Statistica Sinica **17**(2007), 531-547

BAYESIAN DOSE FINDING IN PHASE I CLINICAL TRIALS BASED ON A NEW STATISTICAL FRAMEWORK

Y. Ji, Y. Li and G. Yin

The University of Texas M. D. Anderson Cancer Center

Abstract: Phase I clinical trials aim to find the maximum tolerated dose of an experimental drug. We consider dose escalation, de-escalation, or staying at the current dose as three different stochastic moves over the lattice of a sequence of prespecified dose levels. Each move is chosen by minimizing an expected penalty that determines the dose level for treating the next cohort of patients. We develop a stopping rule under which the termination of the trial ensures that the posterior probability that the current dose is the maximum tolerated dose is larger than a prespecified value. Under a new class of priors, posterior estimates for the dose toxicity probabilities are obtained using the Markov chain Monte Carlo method. We demonstrate the new designs using a real phase I clinical trial.

Key words and phrases: Markov chain Monte Carlo, penalty, stopping rule, stochastic moves, toxicity.

1. Introduction

In phase I oncology clinical trials, various doses of a new drug are screened to search for the maximum tolerated dose (MTD), with a probability of toxicity that is closest to a prespecified value p_T. Patients in these trials are at high risk of death, so they consent to undergo therapies at dose levels that are possibly toxic in order to potentially prolong survival. Since relatively little is known about the appropriate dose level in this early phase of a study, a sequence of doses is screened in order to find the MTD.

Various statistical designs have been proposed to locate the MTD. O'Quigley, Pepe and Fisher (1990) proposed the continual reassessment method (CRM), which was based on a hypothetical function between the dose toxicity and dose level. The CRM was further investigated and refined by Goodman, Zahurak, and Piantadosi (1995) and Shen and O'Quigley (1996), among others. The use of a parametric function between dose toxicity and dose level reflects the need for statistical modeling when only a small number of patients are allowed in phase I clinical trials. More recently, a curve-free dose-finding method was proposed by Gasparini and Eisele (2000), in which the functional form between dose toxicity and dose level was relaxed while strong correlations were imposed among different dose toxicities. Cheung (2002) pointed out the rigidity of the curve-free method,

potentially due to its prior constructions. Whitehead and Brunier (1995) introduced a phase I design which incorporated the elements of decision theory. Their approach, like the CRM, also assumed a parametric function between the dose toxicity and the dose level. Durham, Flournoy and Rosenberger (1997) and Stylianou and Flournoy (2002) proposed a random walk rule also known as the biased coin design (BCD) for phase I clinical trials. Bayesian c- and D-optimal designs were investigated by Haines, Perevozskaya and Rosenberger (2003).

In this article, we propose a new method that provides a statistical framework for designing phase I trials. Our method does not assume a functional relationship between the dose toxicity and the dose level. We consider a stochastic walk over the lattice of the ordered experimental dose levels. That is, starting from the lowest dose level, a phase I trial proceeds through dose escalation, de-escalation, or staying at the current dose level. By walking over different doses based on the toxicity responses from sequentially assigned patients, the MTD is thus located.

Regarding these aspects of phase I trials, the proposed framework consists of three components: a set of penalty functions, stochastic moves, and a stopping rule. The penalty functions specify the penalties for taking a wrong move away from the MTD. For example, if the current dose is too toxic and the trial continues assigning patients to this dose level or even to the next higher dose level (which is more toxic), a penalty is then imposed for such a wrong move. The stochastic moves, including dose escalation, de-escalation, and staying at the current dose level, are determined by minimizing the associated penalties. Based on the toxicity responses observed from patients, a stopping rule is simultaneously applied to decide whether the trial can be terminated to select a tried dose as the MTD. Under a proper set of penalty values, the termination of the trial guarantees that the posterior probability that the selected dose is the MTD is greater than a prespecified threshold value.

The proposed dose-finding algorithm differs from the CRM in that the dose escalation, stay, or de-escalation is simply based on the toxicity of the current dose at which patients are treated. The CRM chooses the next dose by comparing the posterior means of toxicity probabilities for all the doses, including the tried and untried ones. The posterior means for the untried doses are not based on any observed data but on model and prior assumptions. If the assumptions are wrong, then the posterior estimates may be biased. The proposed algorithm focuses on the toxicity estimates for the current dose. These estimates are data driven and are less vulnerable to model mis-specifications. Under the new algorithm, if the current dose is close to the MTD, then the dose level will not be changed; if the current dose is too toxic, then the dose level will be de-escalated; otherwise, the dose level will be escalated.

In Section 2, we introduce the three key components of the proposed statistical framework, on which we base the proposal of our Bayesian dose-finding

design and the development of its theoretic properties in Section 3. We describe a new class of prior distributions in Section 4. In Section 5, we perform simulations based on a cancer phase I clinical trial to demonstrate the operating characteristics of the new Bayesian design. We provide some concluding remarks in Section 6, and outline the technical proofs in the Appendix.

2. Statistical Framework

We consider d dose levels of a certain cytotoxic drug in a phase I trial. Let p_i be the unknown probability of toxicity associated with the ith dose, $i = 1, \ldots, d$. The toxicity probability usually increases with the dose level, and we assume

$$p_1 < \cdots < p_d. \tag{1}$$

Suppose that the trial starts at dose i, and n_i ($n_i \geq 1$) patients are treated, of which x_i experience toxicity. Based on the observed values of x_i and n_i, one of three moves is taken to treat the next cohort of patients: de-escalate (D) to the previous lower dose ($i-1$); stay (S) at the same dose i; or escalate (E) to the next higher dose ($i+1$). According to the chosen move, the next cohort is treated at dose $j \in \{i-1, i, i+1\}$; the values of x_j and n_j are observed for the new cohort, after which an appropriate move is taken based on the cumulated data from both dose i and j. The trial thus continues until at least one dose is selected as the MTD.

2.1. The penalty functions

We propose a set of penalty functions for choosing a proper stochastic move at each step of the trial. For dose i, define the penalty functions

$$L(D, p_i) = \begin{cases} K_D, & \text{if } -\delta_1 \leq p_i - p_T \leq \delta_2; \\ 0, & \text{if } p_i - p_T > \delta_2; \\ N_D, & \text{if } p_i - p_T < -\delta_1; \end{cases}$$

$$L(S, p_i) = \begin{cases} 0, & \text{if } -\delta_1 \leq p_i - p_T \leq \delta_2; \\ M_S, & \text{if } p_i - p_T > \delta_2; \\ N_S, & \text{if } p_i - p_T < -\delta_1; \end{cases}$$

$$L(E, p_i) = \begin{cases} K_E, & \text{if } -\delta_1 \leq p_i - p_T \leq \delta_2; \\ M_E, & \text{if } p_i - p_T > \delta_2; \\ 0, & \text{if } p_i - p_T < -\delta_1. \end{cases}$$

In phase I trials, the toxicity probability of a candidate dose almost never equals p_T. Therefore, the interval $(p_T - \delta_1, p_T + \delta_2)$ allows the design to select the dose at which the toxicity probability is close to p_T. The two δ values can be determined by consultation with physicians and by proper tuning to achieve desirable sample size in computer simulations. For ease of description, we define

a target dose as one with a toxicity probability p_i that falls into the interval $[p_T - \delta_1, p_T + \delta_2]$. Note that whether a dose is a target dose depends on the values of the δ's. For small δ_1 and δ_2, there is usually only one target dose, which is considered to be an MTD candidate.

The six penalties K_D, K_E, M_S, M_E, N_S and N_D are positive real numbers. The value of K_D or N_D is the penalty for taking the move D (de-escalate) when the current dose is a target dose and the right move is S (stay), or when the current dose level is lower than the MTD and the right move is E (escalate), respectively. The values of M_S, N_S, K_E and M_E can be interpreted similarly. We assign zero penalty for taking the right moves.

2.2. The stochastic moves

Let $\mathcal{X} = \{(x_1, n_1), \ldots, (x_d, n_d)\}$ be the cumulated data in which n_i patients have been treated at dose i and x_i of them have experienced toxicities, for $i = 1, \ldots, d$. The information set corresponding to $\mathcal{X}$ is a σ-algebra, $\mathcal{F} = \sigma(\mathcal{X})$. Note that the x and n values in $\mathcal{X}$ increase as the trial continues to accrue new patients. Suppose that the prior for the vector $\boldsymbol{p} = (p_1, \ldots, p_d)'$ has a density $\pi(\boldsymbol{p})$. Define

$$R(D, p_i) = E\{L(D, p_i)|\mathcal{F}\}, \quad R(S, p_i) = E\{L(S, p_i)|\mathcal{F}\},$$
$$\text{and } R(E, p_i) = E\{L(E, p_i)|\mathcal{F}\}$$

to be the three posterior expected penalties corresponding to $\pi(\boldsymbol{p})$. Let

$$q_{Di} = \Pr(p_i - p_T > \delta_2|\mathcal{F}), \quad q_{Si} = \Pr(-\delta_1 \le p_i - p_T \le \delta_2|\mathcal{F}),$$
$$\text{and } q_{Ei} = \Pr(p_i - p_T < -\delta_1|\mathcal{F}).$$

Then

$$R(D, p_i) = K_D q_{Si} + N_D q_{Ei}; \quad R(S, p_i) = M_S q_{Di} + N_S q_{Ei};$$
$$\text{and } R(E, p_i) = K_E q_{Si} + M_E q_{Di}. \tag{2}$$

The stochastic move $\mathcal{B}_i$ at dose i is

$$\mathcal{B}_i = \arg \min_{m \in \{D,S,E\}} R(m, p_i), \tag{3}$$

which chooses the move in $\{D, S, E\}$ that has the smallest posterior expected penalty.

2.3. The stopping rule

Define a stopping region

$$\mathcal{A}_i = \left\{ \mathcal{X} : \frac{\min[R(D, p_i), R(E, p_i)]}{R(S, p_i)} \ge \xi \right\}, \quad i = 1, \ldots, d,$$

for each dose i, where ξ is a positive unknown constant. Then define a random variable $\mathcal{T} = \sum_{i=1}^{d} I_{(\mathcal{X} \in \mathcal{A}_i)}$, where $I_{()}$ is the indicator function. The stopping rule terminates the trial if $\mathcal{T} > 0$. Specifically, when $\mathcal{T} = 1$, we select dose i as the target dose for which $\mathcal{X} \in \mathcal{A}_i$. When $\mathcal{T} > 1$, there is more than one target dose and we can select the dose with the smallest posterior expected penalty $R(S, p_i)$. If $\mathcal{T} = 0$, the trial continues by taking the stochastic move $\mathcal{B}_i$. An alternative way of selecting a target dose when $\mathcal{T} > 1$ is by comparing the posterior probability of q_{Si} and selecting the dose with the largest q_{Si}. Under the configurations of the penalties described later in Section 3.3, these two approaches are equivalent.

According to (2), the criterion

$$\frac{\min\left[R(D, p_i), R(E, p_i)\right]}{R(S, p_i)} \geq \xi \tag{4}$$

is satisfied only when the posterior probability q_{Si} is much larger than q_{Di} and q_{Ei}, i.e., when there is a large probability that the current dose is the target dose.

3. Bayesian Design and Its Configurations

3.1. Bayesian design

Based on the penalty function, the stochastic move, and the stopping rule, we propose a two-step Bayesian design:

(1) Before treating the next cohort of patients, if $\mathcal{T} > 0$, terminate the trial. Specifically, when $\mathcal{T} = 1$, select dose i, for which $\mathcal{X} \in \mathcal{A}_i$, as the target dose. When $\mathcal{T} > 1$ there is more than one target dose, and choose the one with the smallest posterior expected penalty of S, $R(S, p_i)$.

(2) If $\mathcal{T} = 0$, take the stochastic move $\mathcal{B}_i$ for the current dose i at which patients are treated.

For safety reasons, one can add the restriction that any untried dose cannot be selected as the target dose.

To implement the Bayesian design, we need to specify the values of all the design parameters, including the penalties and the unknown parameters δ_1, δ_2 and ξ.

3.2. Specification of ξ

The following theorem provides a specification of parameter ξ.

Theorem 3.1. *For any given $\alpha \in (0, 1]$, let*

$$\xi = \frac{(1-\alpha)\{\max(K_D, K_E) - \max(N_D, M_E)\} + \max(N_D, M_E)}{\alpha \min(M_S, N_S)}. \tag{5}$$

Then

$$\mathcal{A}_i \subseteq \{\mathcal{X} : \ Pr(-\delta_1 \leq p_i - p_T \leq \delta_2 | \mathcal{F}) \geq 1 - \alpha\}. \tag{6}$$

According to (6)*, when the trial is terminated based on the* ξ *in* (5) *and dose* i *is selected based on the proposed design, the posterior probability that dose* i *is a target dose will be at least* $(1-\alpha)$*.*

3.3. Configuration of penalties

We impose the following condition for the penalties.

Condition 3.1. The penalties satisfy $M_S + N_S = K_D + N_D = K_E + M_E$.

This condition is based on the rationale that when the three posterior probabilities q_{Si}, q_{Di} and q_{Ei} are equal, the three posterior penalties $R(S, p_i)$, $R(D, p_i)$ and $R(E, p_i)$ should be equal.

To facilitate both the stochastic moves $\mathcal{B}_i$ and the stopping rule $\mathcal{T}$, the following theoretical results lead to two sets of penalty values, one for the stochastic move and the other for the stopping rule. In dose-finding trials, fast escalation and de-escalation lead to a fast location of the MTD, which is essential in maintaining a small sample size. Therefore, the penalties for the stochastic moves should be configured such that moving over different doses is as free as possible. In contrast, the stopping rule terminates the trial when the target dose is located. Thus, the penalties for the stopping rule should be configured such that the termination of the trial is as easy as possible once the target dose is found. Apparently, the purpose of the stochastic move and that of the stopping rule are complementary; the stochastic move tends to carry on the trial and the stopping rule is formulated to terminate the trial. Consequently, their corresponding penalties are very different.

Penalties for the stopping rule: Lemma 3.1, Theorems 3.2 and 3.3 provide the best penalty values for stopping the trial.

Lemma 3.1. *Under Condition* 3.1*, for any trial data* $\mathcal{X}$*, if dose* i *is selected as a target dose based on the penalty function in which* $K_D \neq K_E$ *or* $M_S \neq N_S$*, dose* i *must be selected based on the penalty function in which* $K_D = K_E$ *and* $M_S = N_S$*.*

Remark 3.1. Under Condition 3.1 and Lemma 3.1, the penalties are in the form of $M_S = N_S = K$, $K_D = K_E = K(1+\beta)$ and $N_D = M_E = K(1-\beta)$, for $K > 0$ and $\beta \in (-1, 1)$.

Without loss of generality, we assume $K = 1$ hereafter. The penalties thus are functions of β, which is directly related to the ease of termination of the trial. If we refer to the definition of the penalty and consider the case when $-\delta_1 < p_i - p_T < \delta_2$, we see that a larger β leads to larger penalties of D and E, thus making the stay S easier, which further leads to easier termination of the trial.

Theorem 3.2. *Let $\mathcal{L}(\beta)$ denote the set of penalty functions in which $M_S = N_S = 1$, $K_D = K_E = 1 + \beta$, and $N_D = M_E = 1 - \beta$, for $\beta \in (-1, 1)$. Then, given $\mathcal{X}$, if dose i is selected as a target dose under $\mathcal{L}(\beta)$ for $\beta \in (-1, 0]$, dose i must be selected under $\mathcal{L}(\beta)$ for $\beta \in (0, 1)$.*

This theorem states that the set of doses selected under $\mathcal{L}(\beta)$ ($\beta \leq 0$) are contained in the set of doses selected under $\mathcal{L}(\beta)(\beta > 0)$. Define the average sample size as the expected sample size of the trial where the expectation is taken with respect to the probability measure on the space $(\mathcal{X}, \mathcal{F})$. The following theorem reveals the relationship between the average sample size of the trial and the value of β.

Theorem 3.3. *The average sample size of the design under $\mathcal{L}(\beta)$ is a decreasing function in β, for $\beta \in (-1, 1)$.*

According to the above theorem, the optimal set of penalties for stopping the trial has the form $M_S = N_S = 1$, $K_D = K_E = 1 + \beta$ and $N_D = M_E = 1 - \beta$, with β taking a value less than, but as close to, one as possible. As indicated by simulations, a design with $\beta = 0.9$, i.e., with the penalty $\mathcal{L}(0.9)$, is close to being optimal in practice, and the reduction in the average sample size is negligible if β is larger.

Penalties for the stochastic move: The penalty function $\mathcal{L}(0.9)$ is not suitable for determining the three stochastic moves. It is easy to show that the posterior expected penalties $R(D, p_i)$ and $R(E, p_i)$ for taking the moves D and E are decreasing in β under $\mathcal{L}(\beta)$ for $\beta > 0$. Therefore, if $\mathcal{L}(0.9)$ is used to compute the stochastic move $\mathcal{B}_i$, it will result in a design that is reluctant to change from the current dose. Consequently, fast escalation and de-escalation would not be possible under this penalty.

Alternatively, we choose the penalty $\mathcal{L}(-0.9)$ for computing the stochastic move $\mathcal{B}_i$ which gives minimum penalty for taking the moves D and E. This penalty function will allow the design to escalate or de-escalate quickly and hence to find the target dose rapidly. In addition, it controls very well for overdose since it moves away from excessively toxic doses more easily than other penalties. However, the design under $\mathcal{L}(-0.9)$ will tend to assign fewer patients at the MTD, since it favors the moves D and E over the move S.

To implement the new Bayesian design, we recommend using $\mathcal{L}(0.9)$ in computing the stopping rule $\mathcal{T}$ and using $\mathcal{L}(-0.9)$ in computing the stochastic move $\mathcal{B}_i$.

4. Probability model

The proposed prior distributions are based on a class of transformations that

represent the toxicity probabilities p_i. Define

$$p_i = h(\epsilon_i) = \frac{\sum_{j=1}^{i} \exp(\epsilon_i)}{1 + \sum_{j=1}^{i} \exp(\epsilon_i)}, \ i = 1, \ldots, d.$$

We assume that ϵ_i are independent and follow $N(0, \sigma^2)$. Under this transformation, the condition that $p_i < p_{i+1}$ is guaranteed automatically. The likelihood function is a product of binomial densities $l(\boldsymbol{p}) \propto \prod_{i=1}^{d} p_i^{x_i}(1-p_i)^{n_i - x_i}$; and the joint posterior density of $\boldsymbol{\epsilon} = (\epsilon_1, \ldots, \epsilon_d)'$ is given by

$$f(\boldsymbol{\epsilon}) \propto \prod_{i=1}^{d} h(\epsilon_i)_i^x (1 - h(\epsilon_i))^{n_i - x_i} \phi(\epsilon_i; 0, \sigma^2), \tag{7}$$

where $\phi(\cdot; \mu, \sigma^2)$ is the density function of a normal distribution with mean μ and variance σ^2.

When the toxicity outcomes of the previous cohort are observed, posterior samples of the toxicity probabilities p_i are drawn from the full conditional distributions using the Gibbs sampler with the adaptive rejection Metropolis sampling algorithm (Gilks, Best and Tan (1995)). After 3,000 burn-in samples, we keep every fifth sample in the 5,000 samples from Markov chains. In our MCMC simulations, the autocorrelations of the final samples were negligible and the Markov chains mixed very well, implying fast convergence.

5. Example

We illustrate the proposed Bayesian design based on a clinical trial described in a study by Goodman, Zahurak, and Piantadosi (1995). For comparison, we also implement the continual reassessment method (CRM) and the biased coin design (BCD). The clinical trial was an open-label, non-comparative, multicenter dose-escalation study and the drug was an orally administered compound for treating patients with advanced cancer. The MTD was defined to be the highest dose level at which no more than $p_T = 25\%$ of the treated patients would exhibit dose-limiting toxicity. There were eight doses ($d = 8$) available at 50, 100, 200, 300, 400, 500, 650, and 800 mg/day. The first cohort of patients started at the lowest dose of 50 mg/day. In our simulations, the cohort size was three for the new Bayesian design and CRM, and was one for the BCD.

For the CRM design, we used the power model (Shen and O'Quigley (1996)) in which

$$p_i = \hat{p}_i^{\alpha}, \quad i = 1, \ldots, d,$$

where α has a normal prior with mean 0 and variance 2. In a sequential manner, the BCD design assigns the next patient to a future dose whenever the toxicity response is observed from the previous patient. When $p_T < 0.5$, the BCD steps

down a dose if toxicity is observed from the previous patient and randomizes with probability $p_T/(1-p_T)$ to the next higher dose and $\{1-p_T/(1-p_T)\}$ to the same dose if no toxicity is observed from the previous patient. When $p_T > .5$, the BCD design is similar. (See the work of Stylianou and Flournoy (2002) for details.) The trial is stopped when the prespecified maximum sample size is reached. We used the isotonic estimator with linear interpolation, as described by Stylianou and Flournoy (2002), to estimate the dose level of the MTD. The isotonic estimator is obtained via the pool adjacent violators algorithm (Robertson, Wright and Dykstra (1988)). Then, the dose closest to the estimated dose level of the MTD is selected as the recommended dose.

To implement the new Bayesian design, we took a vague prior for ϵ_i by assigning the prior variance $\sigma_i^2 = 20$ for $i = 1, \ldots, 8$. We took $\delta_1 = 0.1$, $\delta_2 = 0.15$, and $(1-\alpha) = 0.7$, which determines the value of ξ in the stopping rule. As suggested by a referee, we tried the following alternative stopping rule for the proposed method: the trial is terminated if the posterior probability q_{Si} of staying is greater than a cutoff value q^*, which is taken to be 0.7 in our simulation. Intuitively, this rule stops the trial and selects dose i if there is substantial evidence that p_i is in the neighborhood of p_T. For ease of exposition, we denote the proposed method using the stopping rule (4) as "Bayesian ξ" and the method with the alternative stopping rule as "Bayesian q_S".

We simulated 1,000 trials for 10 different scenarios. For each scenario, eight prespecified true dose toxicity probabilities were assigned to the corresponding doses, and toxicity outcomes were generated based on these probabilities. The maximum sample size was 30. For both the proposed Bayesian design and the CRM, the prior toxicity probabilities $\hat{p}_i = 0.05i$, $i = 1, \ldots, 8$. In practice, ethical concerns require that the trial be terminated early once excessive toxicity is observed for the first dose. For the Bayesian design and the CRM, if the posterior probability of the first dose's toxicity probability being larger than p_T is larger than .95, the trial is stopped and no dose is selected as the MTD. The BCD design in our simulation did not have an early stopping rule and always assigned all 30 patients to the trial doses.

Table 1 lists the operating characteristics of all four designs for the 10 different dose toxicity scenarios. For each scenario, the percentages for selecting the target doses and the average number of patients treated at each dose are listed. The standard deviations of the toxicity percentage and the sample size over the 1,000 simulations, if provided by the software, are given in parentheses in the last two columns. The correct selection percentages of the MTD are in bold font.

The operating characteristics of the four methods are comparable in most scenarios. In Scenario 1, dose 5 and 6 are the MTD. The three methods, "Bayesian ξ", "Bayesian q_S" and the CRM all performed very well, although the CRM has

Table 1. Simulation results comparing the new Bayesian designs, the CRM and the BCD.

	Recommendation percentage at dose level $100(1-\alpha)=70$, $\delta_1=0.1$, $\delta_2=0.15$									Toxicity percentage (SD	Average number of
Design	1	2	3	4	5	6	7	8	none	in percentage)	patients (SD)
Scenario 1	1	5	10	10	25	25	35	45			
Bayesian ξ	0	0.4	5.3	24.0	**37.4**	**23.6**	8.2	1.2	0	13.0 (3.7)	29.6 (1.3)
# pt	3.6	4.9	5.3	6.6	5.2	2.8	0.8	0.3			
Bayesian q_S	0	0.5	5.6	21.1	**42.4**	**21.9**	7.5	1.0	0	13.1 (4.1)	29.6 (1.2)
# pt	3.7	4.6	5.1	7.0	5.4	2.8	0.7	0.2			
CRM	0	0	4	18	**35**	**26**	13	4	0	15.9	30
# pt	3.1	3.8	4.3	6.0	6.3	4.0	1.9	0.5			
BCD	0	2.2	5.6	23.4	**31.1**	**20.3**	13.6	3.8	0	14.3 (5.3)	30
# pt	3.6	4.7	5.1	6.1	4.8	3.2	1.8	0.7			
Scenario 2	5	10	15	25	50	55	70	80			
Bayesian ξ	0.5	4.7	21.8	**60.1**	11.9	1.0	0	0	0	17.8 (4.9)	29.3 (2.0)
# pt	4.9	5.9	7.5	8.0	2.7	0.3	0	0			
Bayesian q_S	0.5	7.2	28.2	**54.7**	9.1	0.3	0	0	0	17.9 (4.5)	28.8 (2.5)
# pt	4.9	6.4	7.7	7.2	2.3	0.3	0	0			
CRM	0	6	29	**53**	11	0	0	0	0	19.7	30
# pt	4.0	5.9	7.5	9.0	3.0	0.7	0.1	0			
BCD	0.3	5.6	27.4	**54.2**	10.6	1.7	0	0	0	18.7 (5.3)	30
# pt	5.0	6.1	7.9	7.0	3.0	0.8	0.2	0			
Scenario 3	5	15	25	35	45	55	65	75			
Bayesian ξ	1.3	23.7	**46.5**	24.7	3.6	0.2	0	0	0	19.9 (4.9)	28.4 (1.2)
# pt	5.9	9.2	8.1	4.1	1.0	0.2	0	0			
Bayesian q_S	1.1	27.4	**45.4**	21.7	4.2	0.2	0	0	0	20.0 (4.8)	28.1 (3.0)
# pt	6.0	9.2	7.8	4.0	1.0	0.2	0	0			
CRM	0	25	**50**	21	4	0	0	0	0	21.8	30
# pt	4.2	9.5	9.5	5.1	1.4	0.3	0	0			
BCD	0.9	31.6	**38.8**	22.8	5.2	0.4	0	0	0	20.8 (5.5)	30
# pt	7.1	8.9	7.4	4.3	1.8	0.5	0.1	0			
Scenario 4	25	50	55	60	65	70	75	80			
Bayesian ξ	**71.2**	13.2	0.9	0.3	0	0	0	0	14.4	32.0 (7.0)	22.3 (7.3)
# pt	15.6	5.8	0.8	0.1	0	0	0	0			
Bayesian q_S	**74.7**	12.2	0.4	0	0	0	0	0	12.7	31.8 (13.7)	22.1 (7.1)
# pt	16.1	5.5	0.5	0	0	0	0	0			
CRM	**85**	9	0	0	0	0	0	0	7	30.6	29
# pt	22.4	5.9	0.5	0.1	0	0	0	0			
BCD*	**90.9**	8.8	0.2	0	0	0	0	0	0	35.3 (6.5)	30
# pt	18.8	8.3	2.3	0.6	0.1	0	0	0			
Scenario 5	1	1	1	25	60	70	80	90			
Bayesian ξ	0	0	2.3	**88.1**	9.5	0.1	0	0	0	18.4 (5.2)	29.8 (1.0)
# pt	3.2	3.2	7.6	11.8	3.8	0.2	0	0			
Bayesian q_S	0	0	2.7	**87.6**	9.7	0.0	0	0	0	18.2 (4.9)	29.7 (1.1)
# pt	3.2	3.2	7.7	11.7	3.7	0.2	0	0			
CRM	0	0	9	**80**	11	0	0	0	0	22.6	30
# pt	3.1	3.2	5.3	12.8	5.0	0.6	0	0			
BCD	0	0	8	**84.8**	6.7	0.4	0.1	0	0	17.7 (5.8)	30
# pt	0.0	0.3	10.0	13.0	5.5	1.1	0.1	0			

Table 1 (continued)

	Recommendation percentage at dose level $100(1-\alpha)=70$, $\delta_1=0.1$, $\delta_2=0.15$									Toxicity percentage (SD	Average number of
Design	1	2	3	4	5	6	7	8	none	in percentage)	patients (SD)
Scenario 6	1	1	1	20	20	20	60	70	none		
Bayesian ξ	0	0	1.8	**32.5**	**29.1**	**31.7**	4.9	0	0	13.4 (4.2)	29.7 (1.1)
# pt	3.3	3.3	6.3	8.0	4.6	3.2	0.9	0.1			
Bayesian q_S	0	0	1.6	**32.2**	**26.7**	**35.6**	3.9	0	0	13.4 (4.0)	29.7 (1.1)
# pt	3.2	3.2	6.5	7.9	4.7	3.2	0.8	0.2			
CRM	0	0	1	**18**	**27**	**44**	11	0	0	16.9	30
# pt	3.1	3.3	3.6	5.8	6.2	5.6	2.2	0.2			
BCD	0	0	3.3	**30.2**	**21.4**	**38.7**	6.0	0.4	0	14.7 (5.4)	30
# pt	3.1	3.2	6.5	6.2	4.9	4.2	1.6	0.2			
Scenario 7	50	70	80	87	88	89	90	90			
Bayesian ξ	11.5	0	0	0	0	0	0	0	**88.2**	51.7 (10.9)	11.5 (7.3)
# pt	10.7	0.7	0	0	0	0	0	0			
Bayesian q_S	12.0	0	0	0	0	0	0	0	**88.0**	51.6 (19.6)	11.8 (7.1)
# pt	10.9	0.8	0	0	0	0	0	0			
CRM	15	0	0	0	0	0	0	0	**85**	51.6	15.7
# pt	15	0.7	0	0	0	0	0	0			
BCD*	98.8	1.2	0	0	0	0	0	0	**0**	54.0 (7.9)	30
# pt	24.0	5.4	0.6	0	0	0	0	0			
Scenario 8	1	2	3	4	5	5	6	6			
Bayesian ξ	0	0	0.1	0.8	3.3	12.4	13.0	**70.3**	0	4.0 (3.0)	30.0 (0.2)
# pt	3.3	3.5	3.7	3.9	3.9	3.9	3.1	4.8			
Bayesian q_S	0	0	0.2	1.4	2.8	11.6	12.2	**71.8**	0	4.0 (3.5)	30.0 (0.3)
# pt	3.3	3.4	3.8	3.9	3.8	3.8	3.2	4.9			
CRM	0	0	0	2	4	7	8	**79**	0	4.2	30
# pt	3.1	3.3	3.3	3.8	3.4	3.7	3.1	6.2			
BCD	0	0	0.9	1.8	4.7	8.9	13.4	**70.3**	0	4.2 (3.6)	30
# pt	3.4	3.3	3.3	3.7	3.5	3.3	3.2	6.2			
Scenario 9	1	2	3	5	10	20	35	45			
Bayesian ξ	0	0.1	0.1	2.4	21.4	**49.5**	24	2.5	0	11.1 (3.5)	29.9 (0.5)
# pt	3.3	3.5	3.8	4.7	5.7	6.0	2.2	0.7			
Bayesian q_S	0	0	0	2.8	19.8	**50.0**	24.2	3.2	0	11.1 (3.7)	29.9 (0.6)
# pt	3.3	3.5	3.9	4.8	5.6	5.9	2.2	0.7			
CRM	0	0	1	3	14	**40**	30	13	0	14.1	30
# pt	3.1	3.4	3.3	3.9	4.7	5.8	3.8	2.0			
BCD	0	0	1.1	4.1	22.2	**39.2**	25.9	7.3	0	12.0 (5.0)	30
# pt	3.3	3.6	3.7	4.6	5.7	5.1	3.0	1.2			
Scenario 10	5	10	50	60	70	75	78	80			
Bayesian ξ	1.1	**68.8**	29.0	1	0	0	0	0	0.1	21.7 (5.6)	29.1 (3.2)
# pt	4.8	15.4	8.0	0.9	0.1	0	0	0			
Bayesian q_S	0.9	**54.8**	43.7	0.5	0	0	0	0	0.1	21.6 (5.7)	29.2 (2.3)
# pt	4.9	15.5	7.9	0.8	0	0	0	0			
CRM	0	**59**	40	0	0	0	0	0	0	25.6	30
# pt	3.9	14.3	10.2	1.5	0.1	0	0	0			
BCD	0.5	**69.2**	29.2	0.8	0.2	0	0	0	0	22.5 (5.3)	30
# pt	6.9	13.3	7.5	1.9	0.3	0	0	0			

*BCD does not stop the trial early and therefore is not comparable under Scenarios 4 and 7.

slightly higher toxicity percentage than the other two Bayesian methods. In Scenario 2, dose 4 is the MTD and "Bayesian ξ" has the highest selection percentage for the MTD. The performances of the other three methods are very similar. The CRM is the best method for Scenario 3, in which the toxicity increases at a constant rate and dose 3 is the MTD. In Scenario 4, the BCD does not stop the trial early and is thus not comparable to the other three methods. The CRM is the best method for this scenario. In Scenario 5, dose 4 is the MTD with doses higher being much more toxic and doses lower being much less toxic. "Bayesian ξ" and "Bayesian q_S" performed very well as they not only selected the MTD with higher percentages, they treated fewer patients at highly toxic doses. Scenario 6 is similar to 5 except that the middle three doses, dose 4, 5 and 6, are the MTD. The two Bayesian methods are better for this scenario. All the doses are too toxic in Scenario 7, and the BCD is not comparable here since it does not terminate the trial early. The other three methods have similar results, although the two Bayesian methods, on average, treated about four fewer patients in total than the CRM. In Scenario 8, all the doses are lower than the MTD and the CRM performs a little better than the other three methods. The two Bayesian methods performed better than the CRM and the BCD in Scenario 9. Finally, in Scenario 10, where dose 2 is very nontoxic and dose 3 is very toxic, "Bayesian ξ" and the BCD are the clear winners. Overall, the new Bayesian designs with both stopping rules have nice properties and are comparable to the CRM and the BCD.

Table 2 presents the simulation results of the proposed "Bayesian ξ" method under three representative prior distributions for Scenario 1. Specifically, we varied the values of σ^2 to be 2, 20 and 200. When $\sigma^2 = 2$, the priors of ϵ_i are informative and the induced priors of p_i are unimodal and rigid, which leads to undesirable simulation results. When the value of σ^2 is large (e.g., 20 or 200), the priors of ϵ_i are vague and the induced priors of p_i are U-shaped, assigning a small probability mass on most values between 0 and 1. Based on the work by Zhu and Lu (2004), the U-shaped prior is noninformative to the Bernoulli family. However, our model is more complicated since we are modeling the priors of d toxicity probabilities that are order restricted. Therefore, it is not always the case that a larger prior variance of p_i will lead to better results. In Table 2, the operating characteristics under $\sigma^2 = 20$ and 200 are both reasonable, partly because they correspond to U-shaped priors for toxicity probabilities. The design under the $\sigma^2 = 20$ prior has a better overdose control, i.e., it selects the excessive toxic doses with smaller percentages. This is mainly because the $\sigma^2 = 200$ prior assigns more probability mass to values close to 0 than does the $\sigma^2 = 20$ prior. Therefore, the $\sigma^2 = 200$ prior tends to underestimate the toxicity probabilities, especially when most patients do not experience toxicity. Consequently, the dose

may be considered safer than it actually is, and doses with excessive toxicity are more likely to be selected. We see this phenomenon in Table 2 where dose 7 is selected with a larger percentage under the $\sigma^2 = 200$ prior than under the $\sigma^2 = 20$ prior.

Table 2. Sensitivity analysis for the new Bayesian design using different values of σ^2.

	Recommendation percentage at dose level $100(1-\alpha) = 40$, $\delta_1 = 0.1$, $\delta_2 = 0.15$									Toxicity percentage (SD in percentage)	Average number of patients (SD)
	1	2	3	4	5	6	7	8	none		
Scenario 1	1	5	10	10	25	25	35	45			
$\sigma^2 = 2$	0.8	48.0	51.2	0	**0**	**0**	0	0	0	4.5 (1.3)	21.0 (4.4)
# pt	7.2	9.7	4.1	0	0	0	0	0			
$\sigma^2 = 20$	0	0.4	5.3	24.0	**37.3**	**23.6**	8.2	1.2	0	13.0 (3.7)	29.6 (1.3)
# pt	3.6	4.9	5.3	6.6	5.2	2.8	0.8	0.3			
$\sigma^2 = 200$	0.2	0.2	2.9	13.1	**24.1**	**27.5**	23.2	8	0	17.4 (4.9)	29.5 (1.2)
# pt	3.3	3.5	3.8	5.2	5.3	4.3	2.6	1.5			

6. Discussion

We have proposed a new statistical framework and a Bayesian design for phase I clinical trials and have demonstrated its operating characteristics based on a real trial. Implementation of the design is straightforward. According to the theoretical results that specify most design parameter values, we only need to tune the values of δ_1 and δ_2 to achieve a desirable sample size. We also find that the operating characteristics of the Bayesian design are not sensitive to the values of $\hat{p}_i$, the prior guess of the toxicity probabilities (results not shown).

Babb, Rogatko and Zacks (1998) and Tighiouart, Rogatko and Babb (2005) used a different set of penalties that are proportional to the distance between the dose toxicity probabilities and p_T. This type of penalty does not fit into our proposed dose-finding strategy, in which the dose-assignment rule does not involve estimation of toxicity probabilities for untried doses. They also imposed larger penalties for moving toward excessive doses to achieve better overdose control. Our new Bayesian design appears to have good control for overdosing, as demonstrated by the simulation results, due to the configuration of the penalties.

The proposed Bayesian design only deals with phase I clinical trials considering toxicity. We are currently extending the method to accommodate both efficacy and toxicity simultaneously. The extension involves deriving the bivariate penalty function in the two-dimensional space of toxicity and efficacy. Current research on combining toxicity and efficacy in phase I/II clinical trials includes the work by Thall and Cook (2004) and Thall and Russell (1998), among others.

Acknowledgement

We thank Mario Stylianou for providing computer code for implementing the biased-coin design. We thank an associate editor and two anonymous referees for their very helpful comments that greatly improved the manuscript.

Appendix

Proof of Theorem 3.1. First, note that $R(D,p_i) \le K_D q_{Si} + N_D(1-q_{Si})$ and $R(E,p_i) \le K_E q_{Si} + M_E(1-q_{Si})$, and therefore

$$\begin{aligned}&\min\left[R(D,p_i),R(E,p_i)\right]\\ &\quad\le \max(K_D,K_E)q_{Si} + \max(N_D,M_E)(1-q_{Si})\\ &\quad= \left[\max(K_D,K_E) - \max(N_D,M_E)\right]q_{Si} + \max(N_D,M_E).\end{aligned}$$

Also, $R(S,p_i) \ge \min(M_S,N_S)(1-q_{Si})$. Define

$$f(q_{Si}) = \frac{\left[\max(K_D,K_E) - \max(N_D,M_E)\right]q_{Si} + \max(N_D,M_E)}{\min(M_S,N_S)(1-q_{Si})}.$$

By the above inequalities, it immediately follows that

$$\frac{\min\left[R(D,p_i),R(E,p_i)\right]}{R(S,p_i)} \le f(q_{Si}).$$

Condition (5) in this theorem leads to $\xi = f(1-\alpha)$. Since $f(q_{Si})$ is increasing in q_{Si}, if

$$\frac{\min\left[R(D,p_i),R(E,p_i)\right]}{R(S,p_i)} \ge \xi,$$

then $f(q_{Si}) \ge \xi = f(1-\alpha)$. Therefore, $q_{Si} \ge 1-\alpha$. That is, if dose i is accepted based on the stopping region $\mathcal{A}_i$, it is accepted based on $\{\mathcal{X}:\ \Pr(-\delta_1 \le p_i - p_T \le \delta_2|\mathcal{F}) \ge 1-\alpha\} = \{\mathcal{X}:\ q_{Si} \ge 1-\alpha\}$.

Proof of Lemma 3.1. From the formula for computing ξ and the Bayesian design, it follows immediately that when $M_S \ne N_S$, the value of ξ is larger than that when $M_S = N_S$. Therefore, with all other penalties fixed, if the trial is terminated and dose i is selected given data $\mathcal{X}$ with $M_S \ne N_S$, the trial will always be terminated and dose i will be selected given the same data $\mathcal{X}$, but with $M_S = N_S$.

Without loss of generality, let $M_S = N_S = 1$. Under the penalties in which $K_D = K_E$, and according to Condition 3.1, let $K_D = K_E = 1+\beta$ and $N_D = M_E = 1-\beta$ for some $\beta \in (-1,1)$. Under this penalty function, ξ can be expressed in the form $\xi_1 = 1/\alpha + \beta(1-2\alpha)/\alpha$.

If $K_D \neq K_E$, without loss of generality assume that $K_D > K_E$. Under Condition 3.1, $N_D < M_E$. Then, for $\gamma \in (0,1)$, K_E, M_E, K_D and N_D can be expressed as $M_E = 1-\beta$, $K_E = 1+\beta$, $N_D = \gamma(1-\beta)$, and $K_D = 2-\gamma(1-\beta)$, for $\beta \in (-1,1)$, which guarantees $K_D > K_E$. Therefore, ξ can be written as

$$\xi_2 = \frac{2(1-\alpha) - (1+\gamma)(1-\alpha)(1-\beta) + 1 - \beta}{\alpha}.$$

It can be easily shown that ξ_2 converges monotonically down to ξ_1 as $\gamma \to 1$, by examining the first derivative of ξ_2. So $\xi_2 \downarrow \xi_1$ as $\gamma \uparrow 1$.

Therefore, for any trial data $\mathcal{X}$, if dose i is selected as the target dose under the penalties with $K_D \neq K_E$, the dose must be selected under the penalties with $K_D = K_E$.

Proof of Theorem 3.2. Without loss of generality, suppose $R(D,p_i) \leq R(E,p_i)$. Under the penalty function $\mathcal{L}(\beta)$, the stopping criterion (4) for dose i is

$$\frac{R(D,p_i)}{R(S,p_i)} \geq \frac{1}{\alpha} + \beta\frac{1-2\alpha}{\alpha},$$

after plugging the penalties into ξ. Next, plugging the penalties $\mathcal{L}(\beta)$ into $R(D,p_i)$ and $R(E,p_i)$ with their explicit expressions given at (2), we have

$$\frac{q_{Si}+q_{Ei}}{q_{Di}+q_{Ei}} + \beta\frac{q_{Si}-q_{Ei}}{q_{Di}+q_{Ei}} \geq \frac{1}{\alpha} + \beta\frac{1-2\alpha}{\alpha}.$$

Since dose i is selected under the penalty function $\mathcal{L}(\beta)$ for $\beta \leq 0$, we need only show that

$$\frac{q_{Si}-q_{Ei}}{q_{Di}+q_{Ei}} \geq \frac{1-2\alpha}{\alpha},$$

which is equivalent to $q_{Di} \geq (1-\alpha-q_{Si})/\alpha$. This inequality is true since $q_{Si} \geq 1-\alpha$ and $q_{Di} \geq 0$. Therefore, dose i is also selected as the target dose under $\mathcal{L}(\beta)$.

Proof of Theorem 3.3. Based on the proof of Theorem 3.2, given any trial data $\mathcal{X}$ and under penalty function $\mathcal{L}(\beta)$, if dose i is selected as the target dose for some β, it must be true that

$$\frac{\min[R(D,p_i), R(E,p_i)]}{R(S,p_i)} \geq \frac{1}{\alpha} + \beta\frac{1-2\alpha}{\alpha}.$$

Without loss of generality, suppose that $R(D,p_i) < R(E,p_i)$. The inequality above becomes

$$\frac{q_{Si}+q_{Ei}}{q_{Di}+q_{Ei}} - \frac{1}{\alpha} + \beta\left(\frac{q_{Si}-q_{Ei}}{q_{Di}+q_{Ei}} - \frac{1-2\alpha}{\alpha}\right) \geq 0. \tag{A.1}$$

According to Theorem 3.1, $q_{Si} \geq 1-\alpha$ for dose i since it is selected as the target dose. So $q_{Di} > (1-\alpha-q_{Si})/\alpha$, under which it must be true that $(q_{Si}-q_{Ei})/(q_{Di}+q_{Ei}) - (1-2\alpha)/\alpha > 0$. Hence, the third term in the left hand side of (A.1), $\beta((q_{Si}-q_{Ei})/(q_{Di}+q_{Ei}) - (1-2\alpha)/\alpha)$, is an increasing function in β.

Therefore, for any $\beta^* \geq \beta$, $(q_{Si}+q_{Ei})/(q_{Di}+q_{Ei}) - 1/\alpha + \beta^*((q_{Si}-q_{Ei})/(q_{Di}+q_{Ei}) - (1-2\alpha)/\alpha) \geq 0$. That is, for any $\beta^* \geq \beta$ and any trial data $\mathcal{X}$, the dose i is accepted under the penalty function $\mathcal{L}(\beta^*)$ if it is accepted under the loss $\mathcal{L}(\beta)$. Therefore, the average sample size under the penalty function $\mathcal{L}(\beta^*)$ is less than or equal to the average sample size under the penalty function $\mathcal{L}(\beta)$, completing the proof.

References

Babb, J., Rogatko, A. and Zacks, S. (1998). Cancer phase I clinical trials: efficient dose escalation with over dose control. *Statist. Medicine* **17**, 1103-1120.

Cheung, Y. (2002). On the use of nonparametric curves in phase I trials with low toxicity tolerance. *Biometrics* **58**, 237-240.

Durham, S., Flournoy, N. and Rosenberger, W. (1997). A random walk rule for phase I clinical trials. *Biometrics* **53**, 745-760.

Gasparini, M. and Eisele, J. (2000). A curve-free method for phase I clinical trials. *Biometrics* **56**, 609-615.

Gilks, W., Best, N. and Tan, K. (1995). Adaptive rejection Metropolis sampling within Gibbs sampling. *Appl. Statist.* **44**, 455-472.

Goodman, S., Zahurak, M. and Piantadosi, S. (1995). Some practical improvements in the continual reassessment method for phase I studies. *Statist. Medicine* **14**, 1149-1161.

Haines, L., Perevozskaya, I. and Rosenberger, W. (2003). Bayesian optimal designs for phase I clinical trials. *Biometrics* **59**, 591-600.

O'Quigley, J., Pepe, M. and Fisher, L. (1990). Continual reassessment method: A practical design for phase 1 clinical trials in cancer. *Biometrics* **46**, 33-48.

Robertson, T., Wright, F. and Dykstra, R. (1988). *Order Restricted Statistical Inference*. Wiley. New York.

Shen, L. and O'Quigley, J. (1996). Consistency of continual reassessment method under model misspecification. *Biometrika* **83**, 395-405.

Stylianou, M. and Flournoy, N. (2002). Dose finding using the biased coin up-and-down design and isotonic regression. *Biometrics* **58**, 171-177.

Thall, P. and Cook, J. (2004). Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* **60**, 684-693.

Thall, P. and Russell, K. (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.

Tighiouart, M., Rogatko, A. and Babb, J. (2005). Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Statist. Medicine* **24**, 2183-2196.

Whitehead, J. and Brunier, H. (1995). Bayesian decision procedures for dose determining experiments. *Statist. Medicine* **14**, 885-893.

Zhu, M. and Lu, A. (2004). The counter-intuitive non-informative prior for the Bernoulli family. *J. Statist. Education* **12**, 1-10.

Department of Bioinformatics and Computational Biology, The University of Texas M. D. Anderson Cancer Center, Houston, TX 77030, U.S.A.

E-mail: yuanji@mdanderson.org

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

E-mail: ysli@mdanderon.org

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

E-mail: gsyin@mdanderson.org

(Received March 2005; accepted February 2006)