Incorporating Individual and Collective Ethics into

Phase I Cancer Trial Designs

Jay Bartroff*

Department of Mathematics, University of Southern California,

3620 South Vermont Avenue, KAP 108, Los Angeles, CA 90089, U.S.A.

**email:* bartroff@usc.edu

and

Tze Leung Lai**

Department of Statistics, Sequoia Hall, Stanford University, Stanford, CA 94305, U.S.A.

***email:* lait@stat.stanford.edu

SUMMARY: A general framework is proposed for Bayesian model-based designs of Phase I cancer trials, in which a general criterion for coherence (Cheung, 2005) of a design is also developed. This framework can incorporate both "individual" and "collective" ethics into the design of the trial. We propose a new design which minimizes a risk function composed of two terms, with one representing the individual risk of the current dose and the other representing the collective risk. The performance of this design, which is measured in terms of the accuracy of the estimated target dose at the end of the trial, the toxicity and overdose rates, and certain loss functions reflecting the individual and collective ethics, is studied and compared with existing Bayesian model-based designs and is shown to have better performance than existing designs.

KEY WORDS: Cancer trials; Coherence; Dose-finding; Logistic regression; Markov decision problem; Phase I.

1. Introduction

A Phase I trial for a new treatment is generally intended to determine a dose to use in subsequent Phase II and III testing. Phase I cancer trials have the additional complexity that the treatment in question is usually a cytotoxic agent and the efficacy usually increases with dose, and therefore it is widely accepted that some degree of toxicity must be tolerated to experience any substantial therapeutic effects. Hence, an acceptable proportion p of patients experiencing *dose limiting toxicities* (DLTs) is generally agreed on before the trial, which depends on the type and severity of the DLT; the dose resulting in this proportion is thus referred to as the *maximum tolerated dose* (MTD). In addition to the explicitly stated objective of determining the MTD, a Phase I cancer trial also has the implicit goal of safe treatment of the patients in the trial. However, the aims of treating patients in the trial and generating an efficient design to estimate the MTD for future patients often run counter to each other. Commonly used designs in Phase I cancer trials implicitly place their focus on the safety of the patients in the trial, beginning from a conservatively low starting dose and escalating cautiously. Escalation is further slowed by the assignment of the same dose to groups of consecutive patients, as in the widely used 3-plus-3 design, which is convenient to administer and shortens trial duration by simultaneously following patients in groups of 3. Von Hoff and Turner (1991) have documented that the overall response rates in these Phase I trials are low, and substantial numbers of patients are treated at doses that are retrospectively found to be non-therapeutic. Moreover, as pointed out by O'Quigley, Pepe and Fisher (1990), these designs are very inefficient for estimating the MTD, which is implied by the 3-plus-3 design to correspond to the case $p = 1/3$. They proposed a Bayesian model-based design, called the "continual reassessment method" (CRM), to choose the dose levels sequentially, making use of all past data at each stage.

More than ninety new Phase I methods were published between 1991 and 2006 (Rogatko, Schoeneck, Jona

2007), and there have been several reviews of the new methods (e.g., Rosenberger and Haines, 2002). In this paper we focus on Bayesian model-based designs and Section 2 describes a general framework to develop and analyze them. As shown in Sections 2 and 3, this framework allows one to incorporate the competing aims of a Phase I cancer trial by choosing the loss function accordingly. It also enables one to derive certain desirable properties of the design, such as coherence (Cheung, 2005), from the loss function, or to enforce them by using simple reformulations in this framework. Section 4 provides implementation details and gives a simulation study comparing Bayesian designs that correspond to different loss functions in the setting of a colon cancer trial considered by Babb, Rogatko and Zacks (1998).

2. Posterior Distributions, Loss Functions and Sequential Dose Determination

A commonly used model-based approach to Phase I cancer clinical trial design assumes the usual logistic regression model for the probability $F_\theta(x)$ of DLT at dose level x:

$$F_\theta(x) = 1/(1 + e^{-(\alpha + \beta x)}), \tag{1}$$

in which $\beta > 0$ and $\theta = (\alpha, \beta)$ is unknown and to be estimated from the observed pairs (x_i, y_i), where $y_i = 1$ if the ith subject, treated at dose x_i, experiences DLT and $y_i = 0$ otherwise. The frequentist approach to inference on θ uses the likelihood function and estimates θ by maximum likelihood, while the Bayesian approach assumes a prior distribution of θ and uses the posterior distribution for inference on θ.

Denote the MTD by $\eta = F_\theta^{-1}(p)$ and the posterior distribution of θ based on $(x_1, y_1), \ldots, (x_k, y_k)$ by Π_k, and let Π_0 denote the prior distribution. The Bayes estimate of η with respect to squared error loss is the posterior mean $E_{\Pi_k}(\eta)$, and the CRM proposed by O'Quigley, Pepe and Fisher (1990) uses this posterior mean to set the dose for the next patient, i.e., $x_{k+1} = E_{\Pi_k}(\eta)$. Instead of the posterior mean, Babb, Rogatko and Zacks (1998) proposed to set x_{k+1} equal to the ω-quantile of the posterior distribution, where $0 < \omega < 1/2$ is chosen to be slightly less

than p in their examples. This design is called "escalation with overdose control" (EWOC) and ω is called the "feasibility bound." A sequence of doses x_n is called "Bayesian feasible" at level $1 - \omega$ if $P_{\Pi_{n-1}}(\eta \geqslant x_n) \geqslant 1 - \omega$ for all $n \geqslant 1$, and the EWOC doses are optimal among Bayesian-feasible ones; see Zacks, Rogatko and Babb (1998).

Note that the dose for the nth patient in CRM or EWOC depends only on the posterior distribution Π_{n-1}, i.e., x_n is a functional $f(\Pi_{n-1})$ of Π_{n-1}. This functional defines $\{\Pi_k : k \geqslant 0\}$ as a Markov chain whose states are distributions on the parameter space Θ and whose state transitions are given by the following.

Bayesian updating scheme: Given current state Π (which is a prior distribution of θ), let $x = f(\Pi)$ and generate first θ from Π and then $y \sim \text{Bern}(F_\theta(x))$. The new state is the posterior distribution of θ given (x, y).

The functional $x = f(\Pi)$ for CRM is $E_\Pi(\eta)$, which minimizes the expected squared error loss $E_\Pi[(\eta - x)^2]$. As pointed out by Babb, Rogatko and Zacks (1998), the symmetric nature of the squared error loss may not be appropriate for modeling the toxic response to a cancer treatment. Instead of squared error loss, EWOC with feasibility bound ω uses the functional $x = x(\Pi)$ that minimizes the asymmetric loss function $E_\Pi[\ell(\eta, x)]$, where

$$\ell(\eta, x) = \begin{cases} \omega(\eta - x), & \text{if } x \leqslant \eta \\ (1 - \omega)(x - \eta), & \text{if } x \geqslant \eta. \end{cases} \tag{2}$$

More generally, we can consider other loss functions $\ell(\theta, x)$ and define $x(\Pi)$ that attains $\min_x E_\Pi[\ell(\theta, x)]$. In particular, the following example gives a response-based version of EWOC.

Example 1: Inverted overdose control. The EWOC loss function (2) penalizes an overdose $x > \eta$ by the amount $(1 - \omega)(x - \eta)$, and an under-dose $x < \eta$ by the amount $\omega(x - \eta)$. However, a dose x deemed "too large" on this scale may actually correspond to a probability of DLT not much larger than the target rate p depending on the dose-response curve, making x a relatively desirable dose. Likewise, a small value of $|x - \eta|$ may correspond to a large

discrepancy between the actual DLT probability $F_\theta(x)$ and p. Hardwick and Stout (2001) suggest to measure the excess/deficit of the DLT rate on the "probability scale." Taking $0 < \gamma < 1/2$, this leads to the "inverted" loss function

$$\ell(\theta, x) = \begin{cases} \gamma(p - F_\theta(x)), & \text{if } x \leqslant \eta \\ (1-\gamma)(F_\theta(x) - p), & \text{if } x \geqslant \eta. \end{cases} \tag{3}$$

Remark: The loss when the dose falls below η, measured by the difference $\omega(\eta - x)$ in (2) and $\gamma(p - F_\theta(x))$ in (3), should ideally be measured by the difference in response rates at η and x, respectively, when efficacy data, taking the value 1 if the patient responds to the treatment and 0 otherwise, are also available besides the toxicity data. Note, however, that this involves bivariate efficacy-toxicity data. While many existing designs solely consider toxicity outcomes and the MTD, designs that incorporate efficacy responses as well have been proposed by a number of authors, including Li, Durham and Flournoy (1995), Hardwick and Stout (2001), Kpamegan and Flournoy (2001), Thall and Cook (2004), Dragalin and Fedorov (2006), Dragalin, Fedorov (2008), and Pronzato (2010). When efficacy responses are available, the *minimum effective dose* (MED) is of interest, i.e., the lowest dose at which some desired proportion of positive efficacy responses is attained. When both efficacy and toxicity data are available, the *optimal safe dose*, which is the dose between the MED and the MTD maximizing the probability of simultaneous efficacy and non-toxicity, is of interest. Since this paper focuses on univariate toxicity data, we consider elsewhere better alternatives to (3) for $x \leqslant \eta$ that also require efficacy data.

Noting that the explicitly stated objective of a Phase I cancer trial is to estimate the MTD, Whitehead and Brunier (1995) considered Bayesian sequential designs that are optimal, in some sense, for this estimation problem. Haines, Perevozskaya and Rosenberger (2003) made use of the theory of optimal design of experiments (Fedorov, 1972; Atkinson and Donev, 1992; Dette, Melas and Pepelyshev, 2004) to construct Bayesian c- and D-optimal designs,

and further imposed a relaxed Bayesian feasibility constraint on the design to avoid highly toxic doses. Optimal design theory involves a design measure ξ on the dose space $\mathcal{X}$, and a sequential design updates the empirical design measure ξ_{n-1} at stage n by changing it to ξ_n with the addition of the dose x_n. The empirical measure ξ_n of the doses $x_1, \ldots, x_n$ up to stage n can be represented by $\xi_n = n^{-1} \sum_{i=1}^{n} \delta_{x_i}$, where δ_x is the probability measure degenerate at x. We let $||\xi||$ denote the number of x_i (not necessarily distinct) in the support of ξ. Thus $||\xi_n|| = n$ and $||\xi_0|| = 0$, with ξ_0 being the zero measure on $\mathcal{X}$. To include the construction of sequential Bayesian optimal designs as a special case of our general approach, we can modify the preceding procedure that minimizes $E_\Pi[\ell(\theta, x)]$ to choose the next dose based on the current posterior distribution Π, by including the current design measure ξ in the loss function.

Example 2: Bayesian c- or D-optimal designs. As described by Haines, Perevozskaya and Rosenberger (2003), optimal design theory is concerned with choosing a design measure ξ to minimize a convex function Ψ of the information matrix $M(\theta, \xi) = \int I(\theta, x) d\xi(x)$, where $I(\theta, x)$ is the Fisher information matrix at design point x:

$$I(\theta, x) = \frac{e^{\alpha + \beta x}}{(1 + e^{\alpha + \beta x})^2} \begin{pmatrix} 1 & x \\ x & x^2 \end{pmatrix}.$$

The convex function Ψ is associated with the optimality criterion, e.g., $\Psi(M) = -\log \det(M)$ for D-optimality and $\Psi(M) = c'M^{-1}c$ for c-optimality. Since $\theta = (\alpha, \beta)$ is unknown, the frequentist approach uses a sequential design that replaces θ in $M(\theta, \xi_t)$ by its maximum likelihood estimate at every stage t. The Bayesian approach puts a prior distribution Π_0 on θ and minimizes $\int \Psi(M(\theta, \xi_t)) d\Pi_0(\theta)$. Noting that this Bayesian approach does not accomodate the fact that patients are assigned doses sequentially in Phase I trials, Haines, Perevozskaya and Rosenberger (2003, Section 5) propose to start the optimal design after an initial sample of k patients so that the dose x of a patient after this initial sample

can be determined by minimizing

$$\int \Psi\left(\{kM(\theta, \xi_k) + I(\theta, x)\}/(k+1)\right) d\Pi_k(\theta), \tag{4}$$

where ξ_k is the empirical measure of the initial sample of design points and Π_k is the posterior

distribution of θ based on the initial sample.

We can easily extend our loss function approach to Bayes sequential designs by including ξ

as an argument of the loss function in this setting. Let Π be the current posterior distribution

of θ and ξ be the current empirical design measure. Define

$$\ell(\theta, x; \xi) = \Psi(M(\theta, \xi_{+\{x\}})), \quad \text{where} \quad \xi_{+\{x\}} = \frac{||\xi||\xi + \delta_x}{||\xi|| + 1}. \tag{5}$$

The sequential Bayes optimal design chooses the next design level x that minimizes $E_\Pi \ell(\theta, x; \xi)$.

The measure $\xi_{+\{x\}}$ in (5) represents the new empirical measure obtained by adding x to the

support of ξ, with $||\xi_{+\{x\}}|| = ||\xi|| + 1$. We can also impose a relaxed feasibility constraint in

the choice of x:

$$\text{Minimize} \quad E_\Pi \ell(\theta, x; \xi) \quad \text{subject to} \quad P_\Pi(\widetilde{\eta} < x) \leqslant \omega, \tag{6}$$

as in Haines, Perevozskaya and Rosenberger (2003), where $\widetilde{\eta} = F_\theta^{-1}(q)$ with $q \geqslant p$ and ω

is a prescribed positive constant. If $q = p$, then $\widetilde{\eta} = \eta$ and the constraint corresponds to

requiring the doses to be Bayesian feasible (see the description of EWOC above).

3. Coherence and Dilemma Between Individual and Collective Ethics

The preceding section has focused on determining the next dose by minimizing $E_\Pi[\ell(\theta, x)]$,

where Π is the current posterior distribution and ℓ is a loss function incorporating the trial's

main objective into the Bayes sequential design. In Example 2 we have shown how additional

information, such as the empirical measure of previous design points, can be included in the

minimization problem to determine the dose. The following subsections extend this idea to

address two important issues in Phase I cancer clinical trial designs.

3.1 *Coherence and Its Enforcement*

Motivated by ethical concerns, Cheung (2005) introduced coherence principles for sequential dose escalation or de-escalation. A dose sequence is said to be "coherent" if a higher (respectively, lower) dose is not given to the next patient when the current patient experiences (respectively, does not experience) DLT. In particular, CRM and EWOC are coherent and the following theorem, whose proof is given in the Appendix, provides conditions for the coherence of a Bayes sequential design that minimizes the posterior loss at every stage.

THEOREM 1: *Suppose that the dose space is a finite interval and that $\ell(\eta, x)$ is convex in x for every fixed η. Assume that for fixed $x > x'$, $\ell(\eta, x) - \ell(\eta, x')$ is non-increasing in η. Then the dose sequence $x_n = \arg\min_x E_{\Pi_{n-1}} \ell(\eta, x)$ is coherent.*

Theorem 1 shows that CRM is coherent since $\ell(\eta, x) = (\eta - x)^2$ is convex and

$$\ell(\eta, x) - \ell(\eta, x') = -2\eta(x - x') + x^2 - (x')^2$$

is non-increasing in η for $x > x'$. The loss function (2) associated with EWOC also satisfies the assumption of Theorem 1, which therefore shows the coherence of EWOC. The loss functions in Examples 1 and 2, however, may not satisfy the assumptions of Theorem 1. Moreover, a modification of EWOC recommended by its proponents (Babb and Rogatko, 2004), in which the feasibility bound is escalated throughout the trial from a low starting value to $1/2$ at the end of the trial, does not satisfy the assumptions of Theorem 1 and it indeed exhibits slight incoherence in the simulation studies in Section 4. This can be understood by noting that, toward the end of the trial, the posterior distribution does not change much from patient to patient, and that an increase in the feasibility bound may overwhelm the slight downward shift in the posterior following an outcome $y = 0$, causing a dose higher than the previous to be assigned. Cheung (2005, p. 865) also found a certain two-stage modification of CRM to be incoherent. On the other hand, we can enforce coherence by modifying $x_n = f(\Pi_{n-1})$

into $x_n = f(\Pi_{n-1}, x_{n-1}, y_{n-1})$, where

$$
f(\Pi, x^*, y) = \begin{cases} \arg\min_{x \leqslant x^*} E_\Pi \ell(\eta, x) & \text{if } y = 1, \\[2ex] \arg\min_{x \geqslant x^*} E_\Pi \ell(\eta, x) & \text{if } y = 0. \end{cases}
\tag{7}
$$

3.2 *Treatment of Current Patient versus Information for Future Patients*

We have noted in Section 2 that CRM or EWOC treats the next patient at the dose x that minimizes $E_\Pi[\ell(\theta, x)]$ for $\ell(\eta, x)$ given by $(\eta - x)^2$ or by (2), where Π is the current posterior distribution. This is tantamount to dosing the next patient at the best guess of η, where "best" means "closest" according to some measure of distance from η. On the other hand, a Bayesian c- or D-optimal design aims at generating doses that provide most information, as measured by the Fisher information matrix of a design measure, for estimating the dose-toxicity curve to benefit future patients. To resolve this dilemma between treatment of patients in the trial and efficient experimental design for post-trial parameter estimation, Bartroff and Lai (2010) considered the finite-horizon optimization problem of choosing the dose levels $x_1, x_2, \ldots, x_n$ sequentially to minimize the "global risk"

$$
E_{\Pi_0}\left[\sum_{i=1}^n h(\eta, x_i) + g(\widehat{\eta}_n, \eta)\right],
\tag{8}
$$

in which Π_0 denotes the prior distribution of θ, $h(\eta, x_i)$ represents the loss for the ith patient in the trial, $\widehat{\eta}_n$ is the terminal estimate of the MTD and g represents a terminal loss function. The optimizing doses x_i depend on $n - i$, where the horizon n is the sample size of the trial, and therefore are not of the form $x_i = f(\Pi_{i-1})$ considered in Section 2. In terms of "individual" and "collective" ethics, note that (8) measures the individual effect of the dose x_k on the kth patient through $h(\eta, x_k)$, and its collective effect on future patients through $\sum_{i>k} h(\eta, x_i) + g(\widehat{\eta}_n, \eta)$.

By using a discounted infinite-horizon version of (8), we can still have solutions of the form $x_i = f(\Pi_{i-1})$ for some functional f that only depends on Π_{i-1}. Specifically, take a discount

factor $0 < \delta < 1$ and replace (8) by

$$E_{\Pi_0}\left[\sum_{i=1}^{\infty} h(\eta, x_i)\delta^{i-1}\right] \qquad (9)$$

as the definition of global risk. Note that this global risk measures the individual effect of the

dose x_k on the kth patient through $h(\eta, x_k)$, and its collective effect on future patients through

$\sum_{i>k} h(\eta, x_i)\delta^{i-k}$. This means the myopic dose x_k that minimizes $E_{\Pi_{k-1}}[h(\eta, x)]$ for treating

the kth patient has to be perturbed such that it also helps to create a more informative

posterior distribution Π_k that is used for dosing future patients. Note that (9) does not have

the term $g(\widehat{\eta}_n, \eta)$ appearing in the finite-horizon problem (8), but even without this term,

the global risk (9) still captures the collective effect of the doses, as indicated above. As we

have pointed out in Section 2, if x_i is of the form $f(\Pi_{i-1})$ for all i, then $\{\Pi_k : k \geqslant 0\}$ is a

Markov chain whose states are distributions of θ and undergo Markovian dynamics described

by the updating scheme for posterior distributions. In the context of the present problem

of minimizing (9), the optimal expected loss $V(\Pi)$ at state Π (posterior distribution of θ)

satisfies Bellman's dynamic programming equation

$$V(\Pi) = \inf_{x} E_{\Pi}\{h(\eta, x) + \delta E_{\Pi} V(\Pi_{+\{x\}})\}, \qquad (10)$$

where $\Pi_{+\{x\}}$ is the new posterior distribution of θ after (x, y) is observed, with $y \sim \text{Bern}(F_\theta(x))$

and $\theta \sim \Pi$; see the Bayesian updating scheme in Section 2. For finite-state controlled Markov

chains, iteration is a commonly used method to solve (10); see Bertsekas (2007, Section 1.3).

In the present case, not only is the state space infinite, but it is also infinite-dimensional

(space of all posterior distributions of θ), making dynamic programming intractable.

The main complexity of the infinite-horizon problem is that the dose x for the next patient

involves also consideration for future patients who will receive optimal doses themselves;

these future doses depend on the future posterior distributions. A simple way to reduce the

complexity is to consider two (instead of infinitely many) future patients. This amounts to

choosing the next dose x to minimize $E_{\Pi}\ell(\eta, x; \Pi)$ when the current posterior distribution

of θ is Π, where

$$\ell(\eta, x; \Pi) = h(\eta, x) + \lambda E_\Pi \{ E_\Pi [h(\eta', x') | x_1 = x, y_1] \}, \tag{11}$$

in which $\eta' = F_{\theta'}^{-1}(p)$ with $\theta' \sim \Pi'$, and Π' and x' are defined below. The first summand in (11) measures the (toxicity) effect of the dose x on the patient receiving it. The second summand considers the patient who follows and receives a myopic dose x' which minimizes the patient's posterior loss; the myopic dose is optimal because there are no more patients involved in (11). The effect of x on this second patient is through the posterior distribution Π' that updates Π after observing (x_1, y_1), with $x_1 = x$. Since y_1 is not yet observed, the expectation outside the curly brackets is taken over $y_1 \sim \text{Bern}(F_\theta(x))$, with $\theta \sim \Pi$. For example, when implemented with $h(\eta, x)$ given by the EWOC loss function (2), this proposal can be viewed as a modification of EWOC since it utilizes its loss function but adds an additional term to represent the effect on future patients.

Unlike $0 < \delta < 1$ in the discounted infinite-horizon problem, the choice of $\lambda > 0$ in (11) can exceed 1 and reflects the balance between the collective ethics in generating information for future patients and the individual ethics for the patient receiving the dose. Although we use here a single patient to represent all patients following the one receiving the next dose, because the posterior distributions also change successively, the doses are functionals of these posterior distributions.

4. Implementation and a Simulation Study

In this section we first describe three main components in the implementation of the above Bayesian sequential designs and then evaluate their performance in a simulation study.

4.1 *Updating the Posterior Distribution*

Letting η denote the MTD and $\rho = F_\theta(x_{\min})$, we follow Babb, Rogatko and Zacks (1998) to transform (α, β) in (1) to (ρ, η) via the formulas

$$
\begin{aligned}
\alpha &= \frac{x_{\min} \log(p^{-1} - 1) - \eta \log(\rho^{-1} - 1)}{\eta - x_{\min}}, \\
\beta &= \frac{\log(\rho^{-1} - 1) - \log(p^{-1} - 1)}{\eta - x_{\min}},
\end{aligned}
$$

and therefore

$$
\alpha + \beta x = \frac{(x - \eta) \log(\rho^{-1} - 1) - (x - x_{\min}) \log(p^{-1} - 1)}{\eta - x_{\min}} = G(x, \rho, \eta).
$$

We assume that the joint prior distribution of (ρ, η) has density $\pi(\rho, \eta)$ with support on $[0, p] \times [x_{\min}, x_{\max}]$. Therefore the $\mathcal{F}_{k-1}$-posterior density of (ρ, η) is

$$
\pi_{k-1}(\rho, \eta) = C \prod_{i=1}^{k-1} \left[\frac{1}{1 + e^{-G(x_i, \rho, \eta)}} \right]^{y_i} \left[\frac{1}{1 + e^{G(x_i, \rho, \eta)}} \right]^{1 - y_i} \pi(\rho, \eta), \tag{12}
$$

where

$$
C^{-1} = \int_{x_{\min}}^{x_{\max}} \int_0^q \prod_{i=1}^{k-1} \left[\frac{1}{1 + e^{-G(x_i, \rho, \eta)}} \right]^{y_i} \left[\frac{1}{1 + e^{G(x_i, \rho, \eta)}} \right]^{1 - y_i} \pi(\rho, \eta) d\rho d\eta.
$$

The marginal $\mathcal{F}_{k-1}$-posterior distribution of η is then $\int_0^p \pi_{k-1}(\rho, \eta) d\rho$, and the CRM and EWOC doses based on $\mathcal{F}_{k-1}$ are the mean and ω-quantile of this distribution, respectively.

4.2 *Computation of $E_\Pi \ell(\eta, x)$ and its Minimizer in Sections 2 and 3.1*

The integrals in (12) can be evaluated by using a numerical double-integration routine involving Gaussian quadrature in MATLAB. This can be used to evaluate $E_\Pi \ell(\eta, x)$ for a posterior distribution Π. We can find the minimum of $E_\Pi \ell(\eta, x)$ over x by a grid search in $[x_{\min}, x_{\max}]$, or by using gradient descent if ℓ is smooth. For computation of the constrained Bayesian optimal design (6), a constrained nonlinear optimization routine in MATLAB can be used in conjunction with numerical integration, as outlined in Haines, Perevozskaya and Rosenberger (2003, p. 593).

4.3 *Minimization of $E_\Pi \ell(\eta, x; \Pi)$ in Section 3.2*

While MCMC or rejection sampling can be used to compute (11) for any candidate dose x,

importance sampling (e.g., Robert and Casella, 2004, Chapter 3.3) is a simple, robust alter-

native that takes advantage of the fact that just an expectation with respect to the posterior

distribution is needed. Letting Π_0 denote the uniform distribution of the transformed coor-

dinates (ρ, η) over $[0, p] \times [x_{\min}, x_{\max}]$, we have

$$E_\Pi \ell(\eta, x; \Pi) \approx B^{-1} \sum_{b=1}^{B} \ell(\eta_b, x; \Pi) \frac{\pi(\rho_b, \eta_b)}{\pi_0(\rho_b, \eta_b)} \tag{13}$$

for large B, where (ρ_b, η_b), $b = 1, \ldots, B$, are i.i.d. and generated from Π_0. Letting $\Pi_{+\{x,y\}}$

denote the posterior distribution obtained from Π by including (x, y) and letting $x' =$

$x'(\Pi_{+\{x,y\}})$, the nested expectation in (11) can be similarly approximated by using

$$E_\Pi[h(\eta', x')|x_1 = x, y] = E_{\Pi_{+\{x,y\}}} h(\eta', x') \approx B^{-1} \sum_{b=1}^{B} h(\eta_b', x') \frac{\pi_{+\{x,y\}}(\rho_b', \eta_b')}{\pi_0(\rho_b', \eta_b')}, \tag{14}$$

$$P_\Pi(y = 1|x) = \int F_\theta(x) d\Pi(\theta) \approx B^{-1} \sum_{b=1}^{B} F_{\theta_b''}(x) \frac{\pi(\rho_b'', \eta_b'')}{\pi_0(\rho_b'', \eta_b'')}, \tag{15}$$

where (ρ_b', η_b'), $(\rho_b'', \eta_b'') \sim \Pi_0$ and $\theta_b'' = \theta(\rho_b'', \eta_b'')$. Let $H_\Pi(x, y)$ and $Q_\Pi(x)$ denote the right-

hand sides of (14) and (15), respectively. Combining (13)-(15) gives

$$E_\Pi \ell(\eta, x; \Pi) \approx B^{-1} \sum_{b=1}^{B} \{h(\eta_b, x) + \lambda [H_\Pi(x, 0)(1 - Q_\Pi(x)) + H_\Pi(x, 1)Q_\Pi(x)]\} \frac{\pi(\rho_b, \eta_b)}{\pi_0(\rho_b, \eta_b)}. \tag{16}$$

We can minimize the right-hand side of (16) over $x \in [x_{\min}, x_{\max}]$ by using a bounded

minimization routine in MATLAB.

4.4 *Simulation Study*

To compare the proposed procedure in Section 3.2 to EWOC, CRM, and the inverted

overdose control (IVOC) design in Example 1, a simulation study was performed in the

setting of the trial to determine the MTD of the antimetabolite 5-fluorouracil (5-FU) for

treating solid tumors in the colon, as described in Babb, Rogatko and Zacks (1998). Based

on previous studies of 5-FU, a dose of 140 mg/m^2 of 5-FU was believed to be safe, and the

MTD was believed to be no greater than 425 mg/m^2, thus the dose space was taken to be the interval $[x_{\min}, x_{\max}] = [140, 425]$. The two-parameter logistic model (1) was chosen based on previous experience with the agent, and the uniform distribution over $[0, p] \times [x_{\min}, x_{\max}]$ was chosen as the prior distribution Π_0 for (ρ, η), with $p = 1/3$. The feasibility bound of $\omega = .25$ was chosen, which was also used here for the IVOC weight γ in (3). In a trial of length $n = 24$, Table 1 compares EWOC that uses a linearly escalated feasibility bound (Babb and Rogatko, 2004), denoted by EWOC*, with IVOC, CRM, and the proposed design in Section 3.2 with h in (11) given by the EWOC loss function (and denoted by EWOC$_+$, in which + signifies an additional future patient considered by (11)), for two different values of the discount factor λ in (11). Each entry in the table was calculated from 10,000 simulated trials. The first set of rows is a Bayesian setting in which, for each replication, a pair (ρ, η) is drawn from Π_0, and the next three sets of rows are frequentist settings (denoted Freq$_1$, Freq$_2$, Freq$_3$) where the true values (ρ, η) are set at fixed values for all 10,000 replications; these three pairs of fixed values were drawn from Π_0. A comprehensive comparison of EWOC, CRM, sequential c-optimal, constrained D-optimal, ADP and other designs has been given by Bartroff and Lai (2010), who use approximate dynamic programming (ADP) to minimize the finite-horizon risk (8).

[Table 1 about here.]

Table 1 reports two different risk measures. Since the length of the trial is fixed at $n = 24$ in the simulation study, Risk$_1$ is the finite-horizon analog of (9), which is (8) with h given by the EWOC loss function (2) and no terminal loss (i.e., $g = 0$), and Risk$_2$ is the same risk function but with h given by the "inverted" loss function (3). Also reported are the bias and RMSE (root mean squared error $\{E(\widehat{\eta}_n - \eta)^2\}^{1/2}$) of the terminal MTD estimate $\widehat{\eta}_n$ (which is the mean of the terminal posterior distribution of η), the DLT rate $P(y = 1)$ (denoted DLT), the overdose rate $P(x > \eta)$ (denoted OD), the excess DLT rate $E[F_\theta(x) - p]^+$ (denoted

OD*), and the coherence violation rate (denoted ChV)

$$(n-1)^{-1}\sum_{i=1}^{n-1}\{P(y_i=0, x_{i+1} < x_i) + P(y_i=1, x_{i+1} > x_i)\}.$$

In these expressions, P and E denote the probability and expectation, respectively, with respect to the prior distribution in the Bayesian setting, or with respect to the appropriate fixed values of (ρ, η) in the frequentist settings, and are computed by Monte Carlo.

In terms of Risk$_1$ and Risk$_2$, EWOC$_+$ performs better in the Bayesian setting than the myopic designs EWOC*, IVOC, and CRM, in that order. This occurs in the frequentist settings as well, although the ordering of the myopic designs varies depending on the particular parameter values. Even though it myopically minimizes the posterior risk at every stage, IVOC performs poorly in terms of the cumulative risk, Risk$_2$, in the Bayesian setting. A possible explanation is that its loss function (3) is a function of $F_\theta(x)$, whose posterior distribution (induced by the posterior distribution of θ) has relatively large variance toward the middle of the interval $(0,1)$ in which $F_\theta(x)$ takes values and, in particular, near $p = 1/3$, resulting in low initial doses observed in the simulations. On the other hand, in the Freq$_3$ setting, where η is relatively small and the dose-response curve is relatively flat (e.g., ρ large), IVOC performs well in terms of the risks. In terms of estimation, EWOC$_{+,\lambda=.4}$ has the smallest RMSE, with EWOC* and EWOC$_{+,\lambda=.1}$ both comparable in the Bayesian setting. Moreover, EWOC$_{+,\lambda=.4}$ has uniformly the smallest RMSE in the frequentist settings, with IVOC comparable to it in Freq$_2$ and IVOC and EWOC$_{+,\lambda=.1}$ comparable to it in Freq$_3$.

5. Conclusion and Discussion

In this paper we present a general formulation of Bayesian sequential design of Phase I cancer trials. This formulation enables us to prove a general coherence result in Theorem 1 applicable to any design that can be defined as the minimizer of the posterior risk when the loss function satisfies some mild conditions. Although the theorem is proved for the widely-

used logistic regression model (1), the last paragraph of its proof in the Appendix shows that it is applicable to any dose-response model that is non-increasing in the MTD, such as the model $F_\theta(x) = \{(\tanh x + 1)/2\}^\theta$, which is also popular.

In Section 3.2 we propose a new design that incorporates both the individual ethics of the current patient begin administered the dose, through a given loss function such as the EWOC loss (2), and the collective ethics of all future patients by including an additional term in the overall loss function to represent the dose's information content for determining another dose for the next patient. The simulation study in Section 4.4 shows that this new design is indeed an improvement over myopic designs in terms of global risk minimization, post-trial estimation of the MTD, and DLT and OD rates. This design provides a practical alternative to the optimal design associated with the intractable Markov decision problem of minimizing (9), which requires at each stage the daunting consideration of all future posterior distributions and calculating their associated optimal doses. For the finite-horizon problem of minimizing (8), Bartroff and Lai (2010) have developed an approximate solution which is a time-varying mixture of myopic and c-optimal designs. The new design in Section 3.2, which can be described by a time-invariant functional of the posterior distribution at each stage, is substantially simpler computationally and provides substantial improvement over the myopic designs. We conjecture that with suitably chosen λ (depending on δ), its global risk (9) can approximate that of the optimal design minimizing (9). Instead of minimizing (9) directly, it may be possible to obtain a good lower bound for (9). Such a bound, which can provide a benchmark for assessing the proposed design, is a topic for future work.

We also consider an "inverted" loss function (3), which measures deviation from the target DLT rate p on the probability scale rather than on the dose scale, and the associated myopic design IVOC. Even though IVOC minimizes the myopic posterior expected loss (3) at each stage, its cumulative global loss Risk$_2$ in Table 1 is far from optimum, exceeding even that of

EWOC which uses a completely different loss function, in the Bayesian setting. On the other hand, the design proposed in Section 3.2 can be applied with the IVOC loss function (3) to yield a substantially improved design IVOC$_+$.

<center>APPENDIX</center>

Proof of Theorem 1. We prove coherence in de-escalation; the proof for escalation is similar. Let $x_{\min} < x_{\max}$ be the boundaries of the dose space, which is assumed to be a finite interval. For fixed η, since $\ell(\eta, x)$ is a convex function of x, its right derivative $\ell_x(\eta, x)$ with respect to x is nondecreasing for $x_{\min} \leqslant x < x_{\max}$, and the same is also true for the left derivative for $x_{\min} < x \leqslant x_{\max}$. Moreover, the left and right derivatives are equal and continuous except for at most countably many points; see Rockafellar (1970, pages 214, 228, 244). Let $x_\Pi = \arg\min_x E_\Pi \ell(\eta, x)$, $\widetilde{\Pi}$ be the posterior distribution obtained from Π and the additional dose-response pair $(x, y) = (x_\Pi, 1)$, and let $L(x) = E_{\widetilde{\Pi}} \ell(\eta, x)$. Since $\ell(\eta, x)$ is convex in x for every η, so is $L(x)$; moreover, its right derivative is given by $\dot{L}_+(x) = E_{\widetilde{\Pi}} \ell_x(\eta, x)$. To show that $x_{\widetilde{\Pi}} \leqslant x_\Pi$, we shall assume that $x_\Pi < x_{\max}$ because the case $x_\Pi = x_{\max}$ is trivial. It suffices to show that $\dot{L}_+(x_\Pi) \geqslant 0$ because L is convex and has minimizer $x_{\widetilde{\Pi}}$. Since $E_\Pi \ell_x(\eta, x_\Pi) \geqslant 0$ and $d\widetilde{\Pi}(\theta) = F_\theta(x_\Pi) d\Pi(\theta) / \int F_{\theta'}(x_\Pi) d\Pi(\theta')$, recalling that $(x, y) = (x_\Pi, 1)$, it follows that

$$
\begin{aligned}
\dot{L}_+(x_\Pi) &\geqslant E_{\widetilde{\Pi}} \ell_x(\eta, x_\Pi) - E_\Pi \ell_x(\eta, x_\Pi) \\
&= \frac{\int \ell_x(\eta, x_\Pi) F_\theta(x_\Pi) d\Pi(\theta)}{\int F_{\theta'}(x_\Pi) d\Pi(\theta')} - \frac{\int \ell_x(\eta, x_\Pi) d\Pi(\theta)}{\int d\Pi(\theta')} \\
&= A \left/ \int F_{\theta'}(x_\Pi) d\Pi(\theta') \right. ,
\end{aligned}
\tag{17}
$$

where $A = \int \int \ell_x(\eta, x_\Pi)[F_\theta(x_\Pi) - F_{\theta'}(x_\Pi)] d\Pi(\theta) d\Pi(\theta')$. A change of variables also yields $A = -\int \int \ell_x(\eta', x_\Pi)[F_\theta(x_\Pi) - F_{\theta'}(x_\Pi)] d\Pi(\theta) d\Pi(\theta')$. Hence

$$
2A = \int \int [\ell_x(\eta, x_\Pi) - \ell_x(\eta', x_\Pi)][F_\theta(x_\Pi) - F_{\theta'}(x_\Pi)] d\Pi(\theta) d\Pi(\theta') \geqslant 0,
\tag{18}
$$

in which the inequality follows from

$$[\ell_x(\eta, x_\Pi) - \ell_x(\eta', x_\Pi)][F_\theta(x_\Pi) - F_{\theta'}(x_\Pi)] \geqslant 0 \qquad (19)$$

for all x, θ and θ', as will be shown below. Combining (17) and (18) yields $\dot{L}_+(x_\Pi) \geqslant 0$, completing the proof of the theorem.

From the assumption that $\ell(\eta, x) - \ell(\eta, x')$ is non-increasing in η for any $x > x'$, it follows that $\ell_x(\eta, x)$ is non-increasing in η for fixed x. It therefore suffices for the proof of (19) to show that $F_\theta(x)$ is non-increasing in $\eta = F_\theta^{-1}(p)$. Since $p^{-1} = 1/F_\theta(\eta) = 1 + e^{-(\alpha+\beta\eta)}$, $F_\theta(x) = 1/[1 + \exp\{\log(p^{-1} - 1) + \beta\eta - \beta x\}]$, which is non-increasing in η since $\beta > 0$.

References

Atkinson, A. C. and Donev, A. N. (1992). *Optimum Experimental Designs.* Oxford University Press.

Babb, J. and Rogatko, A. (2004). Bayesian methods for cancer phase I clinical trials. In Geller, N. L., editor, *Advances in Clinical Trial Biostatistics.* Marcel Dekker.

Babb, J., Rogatko, A., and Zacks, S. (1998). Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine* **17,** 1103–1120.

Bartroff, J. and Lai, T. L. (2010). Approximate dynamic programming and its applications to the design of phase I cancer trials. *Statistical Science*, in press.

Bertsekas, D. P. (2007). *Dynamic Programming and Optimal Control*, volume 2. Athena Scientific, Belmont, MA, 3rd edition.

Cheung, Y. K. (2005). Coherence principles in dose-finding studies. *Biometrika* **92,** 863–873.

Dette, H., Melas, V. B., and Pepelyshev, A. (2004). Optimal designs for a class of nonlinear regression models. *The Annals of Statistics* **32,** 2142–2167.

Dragalin, V. and Fedorov, V. (2006). Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* **136,** 1800–1823.

Dragalin, V., Fedorov, V., and Wu, Y. (2008). Adaptive designs for selecting drug combinations based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* **138,** 352–373.

Fedorov, V. V. (1972). *Theory of Optimal Experiments.* Academic Press, New York.

Haines, L. M., Perevozskaya, I., and Rosenberger, W. F. (2003). Bayesian optimal design for phase I clinical trials. *Biometrics* **59,** 591–600.

Hardwick, J. and Stout, Q. F. (2001). Optimizing a unimodal response function for binary variables. In Atkinson, A., Bogacka, B., and Zhigljavsky, A., editors, *Optimum Design 2000*, pages 195–210. Kluwer Academic Publishers, Dordrecht.

Kpamegan, E. E. and Flournoy, N. (2001). An optimizing up-and-down design. In Atkinson, A., Bogacka, B., and Zhigljavsky, A., editors, *Optimum Design 2000.* Kluwer Academic Publishers, Dordrecht.

Li, Z., Durham, S. D., and Flournoy, N. (1995). An adaptive design for maximization of a contingent binary response. In Flournoy, N. and Rosenberger, W. F., editors, *Adaptive Designs*, pages 179–196. Institute of Mathematical Statistics.

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.

Pronzato, L. (2010). Penalized optimal designs for dose-finding. *Journal of Statistical Planning and Inference* **140,** 283–296.

Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods.* Springer-Verlag, New York, 2nd edition.

Rockafellar, R. T. (1970). *Convex Analysis.* Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J.

Rogatko, A., Schoeneck, D., Jonas, W., Tighiouart, M., Khuri, F., and Porter, A. (2007). Translation of innovative designs into phase I trials. *Journal of Clinical Oncology* **25,** 4982–4986.

Rosenberger, W. F. and Haines, L. M. (2002). Competing designs for phase I clinical trials: A review. *Statistics in Medicine* **21,** 2757–2770.

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

Von Hoff, D. and Turner, J. (1991). Response rates, duration of response, and dose response effects in phase I studies of antineoplastics. *Investigational New Drugs* **9,** 115–122.

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

Zacks, S., Rogatko, A., and Babb, J. (1998). Optimal Bayesian-feasible dose escalation for cancer phase I trials. *Statistics & Probability Letters* **38,** 215–220.

Table 1

Risk$_1$, Risk$_2$, bias and RMSE of the final MTD estimate, DLT rate, MTD overdoes rate (OD), excess DLT rate $E[F_\theta(x) - p]^+$ (OD), and coherence violation rate (ChV), with SEs in parentheses, of various designs.*

Statistic	EWOC*	IVOC	CRM	EWOC$_{+,\lambda=.1}$	EWOC$_{+,\lambda=.4}$
		Bayesian: $(\rho, \eta) \sim \Pi_0$			
Risk$_1$	485.5 (3.6)	723.2 (3.6)	986.1 (45.9)	469.5 (3.0)	454.8 (2.8)
Risk$_2$	1.03 (.01)	1.44 (.02)	1.54 (.02)	.85 (.01)	.73 (.007)
Bias	-9.67 (.6)	-50.2 (.8)	22.3 (1.6)	2.42 (.7)	-5.9 (.6)
RMSE	61.5 (.7)	75.8 (.2)	157.3 (.8)	66.9 (.2)	58.6 (.7)
DLT (%)	33.5 (.001)	26.6 (9×10^{-4})	39.1 (.001)	29.3 (.0009)	29.1 (9×10^{-4})
OD (%)	37.4 (.001)	17.5 (8×10^{-4})	55.6 (.001)	29.6 (9×10^{-4})	27.0 (9×10^{-4})
OD*	.043 (2×10^{-4})	.043 (3×10^{-4})	.078 (3×10^{-4})	.029 (2×10^{-4})	.021 (1×10^{-4})
ChV (%)	0 (0)	.6 (.04)	0 (0)	15.0 (8×10^{-4})	14.4 (8×10^{-4})
		Freq$_1$: $\rho = .07, \eta = 403.9$			
Risk$_1$	585.6 (1.3)	1302.2 (.8)	368.8 (1.1)	175.9 (1.0)	170.3 (1.1)
Risk$_2$	.78 (.002)	1.40 (6×10^{-4})	.53 (.001)	.29 (.001)	.25 (.001)
Bias	-51.1 (.2)	-151.7 (.2)	-30.1 (.4)	-23.2 (.3)	-48.7 (.2)
RMSE	58.8 (.8)	156.1 (.6)	44.9 (.1)	49.0 (.5)	19.2 (.4)
DLT (%)	20.6 (8×10^{-4})	10.3 (7×10^{-4})	24.9 (9×10^{-4})	18.7 (8×10^{-4})	18.6 (8×10^{-4})
OD (%)	0 (0)	0 (0)	2.3 (3×10^{-4})	1.2 (2×10^{-4})	1.1 (2×10^{-4})
OD*	0 (0)	0 (0)	4×10^{-4} (1×10^{-5})	.001 (1×10^{-5})	.001 (1×10^{-5})
ChV (%)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
		Freq$_2$: $\rho = .19, \eta = 269.1$			
Risk$_1$	402.8 (.8)	423.1 (.9)	313.7	275.0 (3.7)	266.2 (3.8)
Risk$_2$	.53 (.003)	.49 (.001)	.99 (.006)	.43 (.001)	.26 (.001)
Bias	15.0 (.5)	-28.9 (.1)	32.5 (.4)	17.3 (.7)	12.3 (.4)
RMSE	47.5 (.2)	32.0 (.7)	52.1 (.7)	38.2 (.3)	25.2 (.3)
DLT (%)	32.6 (.001)	25.4 (9×10^{-4})	38.2 (.001)	28.2 (9×10^{-4})	24.8 (9×10^{-4})
OD (%)	32.8 (.001)	0 (0)	83.1 (8×10^{-4})	27.1 (9×10^{-4})	26.9 (9×10^{-4})
OD*	.02 (7×10^{-5})	0 (0)	.053 (8×10^{-5})	.03 (4×10^{-5})	.01 (4×10^{-5})
ChV (%)	4.3 (4×10^{-4})	3.1 (4×10^{-4})	0 (0)	17.6 (8×10^{-4})	12.7 (7×10^{-4})
		Freq$_3$: $\rho = .30, \eta = 226.7$			
Risk$_1$	674.3 (4.9)	232.1 (2.9)	1444.9 (6.3)	158.9 (4.9)	146.1 (3.5)
Risk$_2$	.27 (.002)	.09 (4×10^{-4})	.59 (.003)	.05 (4×10^{-4})	.05 (6×10^{-4})
Bias	42.9 (.62)	10.3 (.14)	81.4 (.5)	54.9 (.54)	47.4 (.49)
RMSE	74.2 (.7)	17.4 (.5)	96.8 (.8)	18.0 (.5)	13.2 (.2)
DLT (%)	34.2 (.001)	32.7 (.001)	36.6 (.001)	35.8 (.001)	35.7 (.001)
OD (%)	73.2 (.001)	54.3 (.001)	93.5 (5×10^{-4})	88.8 (6×10^{-4})	81.0 (6×10^{-4})
OD*	.014 (3×10^{-5})	.002 (2×10^{-5})	.032 (4×10^{-5})	.031 (4×10^{-5})	.030 (4×10^{-5})
ChV (%)	0 (0)	3.9 (5×10^{-4})	0 (0)	10.0 (6×10^{-4})	9.6 (6×10^{-4})