13},0.67_{0.10})$
$\bar{\pi}_n(\mathbf{x})$, true $\pi(\mathbf{x})$	$0.30_{0.05}$, $0.30_{0.06}$	$0.30_{0.01}$, $0.31_{0.04}$	$0.30_{0.01}$, $0.28_{0.07}$	$0.30_{0.01}$, $0.32_{0.04}$
$\text{Var}\{\pi(\mathbf{x}) \mid \mathbf{Z}_n\} \times 10^2$	$0.82_{0.32}$	$0.33_{0.12}$	$0.43_{0.19}$	$0.37_{0.18}$
		Scenario 2		
Selected dose $\mathbf{x}$	$(0.81_{0.04},0.81_{0.04})$	$(0.80_{0.03},0.80_{0.03})$	$(0.92_{0.05},0.72_{0.05})$	$(0.58_{0.12},0.86_{0.05})$
$\bar{\pi}_n(\mathbf{x})$, true $\pi(\mathbf{x})$	$0.30_{0.01}$, $0.30_{0.08}$	$0.30_{0.01}$, $0.28_{0.05}$	$0.29_{0.01}$, $0.35_{0.05}$	$0.30_{0.01}$, $0.22_{0.05}$
$\text{Var}\{\pi(\mathbf{x}) \mid \mathbf{Z}_n\} \times 10^2$	$0.78_{0.14}$	$0.29_{0.03}$	$0.41_{0.10}$	$0.38_{0.11}$
		Scenario 3		
Selected dose $\mathbf{x}$	$(0.28_{0.14},0.28_{0.14})$	$(0.23_{0.17},0.23_{0.17})$	$(0.46_{0.21},0.07_{0.17})$	$(0.10_{0.14},0.40_{0.22})$
$\bar{\pi}_n(\mathbf{x})$, true $\pi(\mathbf{x})$	$0.37_{0.07}$, $0.36_{0.05}$	$0.30_{0.02}$, $0.33_{0.08}$	$0.29_{0.02}$, $0.29_{0.08}$	$0.30_{0.02}$, $0.31_{0.08}$
$\text{Var}\{\pi(\mathbf{x}) \mid \mathbf{Z}_n\} \times 10^2$	$0.94_{0.28}$	$0.40_{0.22}$	$0.48_{0.23}$	$0.44_{0.26}$
		Scenario 4		
Selected dose $\mathbf{x}$	$(0.57_{0.12},0.57_{0.12})$	$(0.56_{0.07},0.56_{0.07})$	$(0.70_{0.10},0.38_{0.15})$	$(0.39_{0.10},0.68_{0.08})$
$\bar{\pi}_n(\mathbf{x})$, true $\pi(\mathbf{x})$	$0.29_{0.05}$, $0.32_{0.07}$	$0.30_{0.01}$, $0.32_{0.04}$	$0.30_{0.005}$, $0.26_{0.08}$	$0.30_{0.01}$, $0.34_{0.04}$
$\text{Var}\{\pi(\mathbf{x}) \mid \mathbf{Z}_n\} \times 10^2$	$0.75_{0.29}$	$0.35_{0.09}$	$0.43_{0.16}$	$0.39_{0.15}$
		Scenario 5		
Selected dose $\mathbf{x}$	$(0.47_{0.11},0.47_{0.11})$	$(0.45_{0.09},0.45_{0.09})$	$(0.67_{0.13},0.17_{0.16})$	$(0.29_{0.13},0.63_{0.11})$
$\bar{\pi}_n(\mathbf{x})$, true $\pi(\mathbf{x})$	$0.31_{0.03}$, $0.33_{0.10}$	$0.30_{0.01}$, $0.32_{0.10}$	$0.30_{0.01}$, $0.27_{0.13}$	$0.30_{0.01}$, $0.32_{0.10}$
$\text{Var}\{\pi(\mathbf{x}) \mid \mathbf{Z}_n\} \times 10^2$	$0.92_{0.31}$	$0.38_{0.14}$	$0.50_{0.22}$	$0.43_{0.21}$

stage 2. The largest differences between $\{\pi(\mathbf{x}), \bar{\pi}_{n_1}(\mathbf{x})\}$ for $\mathbf{x} = \mathbf{x}_{n_1}^{\text{middle}} = \mathbf{x}_{20}^{\text{middle}}$, at the end of stage 1, and the final values for $\mathbf{x} = \mathbf{x}_N^{\text{middle}} = \mathbf{x}_{60}^{\text{middle}}$, are seen under scenario 3. This seems reasonable, given that the desired contour is located at very-low-dose values in this case.

Table 3 summarizes the numbers of patients treated and toxicities under each scenario, broken down by the true toxicity probability $\pi(\mathbf{x})$. Under the scenarios 1–4, on average, between 36 (60%) and 52 (87%) of the 60 patients are treated at a dose with $|\pi(\mathbf{x}) - 0.30| < 0.10$, at most 19 (15%) are treated at a dose with $0.41 \leq \pi(\mathbf{x}) \leq 0.50$, only 1.4 patients are treated at a dose with $0.51 \leq \pi(\mathbf{x}) \leq 0.70$, and none are treated at a more toxic dose. The dispersion of patients over the range of true toxicity probabilities is slightly higher under scenario 5, with on average 9.54 patients (16%) treated at doses with toxicities above 0.40, although most of these were in the 0.41–0.50 range. The algorithm thus appears to be very safe.

7. Robustness

Parametric models other than that used here may be considered, subject to the admissibility conditions given in Section 3. For example, the interaction term $\alpha_3(x_1^{\beta_1}x_2^{\beta_2})^{\beta_3}$ may be simplified to $\alpha_3(x_1^{\beta_1}x_2^{\beta_2}), \alpha_3(x_1x_2)^{\beta_3}$, or $\alpha_3x_1x_2$. An alternative model is to first define explicit forms for $\pi_1 = \pi_1(x_1, \boldsymbol{\theta}_1)$ and $\pi_2 = \pi_2(x_2, \boldsymbol{\theta}_2)$, such as $\pi_i(x_i, \boldsymbol{\theta}_i) = g^{-1}\{\log(\alpha_i) + \beta_i \log(x_i)\}$ for a link function g, and let $\pi(x_1, x_2, \boldsymbol{\theta})$ be a function of π_1 and π_2, such as $\lambda\{\pi_1 + (1-\pi_1)\pi_2\} + (1-\lambda)\{\pi_2 + (1-\pi_2)\pi_1\}$, for $0 < \lambda < 1$. Given several possible models, model selection may be based on Bayes factors (Kass and Raftery, 1995), or Bayesian model averaging methods (Madigan and Raftery, 1994), added to each step of the adaptive algorithm. While a thorough comparison of the method's behavior under alternative models would be quite useful, it is beyond the scope of the present article.

Some insight into the method's robustness is provided by the simulation results under scenario 5, which does not correspond to any parameter configuration of the underlying probability model. This is a more dangerous case in which the target contour where $\pi(\mathbf{x}) = 0.30$ is similar to that in scenario 1, but the toxicity probability surface is much steeper. Table 2 indicates that the algorithm's average ability to select doses with $\pi(\mathbf{x})$ close to the target is still quite good in this case, although the variability is greater. The average posterior mean and true toxicity probabilities at the selected doses are all between 0.27 and 0.33, and these results are nearly identical to those under scenario 1. Table 3 indicates that the algorithm's interim dose selections are somewhat more variable, although it is still quite safe.

To examine the method's sensitivity to the placement of L_1, we modified the design by using two additional versions of L_1, illustrated as solid lines in the lowest right corner graph of Figure 2, labeled SCENARIO 1*, with the original version of L_1 given as a dotted line. The first new stage 1 line, L_1^a, is steeper, running from (0.155, 0.12) to (0.80, 1.0), with D_1 consisting of these points plus eight intermediate points on L_1^a having the same x_2 entries as those used previously. Use of L_1^a might be motivated by the desire to decrease the upper limit on the dose of the first agent during stage 1, based on considerations of its toxicity as a single agent. In this case, the highest dose of Gem on L_1 is decreased from 1200 to 960 mg/m^2. We also considered the less steep line, L_1^b, runnning from (0.12, 0.296) to (1.0, 1.0), with D_1 having the same x_1 entries as those used previously. Use of L_1^b might be motivated by the desire to increase the lower limit on the dose of the second agent during stage 1, in this case increasing the lowest CTX dose on L_1 from from 72 to 178 mg/m^2. Simulations of these two

Table 3

Number of patients treated and number of toxicities, out of 60 *patients, as functions of the true probability of toxicity. Standard deviations are given as subscripts.*

	True probability of toxicity									
	0–0.10	0.11–0.20	0.21–0.30	0.31–0.40	0.41–0.50	0.51–0.60	0.61–0.70	0.71–0.80	0.81–0.90	0.91–1.0
					Scenario 1					
Treated	$0.4_{1.3}$	$7.0_{8.2}$	$24.9_{12.0}$	$26.8_{16.0}$	$0.7_{2.2}$	$0.3_{1.1}$	0	0	0	0
Toxicity	$0.01_{0.08}$	$0.9_{1.6}$	$5.9_{4.1}$	$10.6_{6.0}$	$0.4_{1.1}$	$0.2_{0.8}$	0	0	0	0
					Scenario 2					
Treated	$8.6_{1.4}$	$11.5_{6.8}$	$18.1_{5.3}$	$18.1_{5.0}$	$3.5_{4.4}$	$0.1_{0.5}$	0	0	0	0
Toxicity	$0.2_{0.4}$	$1.8_{2.1}$	$4.5_{1.7}$	$6.9_{2.1}$	$1.7_{2.1}$	$0.1_{0.4}$	0	0	0	0
					Scenario 3					
Treated	0	$6.9_{5.5}$	$24.8_{10.8}$	$18.2_{9.5}$	$9.1_{10.4}$	$0.8_{2.0}$	$0.2_{1.0}$	0	0	0
Toxicity	0	$0.7_{1.0}$	$6.3_{3.8}$	$6.9_{3.6}$	$4.6_{5.3}$	$0.4_{1.1}$	$0.1_{0.7}$	0	0	0
					Scenario 4					
Treated	$2.9_{2.9}$	$6.8_{4.0}$	$17.4_{7.3}$	$30.2_{9.3}$	$2.2_{3.4}$	$0.1_{0.6}$	$0.3_{1.5}$	0	0	0
Toxicity	$0.1_{0.3}$	$0.7_{1.0}$	$4.4_{2.8}$	$11.6_{3.6}$	$1.2_{1.9}$	$0.04_{0.3}$	$0.2_{0.8}$	0	0	0
					Scenario 5					
Treated	$1.7_{2.1}$	$15.2_{6.2}$	$15.2_{5.3}$	$18.3_{5.9}$	$7.5_{5.1}$	$1.5_{2.3}$	$0.2_{0.6}$	$0.1_{0.5}$	$0.04_{0.3}$	$0.2_{1.0}$
Toxicity	$0.1_{0.3}$	$1.9_{1.6}$	$3.8_{2.3}$	$7.3_{2.5}$	$4.2_{2.8}$	$1.1_{1.8}$	$0.1_{0.5}$	$0.1_{0.4}$	$0.03_{0.3}$	$0.2_{1.0}$

Table 4

Sensitivity of the posterior toxicity probability at each final selected dose to sample size and cohort size. Each case was simulated under scenario 4 *with* $n_1/N = 1/3$.

		Posterior 90% credible interval for $\pi(\mathbf{x}_N)$			$\lvert\bar{\pi}_N(\mathbf{x}) - \text{true } \pi(x)\rvert$		
N	c	$\mathbf{x}_N^{\uparrow\text{left}}$	$\mathbf{x}_N^{\text{middle}}$	$\mathbf{x}_N^{\downarrow\text{right}}$	$\mathbf{x}_N^{\uparrow\text{left}}$	$\mathbf{x}_N^{\text{middle}}$	$\mathbf{x}_N^{\downarrow\text{right}}$
36	1	0.189–0.425	0.182–0.423	0.170–0.429	0.059	0.051	0.081
	2	0.189–0.423	0.185–0.421	0.174–0.429	0.064	0.052	0.078
	3	0.191–0.437	0.181–0.432	0.160–0.446	0.093	0.102	0.106
48	1	0.195–0.419	0.196–0.414	0.181–0.421	0.053	0.039	0.066
	2	0.196–0.419	0.198–0.413	0.181–0.416	0.052	0.041	0.069
	3	0.194–0.424	0.191–0.418	0.177–0.424	0.088	0.093	0.097
60	1	0.202–0.413	0.205–0.404	0.188–0.410	0.047	0.032	0.059
	2	0.205–0.410	0.206–0.401	0.191–0.409	0.051	0.037	0.057
	3	0.196–0.409	0.200–0.407	0.186–0.411	0.097	0.103	0.104
72	1	0.210–0.407	0.210–0.399	0.191–0.412	0.042	0.034	0.068
	2	0.215–0.407	0.213–0.395	0.198–0.403	0.048	0.030	0.056
	3	0.209–0.407	0.209–0.398	0.191–0.407	0.092	0.095	0.091

modified versions of the algorithm under scenario 1 showed that with either L_1^a or L_1^b, all mean values of $\bar{\pi}_n(\mathbf{x})$ and $\pi(\mathbf{x})$ at the selected dose pairs were between 0.29 and 0.31, and each entry of each selected dose differed from the corresponding value under the original L_1 by 0.01 to 0.06. The method thus seems relatively insensitive to the placement of L_1.

Table 4 summarizes a simulation study of the method's sensitivity to N and c. This shows how the posterior variability of each $\pi(\mathbf{x}_N)$ drops with N. Average values of $|\bar{\pi}(\mathbf{x}_N) - 0.30|$ are insensitive to both N and c, hence not tabled; these equal 0.01 for $\mathbf{x}_N^{\text{middle}}$ and 0.02 to 0.04 for $\mathbf{x}_N^{\uparrow\text{left}}$ or $\mathbf{x}_N^{\downarrow\text{right}}$. In contrast, for each type of dose pair, $|\bar{\pi}(\mathbf{x}_N) - \text{true } \pi(\mathbf{x})|$ on average decreases slightly with N but increases sharply as c increases from 2 to 3. This is due to the fact that as c increases for given N, the total number of cohorts, N/c, decreases; hence information is available on fewer dose pairs. It thus appears that, in this setting, $c = 1$ or 2 provides a much more reliable dose pair selection than $c = 3$.

8. Discussion

We have proposed a model-based, two-stage Bayesian adaptive method for determining several acceptable dose pairs of two agents used together in a Phase I cancer chemotherapy trial. Our simulation study indicates that the method has desirable properties and is very safe. Because this type of dose-finding problem arises frequently in clinical oncology, the method is broadly applicable. An advantage of the method is that it is rooted in actual clinical practice. The requirement of having informative priors on $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$, the prior elicitation algorithm, and the goal of choosing several dose pairs are all motivated by our practical experience, as the Gem/CTX trial illustrates. While the method is computationally intensive,

in our experience, this has not presented any substantive problems in implementation.

Several practical extensions or modifications of the method are of interest. If either agent is biologic rather than cytotoxic, then the admissibility conditions and model must be changed to accommodate the possibility that π may not increase in each x_i. Although we have used cancer-killing potential and information as criteria for selecting doses in stage 2, other criteria, such as drug cost or feasibility, could be used. General problems are to optimize c, $(n_1,\ n_2)$ or, given N, the proportion $n_1/(n_1 + n_2)$. A fully Bayesian approach (Berry, 1995) would employ decision theory, with successive doses selected by maximizing the posterior mean of an appropriate gain function. A very useful generalization would be to base dose-finding on both response and toxicity, possibly by combining the methodology described here with the dose-finding method proposed by Thall and Russell (1998).

Résumé

Nous proposons un dispositif bayésien adaptatif à deux étapes pour la recherche d'une ou plusieurs combinaisons acceptables de doses de deux produits cytotoxiques utilisés conjointement dans un essai clinique de phase I. Cette méthode nécessite que chacun des deux produits ait été étudié séparément au préalable, ce qui est presque toujours le cas en pratique. De surcroît on fait l'hypothèse d'un modèle paramétrique pour l'évaluation de la probabilité de toxicité en fonction des deux doses. Des a priori informatifs pour les paramètres caractérisant les courbes de probabilité de la toxicité de chaque produit pris isolément sont soit obtenus du (des) clinicien(s) préparant l'essai, soit issus de données historiques, tandis qu'on ne définit que des a priori vagues pour les paramètres caractérisant les interactions entre produits. Une méthode d'obtention des a priori non informatifs est décrite. Le schéma est appliqué à un essai sur la gemcitabine et le cyclophosphamide, et on présente également une étude de simulation.

References

Berry, D. A. (1995). Decision analysis and Bayesian methods in clinical trials. In *Recent Advances in Clinical Trial Design and Analysis*, P. F. Thall (ed), 125–154. Norwell, Massachusetts: Kluwer.

Chaloner, K. and Larntz, K. (1989). Optimal Bayesian design applied to logistic regression experiments. *Journal of Statistical Planning and Inference* **21,** 191–208.

Gelfand, A. E. and Smith, A. F. M. (1990). Sampling based approaches to calculating marginal densities. *Journal of the American Statistical Association* **85,** 398–409.

Gilks, W. R., Clayton, D. G., Spiegelhalter, D. J., Best, N. G., McNeil, A. J., Sharples, L. D., and Kirby, A. J. (1993). Modelling complexity: Applications of Gibbs sampling in medicine. *Journal of the Royal Statistical Society Series B* **55,** 39–52.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association* **90,** 773–795.

Madigan, D. and Raftery, A. E. (1994). Model selection and accounting for model uncertainty in graphical models using Occam's window. *Journal of the American Statistical Association* **89,** 1535–1546.

Mantel, N. (1974). Therapeutic synergism. *Cancer Chemotherapy Reports* **4,** 147–149.

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.

Simon, R. and Korn, E. L. (1990). Selecting drug combinations based on total equivalent dose (dose intensity). *Journal of the National Cancer Institute* **82,** 1469–1476.

Simon, R. and Korn, E. L. (1991). Selecting combinations of chemotherapeutic drugs to maximize dose intensity. *Journal of Biopharmaceutical Statistics* **1,** 247–258.

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.

Zocchi, S. S. and Atkinson, A. C. (1999). Optimum experimental designs for multinomial logistics models. *Biometrics* **55,** 437–444.

Received July 2002. *Revised March* 2003.
Accepted March 2003.