Published in final edited form as:

Biometrics. 2010 June ; 66(2): 532–540. doi:10.1111/j.1541-0420.2009.01302.x.

Utility-Based Optimization of Combination Therapy Using Ordinal Toxicity and Efficacy in Phase I/II Trials

Nadine Houede¹, Peter F. Thall^{2,*}, Hoang Nguyen², Xavier Paoletti³, and Andrew Kramar⁴

¹ Department of Medical Oncology, Institut Bergonie, Bordeaux, France

² Department of Biostatistics, M.D. Anderson Cancer Center, Houston, USA

³ Department of Biostatistics, Institut Curie, Paris, France

⁴ Biostatistics Unit, CRLC Val d'Aurelle, Montpellier, France

Abstract

An outcome-adaptive Bayesian design is proposed for choosing the optimal dose pair of a chemotherapeutic agent and a biologic agent used in combination in a phase I/II clinical trial. Patient outcome is characterized as a vector of two ordinal variables accounting for toxicity and treatment efficacy. A generalization of the Aranda-Ordaz model (1983, *Biometrika* **68**, 357–363) is used for the marginal outcome probabilities as functions of dose pair, and a Gaussian copula is assumed to obtain joint distributions. Numerical utilities of all elementary patient outcomes, allowing the possibility that efficacy is inevaluable due to severe toxicity, are obtained using an elicitation method aimed to establish consensus among the physicians planning the trial. For each successive patient cohort, a dose pair is chosen to maximize the posterior mean utility. The method is illustrated by a trial in bladder cancer, including simulation studies of the method's sensitivity to prior parameters, the numerical utilities, correlation between the outcomes, sample size, cohort size and starting dose pair.

Keywords

Adaptive design; Bayesian design; Clinical trial; Combination dose-finding; Utility

1 Introduction

Most of the statistical literature on phase I clinical trial designs deals with the simple case in which it is assumed that each patient is treated with a dose, x , of a single agent (Storer, 1989; O'Quigley, Pepe, and Fisher, 1990; Babb, Rogatko, and Zacks, 1998). Typically, patient outcome is characterized by a binary indicator of dose-limiting toxicity (hereafter, "toxicity") that is observed quickly enough to choose doses adaptively for successive patient cohorts. In general, the goal is to choose x from a predetermined set or interval of values to achieve an acceptable probability of toxicity. In actual clinical practice, however, both the treatment regime and patient outcome are much more complex, often including multiple agents given in combination as well as multiple outcomes. To deal with this complexity, in recent years many authors have addressed the dose-finding problem more generally, going well beyond the simple paradigm described above. Some useful extensions characterize clinical outcome more fully, as time-to-toxicity (Cheung and Chappell, 2000), as an ordinal variable accounting for toxicity severity (Yuan, Chappell, and Bailey, 2007), as a multivariate outcome accounting for both

*Corresponding author, rex@mdanderson.org.

efficacy and toxicity (Thall and Russell, 1998; O'Quigley, Hughes, and Fenton, 2001; Braun, 2002; Ivanova, 2003; Thall and Cook, 2004; Bekele and Shen, 2005; Thall, Nguyen and Estey, 2008), often called “phase I/II” designs (Zohar and Chevret, 2008), or as a vector of different types of toxicity (Bekele and Thall, 2004). Another useful extension deals with the problem of optimizing the dose pair of two agents given in combination (Korn and Simon, 1993; Thall, Millikan, Mueller, and Lee, 2003; Ivanova and Wang, 2005).

In this paper, we address the phase I/II design problem of choosing a dose pair corresponding to a biologic agent and a chemotherapeutic agent given in combination based on a bivariate outcome including toxicity and treatment efficacy. Each outcome is ordinal, thereby accounting for the severity of each toxicity and the level of efficacy achieved. This provides a more informative summary of patient outcome than that provided by binary indicators, which are commonly used in dose-finding designs. The underlying probability model is Bayesian, constructed from two marginal distributions, one for each outcome variable, and using a Gaussian copula to obtain a joint distribution that accounts for association between the outcomes. For each marginal distribution, to accommodate a broad range of possible dose-response functions, including functions that are not monotone in dose, we formulate a generalization of the Aranda-Ordaz (AO) model (1981). The extended AO model accommodates two doses and ordinal outcomes, rather than the more commonly used binary outcomes. This model is inherently robust because it includes as special cases a wide range of commonly used dose-outcome models. The desirability of each possible patient outcome is quantified by a numerical utility specified by a team of physicians using the Delphi method (Dalkey, 1969; Brook, Chassin, Fink et al., 1986). During the trial, each patient's dose pair is chosen adaptively from a two-dimensional grid of possible values to optimize the posterior expected utility of the patient's outcome. The two-agent problem that we address is similar to that considered by Mandrekar, Cui and Sargent (2007), with the important differences that we characterize patient outcome as a vector of two ordinal variables rather than one three-category variable, consequently our dose-outcome model is much more structured, and we choose dose pairs based on posterior expected utilities.

The bladder cancer trial that motivated this research is described in Section 2, followed by a description of the probability model in Section 3. Application of the method to the bladder cancer trial is described in Section 4, including the processes of eliciting priors and utilities. A computer simulation study in the context of this trial is summarized in Section 5, including studies of the method's robustness to model assumptions and numerical utility values. We close with a discussion in Section 6.

2 Motivating Application

To provide a concrete frame of reference, we first describe the bladder cancer trial. The goal of the trial is to evaluate the safety and efficacy of a combination therapy including the chemotherapeutic agents gemcitabine and cisplatin, and a biological agent, in patients with previously untreated advanced bladder cancer. For the purpose of dose-finding, patient outcomes will be evaluated over the first two 28-day cycles of therapy. In each cycle, each patient will receive (1) a fixed dose of 70 mg/m² cisplatin on day 2, (2) the biological agent given orally on each day at one of four possible dose levels, $d_1 = 1, 2, 3$ or 4, and (3) gemcitabine on days 1, 8 and 15 at one of three possible levels, 750, 1000 or 1250 mg/m², coded hereafter as $d_2 = 1, 2$ or 3, with $\mathbf{d} = (d_1, d_2)$. A total of 48 patients will be treated in cohorts of 3, with the first cohort treated at $\mathbf{d} = (2, 2)$. The trial starts at this pair, rather than at the lowest pair (1, 1), because dose level 2 of the chemotherapeutic agent as monotherapy is the recommended dose in the standard of care and it is undesirable to undertreat the first cohort of patients enrolled in the trial, while dose level 2 of the biological agent is 50% of the dose when it is used as monotherapy and it is desired to avoid exposing the first cohort to an unacceptably high risk

of toxicity with the combination. This sort of reasoning, based on previous experience with each agent as monotherapy and considering the unavoidable tension between underdosing and overdosing the first few patients in the trial before any outcomes are observed, is used to choose the starting dose pair whenever a new combination is studied. From this starting dose pair, the method chooses each successive cohort's d to optimize the current posterior expected utility, as described in Section 3.3 below, subject to the constraint that untried dose pairs may not be skipped when escalating, as described in Section 3.4.

Toxicity is a three-level ordinal variable representing the worst severity of non-haematologic adverse events such as fatigue, diarrhea and mucositis, which may be considered to be related to the biological agent, and hematologic toxicities including renal dysfunction and neurotoxicity, which may be considered chemotherapy related. Within-patient dose reduction of the biological agent is done after the first cycle if a grade 1 or 2 non-haematologic toxicity is observed. If the patient does not recover from a grade 3 or 4 non-haematologic toxicity within two weeks, the biological agent is stopped, but the patient may continue to receive the chemotherapeutic agent as a clinical decision of the patient's attending physician. That is, as with any dose-finding trial, to protect patient safety there is an established protocol for modifying each patient's doses on the basis of early interim outcomes before toxicity is formally scored. In any case, the patient's assigned dose pair and observed toxicity outcome, as defined formally below in Section 3, are incorporated into the likelihood and used to choose the next dose pair.

Efficacy is characterized by a three-level ordinal variable characterizing changes in a set of measurable lesions compared to baseline using RECIST criteria (Therasse et al., 2000): tumor response, characterized as complete or partial remission (CR/PR), stable disease (SD), or progressive disease (PD). Nearly all chemotherapy trials include rules, based on the patient's history of treatments and interim outcomes, for deciding whether to continue or terminate the patient's therapy, or for modifying the patient's assigned dose. Such rules typically reflect established clinical practice. For example, treatment can be stopped because disease is progressing or because a severe adverse event such as dose-limiting toxicity has occurred, or dose can be decreased during a cycle of therapy if moderate toxicity is observed. In the bladder cancer trial, if either PD occurs or grade ≥ 3 toxicity related to the chemotherapy agent occurs and is not resolved within 2 weeks then the patient's therapy is terminated. Thus, efficacy may be scored at the end of the second cycle of therapy for patients who receive two full cycles, or as PD at the last efficacy evaluation before two cycles, or efficacy may be inevaluable if treatment is terminated early because severe toxicity cannot be resolved.

The problem that we address is how to choose each successive patient cohort's optimal dose pair of the biological agent and chemotherapeutic agent adaptively from the 12 possible pairs based on the above bivariate outcome. Because dose pairs are chosen to optimize posterior expected utility, our proposed method accounts explicitly for the not uncommon outcome in which efficacy is inevaluable due to early excessive toxicity, rather than simply ignoring this as a "missing" outcome or using only the observed toxicities.

3 Probability Model

3.1 A General Model

Let Y_1 and Y_2 denote toxicity and efficacy, respectively. In general, we define $Y_1 = 0$ if toxicity does not occur, with $Y_1 = 1, 2, \dots, m_1$ corresponding to increasing levels of severity. For efficacy, $Y_2 = 0$ for the worst possible efficacy outcome, such as PD, with $Y_2 = 1, 2, \dots, m_2$ corresponding to increasingly desirable outcomes. To account for the possibility that Y_2 may not be evaluable, we define $Z = 1$ if Y_2 is evaluable and $Z = 0$ if not, with $\zeta = \Pr(Z = 1)$. Thus, the outcome vector $\mathbf{Y} = (Y_1, Y_2)$ is observed if $Z = 1$, whereas only Y_1 is observed if $Z = 0$. We

denote the two-dimensional joint outcome probability mass function (pmf) by $\pi(\mathbf{y}|\mathbf{d}, \theta) = \Pr(\mathbf{Y} = \mathbf{y} | Z = 1, \mathbf{d}, \theta)$, where $\mathbf{y} = (y_1, y_2)$ denotes any of the $(m_1 + 1)(m_2 + 1)$ possible values of $\mathbf{Y}$ when $Z = 1$. We denote the marginal of Y_2 by $\pi_{2,y}(\mathbf{d}, \theta) = \Pr(Y_2 = y | Z = 1, \mathbf{d}, \theta)$, and we assume that $\pi_{1,y}(\mathbf{d}, \theta) = \Pr(Y_1 = y | Z, \mathbf{d}, \theta)$, that is, Z does not affect the marginal distribution of toxicity. Let $\delta(\mathbf{y})$ indicate the event $(\mathbf{Y} = \mathbf{y})$ if $Z = 1$ and let $\delta_1(y_1)$ indicate $(Y_1 = y_1)$, so that the likelihood takes the form

$$\mathcal{L}(Y, Z|\mathbf{d}, \theta) = \left[\zeta \prod_{\mathbf{y}} \{\pi(\mathbf{y}|\mathbf{d}, \theta)\}^{\delta(\mathbf{y})} \right]^Z \left[(1 - \zeta) \prod_{y_1} \{\pi_{1,y_1}(\mathbf{d}, \theta)\}^{\delta_1(y_1)} \right]^{1-Z}. \quad (1)$$

In practice, Z may depend on Y_1 , the observed or unobserved Y_2 , or latent variables that may or may not be related to either entry of $\mathbf{Y}$, including patient dropout or the decision algorithms used by physicians for terminating treatment early, which frequently are complex and may differ substantially from trial to trial. Since Y_2 is defined only if $Z = 1$ whereas Y_1 is defined for either value of Z , the marginal probability in the second product in (1) may be expressed more precisely as $\Pr(Y_1 = y_1 | Z = 0, \mathbf{d}, \theta)$. The assumption made earlier, that $\Pr(Y_1 = y_1 | Z = 0, \mathbf{d}, \theta) = \Pr(Y_1 = y_1 | Z = 1, \mathbf{d}, \theta)$, is needed to allow our construction, given below in Section 3.2, of a tractable parametric model wherein the joint distribution $\pi(\mathbf{y}|\mathbf{d}, \theta)$ is defined in terms of the marginals $\pi_{1,y}(\mathbf{d}, \theta)$ and $\pi_{2,y}(\mathbf{d}, \theta)$, regardless of Z . The simplifying assumption that $\pi_{1,y}(\mathbf{d}, \theta)$ does not depend on Z thus is motivated by a desire for tractability. If one wishes to account for a possible association between Z and Y_1 , however, then a more complex model might be used.

For the bladder cancer trial, both outcomes are three-level ordinal variables. Since it is desired to distinguish between severe toxicities that are or are not resolved, we define $Y_1 = 0$ if severe (grade 3,4) toxicity does not occur, $Y_1 = 1$ if severe toxicity occurs but is resolved within two weeks, and $Y_1 = 2$ if severe toxicity occurs and is not resolved within two weeks. The efficacy outcome is $Y_2 = 0$ if PD occurs at any time, $Y_2 = 1$ if the patient has SD at the end of the second cycle, and $Y_2 = 2$ if tumor response (PR or CR) is achieved. Treatment is terminated if PD is observed early, e.g. at the end of the first 28-day cycle.

3.2 Parametric Models

The particular parametric model chosen for $\pi(\mathbf{y}|\mathbf{d}, \theta)$ must be sufficiently tractable to facilitate application, including computation of the Bayesian posterior decision criteria thousands of times when simulating the trial to establish the design's properties. To account for the joint effects of d_1 and d_2 on each entry of $\mathbf{Y}$, as well as association among its elements, we first model the marginal distribution of each outcome as a function of $\mathbf{d}$ and then combine the marginals using a Gaussian copula to obtain a joint distribution that accounts for association among the elements of $\mathbf{Y}$. For the marginals, we will use a very flexible family obtained by generalizing the AO model for the probability of a binary outcome as a function of a single dose d . Given linear term $\eta(d, \alpha)$ that is a real-valued function of d parameterized by α , the AO model is given by

$$\xi\{\eta(d, \alpha), \lambda\} = 1 - (1 + \lambda e^{\eta(d, \alpha)})^{-1/\lambda} \quad (2)$$

where $\lambda > 0$. In practice, $\eta(d, \alpha)$ may be tailored to the particular application at hand. For example, the AO model with $\eta(d, \alpha) = \alpha_0 + \alpha_1 \log(d)$ has the property that $d = 0$ implies $\xi = 0$, that is, the outcome is impossible if no treatment is given, whereas if $\eta(d, \alpha) = \alpha_0 + \alpha_1 d$ then

$d = 0$ implies $\xi = 1 - (1 + \lambda e^{a_0})^{-1/\lambda}$ is a baseline toxicity rate, possibly due to other treatment components that are not varied. A very useful property of the AO model is that ξ is not monotone in d and the curve may take on a wide variety of different possible shapes. The AO model contains the logistic model $\xi\{\eta(d, \alpha), 1\} = e^{\eta(d, \alpha)} / \{1 + e^{\eta(d, \alpha)}\}$ as a special case, and since $\lim_{\lambda \rightarrow 0} \xi\{\eta(d, \alpha), \lambda\} = 1 - \exp\{-e^{\eta(d, \alpha)}\}$ it contains the generalized linear model with complementary log-log link as a limiting case.

To account for effects of the dose pair $\mathbf{d} = (d_1, d_2)$, we generalize (2) to accommodate two linear terms, one for each dose. Denoting the linear term $\eta_j = \eta_j(d_j, \alpha_j)$ for the effects of the dose of agent $j = 1, 2$, we define the more general probability function

$$\xi^*\{\eta_1, \eta_2, \lambda, \gamma\} = 1 - \{1 + \lambda(e^{\eta_1} + e^{\eta_2} + \gamma e^{\eta_1 + \eta_2})\}^{-1/\lambda}. \quad (3)$$

In this extended model, α_j parameterizes the linear term associated with d_j while γ accounts for interaction between the two agents. Values of $\gamma > 0$ ensure that ξ^* is a probability, although $0 > \gamma > -(e^{-\eta_1} + e^{-\eta_2})$ for all α allows a negative interactive effect.

To specify marginal probabilities for the ordinal outcomes, we first define agent-specific linear terms and then incorporate them into a model that assumes the form (3) for the conditional probabilities $\Pr(Y_k \geq y | Y_k \geq y - 1)$ as follows. We assume that, for each outcome $k = 1, 2$, agent $j = 1, 2$, and outcome $y = 1, \dots, m_k$, the linear terms that determine the marginal distribution of Y_k take the form

$$\eta_{k,y}^{(j)}(d_j, \alpha_k^{(j)}) = \alpha_{k,y,0}^{(j)} + \alpha_{k,y,1}^{(j)} d_j, \quad (4)$$

so that $\alpha_{k,y,0}^{(1)}, \alpha_{k,y,0}^{(2)}$ are intercepts and $\alpha_{k,y,1}^{(1)}, \alpha_{k,y,1}^{(2)}$ are dose effect parameters. To stabilize variances, in (4) each d_j is re-coded by centering it at the mean of its possible values. E.g., for the bladder cancer trial, the numerical values of d_1 are $\{-1.5, -0.5, 0.5, 1.5\}$ and the numerical values of d_2 are $\{-1.0, 0, 1.0\}$. Denoting $\alpha_k^{(j)} = (\alpha_{k,1,0}^{(j)}, \alpha_{k,1,1}^{(j)}, \alpha_{k,2,0}^{(j)}, \alpha_{k,2,1}^{(j)})$, the marginal distribution of Y_k is parameterized by the 10-dimensional vector $\theta_k = (\alpha_k^{(1)}, \alpha_k^{(2)}, \lambda_k, \gamma_k)$, that is, $\pi_{k,y}(\mathbf{d}, \theta) = \pi_{k,y}(\mathbf{d}, \theta_k)$. We incorporate the linear terms into the marginal distributions by assuming that, for each $y = 1, \dots, m_k$,

$$\Pr(Y_k \geq y | Y_k \geq y - 1, \mathbf{d}, \theta_k) = \xi^*\{\eta_{k,y}^{(1)}(d_1, \alpha_k^{(1)}), \eta_{k,y}^{(2)}(d_2, \alpha_k^{(2)}), \lambda_k, \gamma_k\} \equiv \xi_{k,y}^*(\mathbf{d}, \theta_k). \quad (5)$$

This implies that

$$\pi_{k,y}(\mathbf{d}, \theta_k) = \{1 - \xi_{k,y+1}^*(\mathbf{d}, \theta_k)\} \prod_{j=1}^y \xi_{k,j}^*(\mathbf{d}, \theta_k). \quad (6)$$

for all $y \geq 1$, and $\pi_{k,0}(\mathbf{d}, \theta_k) = 1 - \xi_{k,1}^*(\mathbf{d}, \theta_k)$. In the three-level case at hand, the unconditional marginal probabilities are given by $\pi_{k,1}(\mathbf{d}, \theta_k) = \xi_{k,1}^*(\mathbf{d}, \theta_k) \{1 - \xi_{k,2}^*(\mathbf{d}, \theta_k)\}$ and $\pi_{k,2}(\mathbf{d}, \theta_k) = \xi_{k,1}^*(\mathbf{d}, \theta_k) \xi_{k,2}^*(\mathbf{d}, \theta_k)$.

The above model for the marginal distribution $\{\pi_{k,1}(\mathbf{d}, \theta_k), \pi_{k,2}(\mathbf{d}, \theta_k)\}$ of Y_k is especially flexible and informative for dose-finding. While each $\eta_{k,y}^{(j)}(d_j, \alpha_k^{(j)})$ is linear in d_j , the probability functions $\xi_{k,1}^*(\mathbf{d}, \theta_k)$ and $\xi_{k,2}^*(\mathbf{d}, \theta_k)$ take on a wide variety of possible shapes as functions of d_1 and d_2 , and they accommodate the possibility that $\Pr(Y_k \geq y | \mathbf{d}, \theta_k)$ may be non-monotone in either d_1 or d_2 . While it may seem self-evident, it also is important to note that defining each outcome as an ordinal variable having three or more levels and recording two outcome variables provides a much richer characterization of patient outcome than would be obtained by reducing the available clinical information to one or two binary variables.

Given the marginal probabilities of the two outcomes, to obtain a joint distribution for $\mathbf{Y}$ that is sufficiently tractable to facilitate practical application, we combine the marginals using a Gaussian copula. Let Φ_ρ denote the cumulative distribution function (cdf) of a bivariate standard normal distribution with correlation ρ , and let unsubscripted Φ denote the univariate standard normal cdf. A Gaussian copula is given by

$$C_\rho(u, v) = \Phi_\rho\{\Phi^{-1}(u), \Phi^{-1}(v)\} \quad \text{for } 0 \leq u, v \leq 1. \quad (7)$$

To apply this structure, we first note that the cdf of each three-level ordinal outcome having support $\{0, 1, 2\}$ is

$$F_k(y|\mathbf{d}, \theta_k) = \begin{cases} 0 & \text{if } y < 0 \\ 1 - \pi_{k,1}(\mathbf{d}, \theta_k) - \pi_{k,2}(\mathbf{d}, \theta_k) & \text{if } 0 \leq y < 1 \\ 1 - \pi_{k,2}(\mathbf{d}, \theta_k) & \text{if } 1 \leq y < 2 \\ 1 & \text{if } y \geq 2 \end{cases}$$

which is computed by recalling that $\pi_{k,1}(\mathbf{d}, \theta_k) + \pi_{k,2}(\mathbf{d}, \theta_k) = \xi_{k,1}^*(\mathbf{d}, \theta_k)$ and $\pi_{k,2}(\mathbf{d}, \theta_k) = \xi_{k,1}^*(\mathbf{d}, \theta_k) - \xi_{k,2}^*(\mathbf{d}, \theta_k)$. Since the support of each marginal is $\{0, 1, 2\}$, the joint pmf of $\mathbf{Y}$ can be expressed as

$$\pi(\mathbf{y}|\mathbf{d}, \theta) = \sum_{a=1}^2 \sum_{b=1}^2 (-1)^{a+b} C_\rho(u_a, v_b) \quad (8)$$

where $u_1 = F_1(y_1 | \mathbf{d}, \theta)$, $v_1 = F_2(y_2 | \mathbf{d}, \theta)$, $u_2 = F_1(y_1 - 1 | \mathbf{d}, \theta)$, and $v_2 = F_2(y_2 - 1 | \mathbf{d}, \theta)$ with u_2, v_2 the left-hand limits of F_k at y_k . Since C is parameterized by ρ , the model parameter vector is $\theta = (\theta_1, \theta_2, \rho)$, and $\dim(\theta) = 21$. Suppressing the arguments $(\mathbf{d}, \theta)$ for brevity, expanding (8) gives

$$\begin{aligned} \pi(\mathbf{y}) = & \Phi_\rho\{\Phi^{-1} \circ F_1(y_1), \Phi^{-1} \circ F_2(y_2)\} \\ & - \Phi_\rho\{\Phi^{-1} \circ F_1(y_1 - 1), \Phi^{-1} \circ F_2(y_2)\} \\ & - \Phi_\rho\{\Phi^{-1} \circ F_1(y_1), \Phi^{-1} \circ F_2(y_2 - 1)\} \\ & + \Phi_\rho\{\Phi^{-1} \circ F_1(y_1 - 1), \Phi^{-1} \circ F_2(y_2 - 1)\}. \end{aligned} \quad (9)$$

Thus, all bivariate probabilities can be derived from the above expression by applying Φ_ρ to the marginal distributions. Bivariate probabilities for which one or both of the y_k 's equal 0 are

slightly simpler due to the facts that $F_k(-1) = 0$, $\Phi^{-1}(0) = -\infty$ and $\Phi_\rho(x, y) = 0$ if either of its arguments is $-\infty$. For example, $\pi(0, 0) = C_\rho(\pi_{1,0}, \pi_{2,0}) = C_\rho(1 - \pi_{1,1} - \pi_{1,2}, 1 - \pi_{2,1} - \pi_{2,2})$ and, similarly, $\pi(1, 0) = C_\rho(1 - \pi_{1,2}, 1 - \pi_{2,1} - \pi_{2,2}) - \pi(0, 0)$.

3.3 Utilities and Dose Pair Selection

Because our method uses the posterior mean expected utility as a criterion for choosing dose pairs, a key assumption is that the utility function provides a meaningful one-dimensional numerical summary representation of the patient's complex multivariate outcome. This approach is especially useful in settings where some values of the outcome vector may be missing, in particular when efficacy is invaluable due to excessive toxicity. This outcome is not unlikely in the bladder cancer trial, and more generally it occurs quite commonly in oncology trials where the aim is to evaluate both efficacy and toxicity.

Let $U(\mathbf{y})$ denote the numerical utility of outcome $\mathbf{y}$. Given the parameter vector θ , we define the *mean utility for a patient treated with $\mathbf{d}$* to be the mean of $U(\mathbf{Y})$ over the possible outcomes,

$$u(\mathbf{d}, \theta) = E_{\mathbf{Y}}\{U(\mathbf{Y})|\mathbf{d}, \theta\} = \sum_{\mathbf{y}} U(\mathbf{y})\pi(\mathbf{y}|\mathbf{d}, \theta). \quad (10)$$

Given the current data from n patients at any point in the trial, $data_n = \{(\mathbf{Y}_1, \mathbf{d}_1), \dots, (\mathbf{Y}_n, \mathbf{d}_n)\}$, the method selects the dose pair $\mathbf{d}^{opt}(data_n)$ that maximizes the posterior mean over θ of $u(\mathbf{d}, \theta)$ given by (10). Formally, denoting the set of all dose pairs being considered by $\mathcal{D}$,

$$\mathbf{d}^{opt}(data_n) = \arg\max_{\mathbf{d} \in \mathcal{D}} E_\theta\{u(\mathbf{d}, \theta)|data_n\} = \arg\max_{\mathbf{d} \in \mathcal{D}} \sum_{\mathbf{y}} U(\mathbf{y}) E_\theta\{\pi(\mathbf{y}|\mathbf{d}, \theta)|data_n\}, \quad (11)$$

with the second equality obtained by reversing the expectation operators E_θ and $E_{\mathbf{Y}}$.

3.4 Additional Safety Rules

Although the utility function accounts for toxicity, there is the possibility that even the $\mathbf{d}$ having maximum utility will be too toxic. While this might be considered unlikely, it may arise due to unanticipated interactions between the two agents. To guard against this, we include the following two safety rules. The first rule constrains escalation, at any point in the trial, so that untried levels of each agent may not be skipped. Formally, if (d_1, d_2) is the current dose pair, then escalation is allowed to as yet untried pairs $(d_1 + 1, d_2)$, $(d_1, d_2 + 1)$, or $(d_1 + 1, d_2 + 1)$, provided that these are in the matrix of dose pairs being studied. Thus, for example, after the first cohort in the bladder cancer trial is treated at (2,2), allowable higher dose pairs for the second cohort are (3,2), (2,3) and (3,3) but not (4,2) or (4,3). There is no similar constraint on de-escalation.

To define a formal stopping rule, let $\pi_{1,2}^{max}$ denote a fixed upper limit on the probability of the most severe level of toxicity. This must be specified by the physicians planning the trial. Let p_U be a fixed upper probability cut-off, e.g. in the range .80 to .99. The trial will be stopped early if

$$\min_{\mathbf{d}} \Pr\{\pi_{1,2}(\mathbf{d}, \theta) > \pi_{1,2}^{max} | data\} > p_U. \quad (12)$$

This rule stops the trial if, based on the current data, there is a high posterior probability for all dose pairs that the most severe level of toxicity exceeds its pre-specified limit. This is similar to the safety monitoring rules used by Thall and Cook (2004) and Braun et al. (2007). We formulate the safety rule (12) to require that all dose pairs are too toxic, rather than only the pair $\mathbf{d}^{min}$ where both doses are at their lowest level, in order to control the false negative rate, that is, the probability of incorrectly stopping the trial when at least one dose pair is safe.

3.5 Computational Methods

The main technical difficulty in obtaining the dose selection criterion (11) is computing the posterior means $E_{\theta}\{\pi(\mathbf{y} | \mathbf{d}, \theta) | data_n\}$ for all dose pairs $\mathbf{d}$ given each interim data set, since then averaging over $U(\mathbf{y})$ and computing the maximum over the 12 dose pairs to obtain $\mathbf{d}^{opt}(data_n)$ is trivial. We used MCMC with Gibbs sampling (Robert and Cassella, 1999) to compute all posterior quantities. We integrated $\pi(\mathbf{y} | \mathbf{d}, \theta)$ over θ by generating a series $\theta^{(1)}, \dots, \theta^{(N)}$ distributed proportionally to the posterior integrand, using the two-level algorithm described in the Appendix of Braun et al. (2007). MCMC convergence was confirmed by a small ratio ($< 3\%$) of the Monte Carlo standard error (MCSE) to the standard deviation of the utilities of the four corner dose pairs. For each $\theta^{(i)}$, we first computed all linear terms

$\eta_{k,y}^{(j)}(\theta^{(i)})$ associated with each agent for each dose pair. For each outcome $k = 1, 2$ and level $y = 1, 2$, we then evaluated $\xi^*(\eta_{k,y}^{(1)}, \eta_{k,y}^{(2)}, \lambda_k, \gamma_k)$ as given by (3) at $\theta^{(i)}$, and applied the formula (6) to obtain $\pi_{k,1}(\mathbf{d}, \theta_k^{(i)})$ and $\pi_{k,2}(\mathbf{d}, \theta_k^{(i)})$. We then computed the Gaussian copula values (7) to obtain the joint pmf values $\pi(\mathbf{y} | \mathbf{d}, \theta^{(i)})$ given by (8) and averaged over the $\theta^{(i)}$'s to obtain their posterior means.

4 Application

4.1 Establishing Priors

To establish priors, we elicited information from the principal investigator of the bladder cancer trial, a co-author of this paper (N.H.). The elicited values consisted of a variety of prior mean outcome probabilities. Toxicity is scored as $Y_1 = 0$ if no grade 3,4 toxicity occurs, $Y_1 = 1$ if grade 3,4 toxicity occurs but is resolved within 2 weeks, and $Y_1 = 2$ if grade 3,4 toxicity occurs and is not resolved within 2 weeks. Efficacy is scored as $Y_2 = 0$ if PD occurs, with $Y_2 = 1$ if the patient has SD and $Y_2 = 2$ if the patient has PR/CR at the end of two cycles. We allow the possibility that Y_2 cannot be scored due to excessive toxicity, an outcome which has non-trivial probability in dose-finding trials. Unconditional prior means of $\pi_{k,1}(\mathbf{d}, \theta)$ and $\pi_{k,2}(\mathbf{d}, \theta)$ were elicited for each k and dose pair $\mathbf{d}$. The elicited values are summarized in Table 1. We used an extension of the least squares method of Thall and Cook (2004) to solve for prior means of the 20 model parameters (θ_1, θ_2) . We calibrated the prior variances to obtain a suitably non-informative prior to ensure that the data will dominate the decisions early in the trial. This was done in terms of the effective sample size (ESS) of the prior of each $\pi_{k,y}(\mathbf{d}, \theta_k)$ by matching its mean and variance with those of a beta(a, b), assuming $E\{\pi_{k,y}(\mathbf{d}, \theta_k)\} \approx a/(a+b)$ and $\text{var}\{\pi_{k,y}(\mathbf{d}, \theta_k)\} \approx ab/(a+b)^2(a+b+1)$, solving for the approximate ESS as $a+b \approx E\{\pi_{k,y}(\mathbf{d}, \theta_k)\}[1 - E\{\pi_{k,y}(\mathbf{d}, \theta_k)\}]/\text{var}\{\pi_{k,y}(\mathbf{d}, \theta_k)\} - 1$ and calibrating the variances of the elements of θ to obtain suitably small ESS values. We assumed that the $\alpha_{k,y,r}^{(j)}$'s and γ_k 's were normally distributed, each λ_k was log normal. For the correlation, we assumed ρ was uniformly

distributed between -1 and $+1$. Setting $\text{var}(\alpha_{k,y,r}^{(j)})=10^2$ and $\text{var}\{\log(\lambda_k)\} = \text{var}(\gamma_k) = 1.5^2$ gave priors for all $\pi_{k,y}(\mathbf{d}, \theta)$'s having ESS values ranging from 0.81 to 2.86 with overall mean ESS 1.45. Thus, in terms of the outcome probabilities, the prior is uninformative. The prior means of each pair $\pi_{k,1}(\mathbf{d}, \theta)$ and $\pi_{k,2}(\mathbf{d}, \theta)$ are given along with their corresponding elicited values in Table 1.

4.2 Utilities

Utilities of all possible outcomes (Y_1, Y_2) are given in Table 2a. These were determined using the Delphi method, a tool for establishing consensus among experts (Dalkey, 1969; Brook, Chassin, Fink et al., 1986). It consists of a series of repeated interrogations, usually by means of questionnaires, to a group of individuals whose opinions or judgments are of interest. After the initial interrogation of each individual, each subsequent interrogation is accompanied by information regarding the preceding round of replies, usually presented anonymously. The individual is thus encouraged to reconsider and, if appropriate, to change his/her previous reply in light of the replies of other members of the group. After two or three rounds, the group position is determined by averaging. In applying this approach, 8 medical oncologists, all members of the Genitourinary French National Group, agreed to answer the questionnaire out of an initial set of 15 oncologists who were contacted. The questionnaire provided a table of the 10 possible outcomes, with accompanying definitions, and asked the oncologist to provide a numerical utility for each outcome on a scale of 0 to 100. After two rounds of the process described above a consensus among the 8 oncologists was reached. In order to assess the method's sensitivity to the numerical utility values, we also considered two other utilities. The individual utilities of N.H. are given in Table 2b, and hypothetical utilities that place comparatively greater value on achieving a tumor response (CR/PR), compared to SD or PD without toxicity, are given in Table 2c.

5 Computer Simulations

For the bladder cancer trial, the maximum sample size $N = 48$ was chosen based on accrual limitations and preliminary simulations examining the design's sensitivity to values of N in the range 36 to 60. Unless otherwise stated, each simulated trial was conducted using the starting dose pair (2,2), with dose pairs chosen to maximize the posterior mean utility subject to the safety rules described in Section 3.4, with the safety stopping rule applied using $\pi_{1,2}^{max} = .33$ and $p_U = .80$. For each simulation, a scenario was specified in terms of fixed values of all marginal probabilities of both outcomes at all doses, and the correlation parameter.

Specifically, a scenario was determined by assumed true probability values $\pi_{k,1}^{true}(\mathbf{d})$ and $\pi_{k,2}^{true}(\mathbf{d})$ for $k = 1, 2$ for each $\mathbf{d} \in \mathcal{D}$, and true correlation ρ^{true} . Under the Gaussian Copula, these determined the joint probabilities $\pi^{true}(\mathbf{y} | \mathbf{d})$ for each dose pair $\mathbf{d}$, which were used to simulate the outcome of each patient, given the dose assigned to that patient by the method. Scenario 1 corresponds to the elicited prior, and its fixed marginal probabilities $\pi_{k,j}^{true}(\mathbf{d})$ are given in the rows labeled "Elicited" in Table 1. In Scenario 2, the dose pairs $\mathbf{d} = (1,2), (2,2), (3,2)$ in the middle row of the dose pair matrix have the highest utilities, with $\mathbf{d} = (2,2)$ most desirable. Scenario 3 has lower toxicity and higher synergistic effects between the two agents compared to Scenario 1. Scenario 4 has the same toxicities as Scenario 1, but with antagonistic effects between the two agents on the response probabilities $\pi_{2,y}^{true}$. Scenario 5 is a case where none of the dose pairs are safe, with all severe toxicity probabilities $\pi_{1,2}^{true}(\mathbf{d})$ varying from .60 to .75. The values of $\pi_{k,j}^{true}(\mathbf{d})$ for all scenarios are given in Supplementary Tables 1 – 5, along with the corresponding true utilities.

In order to evaluate how well the method performs in picking a dose pair to optimize the utility function, we define the following summary statistic. Recall from (11) that $\mathbf{d}^{opt}(\mathbf{data}_n)$ is the dose pair selected by the method at an interim point in the trial, since it maximizes the posterior mean utility based on the current data from n patients. In each simulation, since the outcome probabilities $\pi^{true}(\mathbf{y} | \mathbf{d})$ at each dose pair are specified, using these in place of the unknown probabilities $\pi(\mathbf{y} | \mathbf{d}, \theta)$ in (10) gives the corresponding true utility of treating a patient with dose $\mathbf{d}$ under the assumed true distribution $\{\pi^{true}(\mathbf{y} | \mathbf{d})\}$,

$$u^{true}(\mathbf{d}) = \sum_{\mathbf{y}} U(\mathbf{y}) \pi^{true}(\mathbf{y}|\mathbf{d}). \quad (13)$$

While probabilistic averages of the form (13) are used routinely in statistics, it is important to bear in mind that, given a utility function $U(\mathbf{y})$, it is not obvious what numerical values of $u^{true}(\mathbf{d})$ will be obtained from an assumed simulation scenario $\{\pi^{true}(\mathbf{y}|\mathbf{d})\}$. Denoting the final selected dose pair based on the data from N patients by $\mathbf{d}_N^{opt}$, a summary statistic to evaluate the method's performance is

$$R = \frac{u^{true}(\mathbf{d}_N^{opt}) - u_{min}}{u_{max} - u_{min}}, \quad (14)$$

where $u_{max} = \max\{u^{true}(\mathbf{d}): \mathbf{d} \in \mathcal{D}\}$ is the maximum and $u_{min} = \min\{u^{true}(\mathbf{d}): \mathbf{d} \in \mathcal{D}\}$ is the minimum possible true utility over all dose pairs in $\mathcal{D}$. Thus, R is the proportion of $u_{max} - u_{min}$ achieved by the selected dose pair, with $R = 1$ if the dose pair having maximum possible true utility under $\{\pi^{true}(\mathbf{y}|\mathbf{d})\}$ is chosen.

All simulations were based on 1000 repetitions of each case. The first simulation study assumes the consensus utilities elicited using the Delphi method (Table 2a) and assumes true correlation $\rho^{true} = 0.10$. The results are summarized in Table 3. When interpreting these results, it is important to bear in mind that each scenario's true utilities are only summary values, and they may hide the complex structure of the bivariate model probabilities. Overall, the method performed well in all scenarios, as summarized by the R statistic in each case. In each of Scenarios 1–4, a dose pair having utility close to the maximum is very likely to be chosen. In Scenario 2, where the utility surface is an inverted “U” shape with largest $u^{true}(\mathbf{d})$ values in the middle of the dose domain at $(d_1, d_2) = (1,2), (2,2)$ and $(3,2)$, the method correctly selects one of these pairs with high probability, and the best pair $(2,2)$ with highest probability. In Scenario 4, the method is very likely to correctly select a high utility dose pair with $d_1 = 1$, in the left column. The fact that the method has high probabilities of choosing dose pairs with high utilities under each of the scenarios, 1–4, where the pairs with highest utilities have very different locations in the dose domain for the different scenarios, shows that the data rather than the prior drive the method's decisions. In terms of safety, the method correctly stops the trial early and chooses no dose pair 85% of the time in Scenario 5, where all pairs were excessively toxic. The cut-off $p_U = .80$ for the safety rule was chosen after preliminary simulations using $p_U = .90$ stopped the trial early only 64% of the time under Scenario 5. Note that, in Scenario 5 where all 12 dose pairs are excessively toxic and the design stops the trial early 85% of the time, the value $R = 0.67$ corresponds to the 15% of cases where the trial was not stopped.

The aim of the second set of simulations was to assess the method's sensitivity to the numerical utilities. Table 4 summarizes the results for each of the three utilities given in Table 2 under each scenario. The R statistic appears to be quite sensitive to the utility chosen for trial conduct. This is a highly desirable property of the method, since if it were not the case then choosing $\mathbf{d}$ to optimize the posterior mean utility would be pointless.

The R statistic is insensitive to using $c = 1$ versus $c = 3$ under all scenarios. However, $c = 1$ is slightly safer under Scenario 5 with early stopping probability 92% compared to 85% for $c = 3$. This is not surprising, since it is well known that continuous monitoring provides the safest design when using outcome-adaptive methods, although it must be borne in mind that using

$c = 1$ may not be logistically feasible in practice. For cohort size $c = 3$, we evaluated a two-stage version of the design with either the first $n_1 = 12$ or 24 patients assigned dose pairs with the chemo agent dose fixed at 1000 mg/m^2 ($d_2 = 2$) and only the biological agent dose varied, and dose pairs chosen from the entire set of 12 for the remaining $48 - n_1$ patients. This has a negligible effect on the values of R for all scenarios. These simulations are summarized in Supplementary Table 6.

The R statistic increases with N_{max} in all scenarios, with small increases in R as N_{max} increases from 36 to 60 and, as a numerical check of consistency, much larger R values for the somewhat unrealistic sample size $N_{max} = 240$. To assess sensitivity of the method to the prior, if a naive prior is used with all prior means 0 and prior standard deviations $\sigma(\alpha) = \sigma\{\log(\lambda)\} = \sigma(\gamma) = 1000$, the R statistic drops substantially under scenarios 1, 2 and 3, but increases by .11 under scenario 4. These simulations are summarized in Supplementary Table 7. We also varied the prior variances, studying several configurations of the standard deviations $\sigma(\alpha_{k,y,r}^{(j)})$ and $\sigma\{\log(\lambda_k)\} = \sigma(\gamma_k)$, summarized in Supplementary Table 8. In terms of R and % no dose selected, the method is insensitive to varying $\sigma(\alpha_{k,y,r}^{(j)})$ from 7 to 15 and $\sigma\{\log(\lambda_k)\} = \sigma(\gamma_k)$ from 1 to 3, with the exception that, under Scenario 4, $R = .68$ for $\sigma(\alpha_{k,y,r}^{(j)}) = 7$ and $\sigma\{\log(\lambda_k)\} = \sigma(\gamma_k) = 1.5$, compared to $R = .78 - .82$ for the other cases. These simulations indicate that it is important to use elicited values to establish prior means and carefully calibrate the prior variances, rather than choosing arbitrary hyperparameter values that appear to be non-informative.

We also studied the method's sensitivity to $\rho^{true} = 0$ to 0.90 (not tabled), and found that there was little or no effect on the design's performance. A surprising result is that, if we assume independence between Y_1 and Y_2 by setting $\rho \equiv 0$, there is virtually no degradation of the design's performance for $0 \leq \rho^{true} \leq .50$. However, for $.60 \leq \rho^{true} \leq .90$, assuming independence gives a design with slightly smaller stopping probability under scenario 5 where all dose pairs are excessively toxic. Finally, Supplementary Table 9 summarizes sensitivity to starting at $\mathbf{d} = (1,1)$ versus $(2,2)$, $(1,2)$ or $(2,1)$. The method is insensitive to the starting pair under scenarios 3 and 5, but under scenario 2 starting at $(1,1)$ gives $R = .66$ versus $R = .81$ when starting at $(2,2)$, whereas under the "antagonistic dose effect" scenario 4 starting at $(1,1)$ gives $R = .94$ versus $R = .78$ when starting at $(2,2)$. So there seems to be a trade-off with regard to choice of starting dose pair.

An Associate Editor has asked the interesting question of how much improvement is obtained by using the data, compared to relying on the prior to choose an optimal $\mathbf{d}$. The prior mean utilities of the 12 dose pairs vary from minimum $u^{prior}(1, 1) = 47.7$ to maximum $u^{prior}(4, 2) = 51.5$. If $\mathbf{d} = (4,2)$ were chosen on this basis without conducting a trial, then the respective values of R for the scenarios 1, 2, 3, and 4 in Table 3 having acceptably safe dose pairs would be 1.00, 0.16, 0.94, and 0.24, compared to the values 0.71, 0.81, 0.85, and 0.78 obtained from running the trial. Thus, essentially, using the prior mean utilities to choose $\mathbf{d}$ would rely on making a lucky guess, since $\mathbf{d} = (4,2)$ would be a very good choice in scenarios 1 and 3, a very poor choice in scenarios 2 and 4, and a disaster in scenario 5 where all $\mathbf{d}$ pairs are unacceptably toxic.

6 Discussion

We have proposed an outcome-adaptive Bayesian design for the common clinical problem of choosing a dose pair of a chemotherapeutic and biologic agent used in combination. While characterizing patient outcome as a vector of two ordinal variables provided a very informative characterization of treatment effect, developing a tractable model and conducting computer simulations were both quite complex and time consuming. However, the computations required for actual trial conduct are quite feasible. The use of numerical utilities that characterize the

desirability of the outcomes yielded a design with very attractive properties. This may be contrasted with the very different, more conventional approach of the continual reassessment method and similar designs, which use only toxicity and choose a dose having posterior mean toxicity probability closest to a given fixed target value. As noted earlier, another advantage of using utilities is that one may account formally for the not uncommon outcome that efficacy is invaluable due to severe toxicity.

In our design patients are not randomized among doses, but rather a single dose pair is chosen for each new cohort to optimize a statistical decision criterion, in our case posterior mean utilities. This common practice in phase I and phase I/II dose-finding trials is motivated by safety concerns, since they are the first trials to study a new agent or combination in humans and higher doses, or higher dose pairs, carry a higher risk of severe toxicity. Thus, doses are chosen in a sequential, outcome adaptive manner primarily due to ethical concerns. However, our design could be modified so that, once an acceptable level of safety has been established for all dose pairs, if two or more pairs have posterior mean utilities that are numerically very close to each other, then the patients in the next cohort could be randomized among these dose pairs. This might spread the sample more evenly over the two-dimensional dose domain, and thus provide improved posterior estimates of $\pi(y|\mathbf{d}, \theta)$ and $u(\mathbf{d}, \theta)$ as functions of $\mathbf{d}$ without sacrificing safety.

The design stops early if all $\mathbf{d}$ are excessively toxic, but not if all $\mathbf{d}$ have a low response rate. In the latter case the method is likely to run the trial to completion and choose the pair $\mathbf{d}$ having highest utility. If all response rates are about equal, this would correspond to choosing the pair having lowest toxicity. The design could be modified to include a “phase II” type rule, similar to (12), to stop the trial if all $\mathbf{d}$ have low response rates, as is done in the phase I/II design of Thall and Cook (2004). In the case where Y_2 is trinary, such a rule could be formulated in terms of either $\pi_{2,2}(\mathbf{d})$ or possibly $\pi_{2,1}(\mathbf{d}) + \pi_{2,2}(\mathbf{d})$.

A menu-driven computer program, named “U2OET”, to implement this methodology is available from the website <http://biostatistics.mdanderson.org/SoftwareDownload>.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Aranda-Ordaz FJ. On two families of transformations to additivity for binary response data. *Biometrika* 1983;68:357–363.
- Babb JS, Rogatko A, Zacks S. Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine* 1998;17:1103–1120. [PubMed: 9618772]
- Bekele BN, Shen Y. A Bayesian approach to jointly modeling toxicity and biomarker expression in a phase I/II dose-finding trial. *Biometrics* 2004;60:343–354.
- Bekele BN, Thall PF. Dose-finding based on multiple toxicities in a soft tissue sarcoma trial. *Journal of the American Statistical Association* 2004;99:26–35.
- Braun T. The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials* 2002;23:240–256. [PubMed: 12057877]
- Braun TM, Thall PF, Nguyen H, de Lima M. Simultaneously optimizing dose and schedule of a new cytotoxic agent. *Clinical Trials* 2007;4:113–124. [PubMed: 17456511]
- Brook RH, Chassin MR, Fink A, Solomon DH, Kosecoff J, Park RE. A method for the detailed assessment of the appropriateness of medical technologies. *International Journal of Technology Assessment and Health Care* 1986;2:53–63.
- Cheung YK, Chappell R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* 2000;56:1177–1182. [PubMed: 11129476]

- Dalkey NC. An experimental study of group opinion. *Futures* 1969;1:408–26.
- Ivanova A. A new dose-finding design for bivariate outcomes. *Biometrics* 2003;59:1001–1007. [PubMed: 14969479]
- Ivanova A, Wang K. Bivariate isotonic design for dose-finding with ordered groups. *Biometrics* 2006;25:2018–2026.
- Joe, H. *Multivariate Models and Dependence Concepts*. New York: Chapman and Hall; 1997.
- Korn EL, Simon RM. Using the tolerable-dose diagram in the design of phase I combination chemotherapy trials. *Journal of Clinical Oncology* 1993;11:794–801. [PubMed: 8478673]
- Mandrekar SJ, Cui Y, Sargent DJ. An adaptive phase I design for identifying a biologically optimal dose for dual agent drug combinations. *Statistics in Medicine* 2007;26:2317–2320. [PubMed: 17016867]
- O’Quigley J, Hughes MD, Fenton T. Dose-finding designs for HIV studies. *Biometrics* 2001;57:1018–1029. [PubMed: 11764240]
- O’Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* 1990;46:33–48. [PubMed: 2350571]
- Robert, CP.; Cassella, G. *Monte Carlo Statistical Methods*. New York: Springer; 1999.
- Spiegelhalter DJ, Best NG, Carlin BP, van der Linde A. Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society, Series B* 2002;64:583–639.
- Storer BE. Design and analysis of phase I clinical trials. *Biometrics* 1989;45:925–937. [PubMed: 2790129]
- Therasse P, Arbuck SG, Eisenhauer EA, Wanders J, Kaplan RS, Rubinstein L, Verweij J, Van Glabbeke M, van Oosterom AT, Christian MC, Gwyther SG. New guidelines to evaluate the response to treatment in solid tumors. *Journal of the National Cancer Institute* 2000;92:205–216. [PubMed: 10655437]
- Thall PF, Cook John D. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* 2004;60:684–693. [PubMed: 15339291]
- Thall PF, Cook John D, Estey EH. Adaptive dose selection using efficacy-toxicity trade-offs: illustrations and practical considerations. *Journal of Biopharmaceutical Statistics* 2006;16:623–638. [PubMed: 17037262]
- Thall PF, Millikan RE, Mueller P, Lee SJ. Dose-finding with two agents in phase I oncology trials. *Biometrics* 2003;59:487–496. [PubMed: 14601749]
- Thall PF, Nguyen H, Estey EH. Patient-specific dose-finding based on bivariate outcomes and covariates. *Biometrics* 2008;64:1126–1136. [PubMed: 18355387]
- Thall PF, Russell KT. A strategy for dose finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* 1998;54:251–264. [PubMed: 9544520]
- Yuan Z, Chappell R, Bailey H. The continual reassessment method for multiple toxicity grades: a Bayesian quasi-likelihood approach. *Biometrics* 2007;63:173–179. [PubMed: 17447942]
- Zohar S, Chevret S. Recent developments in adaptive designs for phase I/II dose-finding studies. *Journal of Biopharmaceutical Statistics* 2008;17:1071–1083. [PubMed: 18027217]

Scenario 1.

Scenario 2.

Scenario 3.

Scenario 4.

Elicited and computed prior means of the probabilities $\pi_{k,1}(\mathbf{d})$, $\pi_{k,2}(\mathbf{d})$ for each dose pair $\mathbf{d}$ and outcome Y_k , where $k = 1$ for toxicity and $k = 2$ for efficacy.

Table 1

d_2 k		d_1 = Dose Level of Biological Agent					
		0	1	2	3	4	
3	1	<i>Elicited</i>	–	.85, .10	.85, .11	.85, .12	.82, .15
		<i>Computed</i>	–	.58, .19	.59, .18	.59, .19	.56, .22
2		<i>Elicited</i>	.50, .33	.50, .33	.50, .33	.55, .35	.60, .38
		<i>Computed</i>	–	.37, .36	.39, .35	.41, .36	.41, .37
2	1	<i>Elicited</i>	–	.80, .04	.80, .05	.80, .06	.80, .08
		<i>Computed</i>	–	.59, .15	.61, .13	.60, .14	.57, .18
2		<i>Elicited</i>	.50, .33	.50, .33	.50, .33	.55, .35	.60, .38
		<i>Computed</i>	–	.37, .34	.39, .34	.40, .35	.41, .36
1	1	<i>Elicited</i>	–	.70, .02	.70, .03	.70, .04	.70, .06
		<i>Computed</i>	–	.56, .14	.57, .13	.57, .14	.54, .17
2		<i>Elicited</i>	.45, .25	.45, .25	.45, .25	.50, .27	.55, .30
		<i>Computed</i>	–	.35, .32	.37, .32	.39, .33	.39, .35

Table 2

Utilities for patient outcomes in the bladder cancer trial. For toxicity, $Y_1 = 0$ if no grade 3,4 toxicity occurs, $Y_1 = 1$ if grade 3,4 toxicity occurs but is resolved within 2 weeks, and $Y_1 = 2$ if grade 3,4 toxicity occurs and is not resolved within 2 weeks.

	$Y_2 = 0$	$Y_2 = 1$	$Y_2 = 2$	Y_2
	<u>Progressive Disease</u>	<u>Stable Disease</u>	<u>Tumor Response</u>	<u>Inevaluable</u>
<i>a. Utilities Elicited Using Delphi Method</i>				
$Y_1 = 0$	25	76	100	–
$Y_1 = 1$	10	60	82	–
$Y_1 = 2$	2	40	52	0
<i>b. Utilities Reflecting One Oncologist</i>				
$Y_1 = 0$	40	90	100	–
$Y_1 = 1$	10	50	80	–
$Y_1 = 2$	0	40	40	0
<i>c. Hypothetical Utilities Reflecting Greater Value on Tumor Response</i>				
$Y_1 = 0$	10	50	100	–
$Y_1 = 1$	5	40	90	–
$Y_1 = 2$	0	30	80	0

Table 3

Simulation results for the design based on the elicited utilities A given in Table 2.

Scenario	d_2	d_1					R	% none
		1	2	3	4			
1	3	$u^{true}(d)$	57.3	56.9	60.4	64.2	0.71	0
		% Sel, # Pats.	1, 0.7	2, 0.7	6, 2.7	14, 5.1		
2	2	$u^{true}(d)$	60.5	60.1	63.6	67.4		
		% Sel, # Pats.	4, 1.4	4, 6.9	14, 8.9	29, 12.0		
1	1	$u^{true}(d)$	54.6	54.2	57.8	61.7		
		% Sel, # Pats.	1, 0.3	0, 0.3	2, 1.0	22, 8.0		
3	3	$u^{true}(d)$	47.2	46.2	44.2	42.2	0.81	1
		% Sel, # Pats.	2, 0.8	1, 0.8	2, 0.9	1, 0.5		
2	2	$u^{true}(d)$	54.6	65.8	61.8	43.4		
		% Sel, # Pats.	8, 2.7	44, 21.9	34, 15.2	1, 1.6		
1	1	$u^{true}(d)$	44.6	44.9	45.1	39.1		
		% Sel, # Pats.	3, 1.0	0, 0.5	2, 1.1	1, 0.9		
3	3	$u^{true}(d)$	48.6	53.1	59.6	64.9	0.85	0
		% Sel, # Pats.	1, 0.7	2, 0.7	9, 3.2	23, 9.4		
2	2	$u^{true}(d)$	43.9	50.8	55.2	63.5		
		% Sel, # Pats.	0, 0.2	1, 4.1	4, 4.9	35, 14.1		
1	1	$u^{true}(d)$	41.4	48.9	53.4	58.8		
		% Sel, # Pats.	0, 0.1	0, 0.1	1, 0.8	25, 9.8		
4	3	$u^{true}(d)$	61.2	51.0	38.0	26.0	0.78	0
		% Sel, # Pats.	12, 3.5	2, 1.0	1, 1.7	1, 3.0		
2	2	$u^{true}(d)$	65.4	53.6	45.3	37.3		
		% Sel, # Pats.	19, 5.6	2, 5.8	4, 6.7	4, 5.7		
1	1	$u^{true}(d)$	73.9	61.9	55.2	44.9		
		% Sel, # Pats.	44, 9.8	0, 0.3	1, 0.8	9, 4.1		
5	3	$u^{true}(d)$	44.6	44.2	46.8	49.6	0.67	85
		% Sel, # Pats.	1, 1.6	1, 1.3	3, 2.4	3, 4.2		

Scenario	d_2	d_1				R	% none
		1	2	3	4		
	2						
		$u^{true}(d)$	45.6	45.0	47.7	50.0	
		% Sel, # Pats.	0, 0.5	1, 3.8	1, 2.0	1, 0.8	
	1						
		$u^{true}(d)$	39.7	39.0	41.7	44.0	
		% Sel, # Pats.	0, 0.7	0, 0.7	1, 2.8	3, 4.7	

Table 4

Comparison of the summary *R* statistics and stopping percentages for conducting trials under the alternative utility definitions given in Table 2.

Scenario	Utility A		Utility B		Utility C	
	R	% none	R	% none	R	% none
1	0.71	0	0.57	0	0.63	0
2	0.81	1	0.68	1	0.45	1
3	0.85	0	0.77	0	0.86	0
4	0.78	0	0.83	0	0.79	0
5	0.67	85	0.66	89	0.80	61