Published in final edited form as:

J Biopharm Stat. 2012 July ; 22(4): 785–801. doi:10.1080/10543406.2012.676586.

Adaptive Randomization to Improve Utility-Based Dose-Finding with Bivariate Ordinal Outcomes

Peter F. Thall and Hoang Q. Nguyen

Department of Biostatistics, University of Texas M.D. Anderson Cancer Center, Houston, Texas
USA 77030

Peter F. Thall: rex@mdanderson.org

Summary

A sequentially outcome-adaptive Bayesian design is proposed for choosing the dose of an experimental therapy based on elicited utilities of a bivariate ordinal (toxicity, efficacy) outcome. Subject to posterior acceptability criteria to control the risk of severe toxicity and exclude unpromising doses, patients are randomized adaptively among the doses having posterior mean utilities near the maximum. The utility increment used to define near-optimality is non-increasing with sample size. The adaptive randomization uses each dose's posterior probability of a set of good outcomes, defined by a lower utility cut-off. Saturated parametric models are assumed for the marginal dose-toxicity and dose-efficacy distributions, allowing the possible requirement of monotonicity in dose, and a copula is used to obtain a joint distribution. Prior means are computed by simulation using elicited outcome probabilities, and prior variances are calibrated to control prior effective sample size and obtain a design with good operating characteristics. The method is illustrated by a phase I/II trial of radiation therapy for children with brain stem gliomas.

Keywords

Adaptive design; Bayesian design; Clinical trial; Dose-finding; Epsilon-greedy algorithm; Phase I/II clinical trial; Utility

1. Introduction

We propose a Bayesian phase I/II procedure for sequentially adaptive dose selection based on a bivariate ordinal (toxicity, efficacy) outcome. The method is based on elicited utilities of the possible outcome pairs (cf. Berger, 1985). Rather than choosing the dose that maximizes the posterior mean utility, we deal with the well-known “exploration versus exploitation” dilemma (cf. Sutton and Barto, 1998) by adaptively randomizing patients among the doses having posterior mean utilities that differ from the maximum by less than a specified increment. We require this increment to be non-increasing with sample size, similarly to an “epsilon decreasing” version of an “epsilon greedy” algorithm. The adaptive randomization (AR) uses each dose's posterior probability of a set of good outcomes, defined by an elicited lower utility cut-off. We obtain a bivariate model by first constructing marginals and using a copula (Nelsen, 1999) to induce association. Aside from link functions, we do not assume functional forms for dose-toxicity or dose-efficacy curves, but rather use marginals with saturated parameterizations. We establish a prior by using elicited outcome probabilities to simulate many very large pseudo samples, with the mean parameter

vector of the pseudo posteriors used as the prior mean parameter vector, and the prior variances calibrated to control prior effective sample size.

Because we consider bivariate ordinal outcomes, our methodology differs substantively from phase I/II designs based on trinary or bivariate binary outcomes (cf. Thall and Russell, 1998; O'Quigley, Hughes, and Fenton, 2001; Braun, 2002; Thall and Cook, 2004; Bekele and Shen, 2005; Zhang, Sargent and Mandrek, 2005; Dragalin and Fedorov, 2006; Thall, Nguyen and Estey, 2008). Our design may be considered a phase I/II, utility-based generalization of several Bayesian phase I methods. Bekele and Thall (2004) use a continual reassessment method (CRM, O'Quigley et al., 1990) type criterion based on posterior means of summed severity scores of multivariate ordinal toxicities; however, their methodology does not incorporate efficacy, maximize a utility, or use AR. For the case of one ordinal toxicity, a method similar to that of Bekele and Thall, but using quasi-likelihood, is proposed by Yuan, Chappell and Bailey (2007), while Van Meter et al. (2011) assume a proportional odds model (McCullagh, 1980) and extend the CRM. Our methodology is similar to the phase I/II design of Houede et al. (2010) for choosing dose pairs of two agents, with the differences that we incorporate AR, require more complex dose admissibility criteria, and use a model with much weaker assumptions.

In Section 2, we describe the motivating trial. The probability model is presented in Section 3. Section 4 gives definitions of the utility function, dose admissibility criteria, AR probabilities, and the design. Application to the RT trial is described in Section 5, including simulation studies of the method's sensitivity to dose admissibility requirements, the use of AR, maximum sample size, prior variability, and number of doses studied. We close with a discussion in Section 6.

2. A Radiation Therapy Trial

Diffuse intrinsic pontine gliomas (DIPGs) are aggressive brain tumors for which no treatment with substantive anti-disease effect currently exists. DIPGs account for about 75% of brain stem gliomas in children, and the median age of DIPG patients is 5 years. Radiation therapy (RT) is the standard treatment, but nearly all patients experience disease progression within eight months after RT and median survival is less than one year. However, the dose-toxicity and dose-efficacy profiles of RT for this disease are not well understood.

This paper was motivated by the desire to design a phase I/II trial of RT for DIPG patients. The trial, which uses the dose-finding design described here and is ongoing at this writing, includes children with DIPGs who previously received RT with or without chemotherapy and currently have progressive disease. RT is administered by separating a total dose of absorbed radiation (in gray units, Gy) into fractions that are given serially. The *biologically equivalent dose* is $BED = (\text{total dose}) \cdot (1 + d/\kappa)$, where $d = \text{dose/fraction}$ and κ is a constant corresponding to the type of tissue being irradiated. This model is based on the empirical observation in animal models that the empirical proportion of cells surviving radiation can be fit closely by a linear-quadratic function of per fraction dose (Fowler, 1989), with the assumption that cell killing of successive fractions is independent. For brain tissue, $\kappa = 3$, and the three combinations of (total dose, d) to be studied in the DIPG trial are (24, 2.0), (26.4, 2.2), and (30.8, 2.2), so the corresponding BEDs are 40.00, 45.76, and 53.39. Toxicity is defined on a 4-level ordinal scale as Low, Moderate, High, or Severe, with each level defined in terms of fatigue, nausea/vomiting, headache, skin inflammation or desquamation, blindness, and brain edema or necrosis, each evaluated during 42 days from the start of therapy. Efficacy is scored at day 42, and is defined as the sum of three indicators of any improvement, compared to baseline, in (i) clinical symptoms, (ii) radiographic appearance,

or (iii) quality of life. Thus, an efficacy score of 0 corresponds to no improvement, 1 to improvement in one of the three categories, and so on.

For our general regime, index the outcomes by $k = 1$ for toxicity and $k = 2$ for efficacy, with $Y_k = 0, 1, \dots, m_k$ identifying the observed ordinal levels. For toxicity, $Y_1 = 0$ denotes the least severe and m_1 the most severe level, while $Y_2 = 0$ denotes the worst and m_2 the best level of efficacy. In the RT trial, $m_1 = m_2 = 3$. Our proposed dose-finding method requires elicited utilities of all possible values $y = (y_1, y_2)$ of the outcome pairs $Y = (Y_1, Y_2)$. The elicited utilities $U(y_1, y_2)$ for all 16 possible outcomes in the RT trial are given in Table 1. The RT trial utility has the required property that $U(y_1, y_2)$ must increase as either toxicity severity level decreases or efficacy increases. While it may seem counterintuitive that outcomes with efficacy score $Y_2 = 0$ should be given positive utilities, because the prognosis of the patients in this trial is very poor and treatment is in large part palliative, the oncologists who organized the trial consider achieving a lower level of toxicity to be more desirable for any level of efficacy. When we questioned the oncologists about this critical point during the utility elicitation process, they explained that, even when the efficacy outcome score equals 0, there still is some palliative effect, that is, the treatment still is useful. Consequently, when the efficacy score is 0, lower toxicity severity levels are far more desirable, as the first row of Table 1 shows.

3. Probability Models

Index doses by $\mathcal{X} = \{1, \dots, J\}$. We will construct a model for $Y = (Y_1, Y_2)$ as a function of $x \in \mathcal{X}$ by formulating marginals for $[Y_1 | x]$ and $[Y_2 | x]$ and using a copula (Nelsen, 1999) to obtain a joint distribution. Let θ denote the model parameter vector, with θ_k the subvector characterizing the marginal of $[Y_k | x]$. Denote

$$\lambda_{k,y,x} = \Pr(Y_k \geq y | Y_k \geq y-1, x, \theta_k) \quad \text{and} \quad \pi_{k,y,x} = \Pr(Y_k = y | x, \theta_k)$$

for $y = 1, \dots, m_k$ and $k=1,2$. Given a monotone increasing link function g , our marginal model assumption is simply

$$g(\lambda_{k,y,x}) = \theta_{k,y,x}, \quad y=1, \dots, m_k, \quad k=1, 2, \quad (1)$$

with all $\theta_{k,y,x}$ real-valued. Thus, $\theta_k = (\theta_{k,1}, \dots, \theta_{k,m_k})$, denoting $\theta_{k,y} = (\theta_{k,y,1}, \dots, \theta_{k,y,J})$. The marginal model (1) is saturated since it has m_k parameters for each x , with $\dim(\theta_k) = Jm_k$, which is the number of $\pi_{k,y,x}$'s needed to specify the J marginals of Y_k for all x .

Equation (1) ensures that the distribution $\{\pi_{k,y,x}, y=0, 1, \dots, m_k\}$ is well defined for each k and x . This follows from the fact that, denoting $\bar{\pi}_{k,y,x} = \Pr(Y_k \leq y | x, \theta_k)$, the unconditional and conditional probabilities are related by the recursive formula

$$\bar{\pi}_{k,y,x} = \prod_{r=1}^y \lambda_{k,r,x} = \prod_{r=1}^y g^{-1}(\theta_{k,r,x}), \quad y=1, \dots, m_k. \quad (2)$$

Consequently, $\bar{\pi}_{k,y,x}$ is decreasing in y for each x , Jm_k -dimensional real-valued θ_k , and monotone link g . The marginal probabilities are given by

$$\begin{aligned}\pi_{k,y,x} &= (1 - \lambda_{k,y+1,x}) \prod_{r=1}^y \lambda_{k,r,x}, \quad y=1, \dots, m_k - 1, \\ \pi_{k,m_k,x} &= \prod_{r=1}^{m_k} \lambda_{k,r,x}.\end{aligned}\quad (3)$$

Given the marginals, we obtain a joint distribution of $[Y_1, Y_2 | x]$ by using a bivariate Gaussian copula, $C_\rho(v_1, v_2) = \Phi_\rho\{\Phi^{-1}(v_1), \Phi^{-1}(v_2)\}$ for $0 \leq v_1, v_2 \leq 1$, where Φ_ρ is the bivariate standard normal cdf with correlation ρ , and Φ is the univariate standard normal cdf. This is used to define the joint distribution as

$$\pi(y|x, \theta) = \Pr(Y=y|x, \theta) = \sum_{a=1}^2 \sum_{b=1}^2 (-1)^{a+b} C_\rho(u_{1,a}, u_{2,b})$$

where $u_{k,1} = \Pr(Y_k = y_k | x, \theta_k)$ and $u_{k,2} = \Pr(Y_k = y_k - 1 | x, \theta_k)$. We chose the Gaussian copula for its tractability. Other copulas, such as the Gumbel or Clayton, may be used (Nelsen, 1999). Denoting the data from the first n patients in the trial, by $D_n = \{(\mathbf{Y}_1, x_{[1]}), \dots, (\mathbf{Y}_n, x_{[n]})\}$ $n = 1, \dots, N_{\max}$, the likelihood is the product

$$L(D_n | \theta) = \prod_{i=1}^n \prod_{y_1=0}^{m_1} \prod_{y_2=0}^{m_2} \{\pi(y_1, y_2 | x_{[i]}, \theta)\}^{I(Y_{i,1}=y_1, Y_{i,2}=y_2)}$$

and the posterior is $p(\theta | D_n) \propto L(D_n | \theta) \text{prior}(\theta)$.

The model parameter vector $\theta = (\theta_1, \theta_2, \rho)$ has dimension $p = J(m_1 + m_2) + 1$, and characterizes a total of $J(m_1 + m_2 + m_1 m_2)$ bivariate probabilities. This parameterization is feasible for many cases arising in practice, with $(J, m_1, m_2) = (3, 2, 2), (3, 3, 3), (4, 3, 2), (4, 3, 3), (5, 2, 2), (5, 3, 2), (5, 3, 3)$ giving corresponding $p = 13, 19, 21, 25, 21, 26, 31$. For $m_1 = m_2 = 1$, which is the bivariate binary outcome case, $p = 2J + 1$ and $\theta_k = (\theta_{k,1}, \dots, \theta_{k,J})$, for $k = 1, 2$.

How $\bar{\pi}_{1,y,x}$ and $\bar{\pi}_{2,y,x}$ may vary with x depends on both the therapeutic modality and the definitions of toxicity and efficacy. An important case is that where it is necessary to assume that $\bar{\pi}_{k,y,x}$ is increasing in x for one or both outcomes. This assumption is appropriate both for cytotoxic agents and for RT, but it may not be realistic for cytostatic or biologic agents. Imposing the constraints $\theta_{k,y,1} \leq \theta_{k,y,2} \leq \dots \leq \theta_{k,y,J}$ implies that $\bar{\pi}_{k,y,x}$ increases in x , for each $y = 1, \dots, m_k$ due to the monotonicity of g , by equation (2). Rather than fitting the model (1) with real-valued $\theta_{k,y,x}$'s subject to these constraints, to reduce computation we obtain monotonicity of $\bar{\pi}_{k,y,x}$ in x by re-parameterizing the model as

$$\theta_{k,y,x} = \mu_{k,y} + \sum_{z=2}^x \gamma_{k,y,z} \quad \text{for all } x=2, \dots, J \quad (4)$$

with real-valued $\mu_{k,y} \equiv \theta_{k,y,1}$ and $\gamma_{k,y,x} \geq 0$ for all k, y , and $x = 2, \dots, J$. Thus, $\theta_{k,y} = (\mu_{k,y}, \gamma_{k,y,2}, \dots, \gamma_{k,y,J})$, and collecting terms we denote $\theta = (\mu, \gamma)$. The parameterization (4) borrows strength between doses quite strongly because, for distinct doses x and z , $\theta_{k,y,x}$ and $\theta_{k,y,z}$ share many parameters.

The model $\pi(y|x, \theta)$ must do a good job of reflecting the way that the posterior mean utility, defined below in Section 4.1, changes as a function of dose. Intuitively, it may seem that the number of parameters p in the range 13 to 31 is impractically large for dose-finding

trials with small sample sizes. This is not the case, essentially because the information in a bivariate ordinal outcome $\mathbf{Y}$ is much greater than that provided by a single ordinal Y or a binary outcome, which is used conventionally in phase I trials. Since our goal is to find a dose with high mean utility, if the model is tractable then the value of p is not critical. For example, in the case $(J, m_1, m_2) = (5, 3, 3)$ where $p = 31$, the algorithm (Section 4.3, below) for computing a prior from elicited values works quite well, implementing the MCMC algorithm for computing posteriors is not problematic, and the design performs well across a large set of different dose-outcome scenarios.

When using a utility $U(\mathbf{Y})$ to quantify the desirability of a bivariate ordinal patient outcome $\mathbf{Y}$, conventional generalized linear models (GLMs, McCullagh and Nelder, 1972) for the marginals may not be sufficiently refined to distinguish reliably between doses. This is especially problematic when a middle dose has the highest utility, which can easily be the case when both $\bar{\pi}_{1,y,x}$ and $\bar{\pi}_{2,y,x}$ increase with x . The family of GLMs for ordinal Y_k given by $g(\pi_{k,y,x}) = \alpha_{k,y} + \beta_k x$ with $\alpha_{k,y}$ decreasing in y , which is the proportional odds model if g is the logit link, may not be sufficiently flexible for dose-finding because a single dose effect β_k is assumed for all levels of Y_k . The more general form $g(\lambda_{k,y,x}) = \alpha_{k,y} + \beta_{k,y}x$, subject to appropriate monotonicity constraints, may provide more flexibility. An alternative model might replace $\lambda_{k,y,x}$ in (1) by the unconditional probability $\bar{\pi}_{k,y,x}$, so that

$g(\bar{\pi}_{k,y,x}) = \theta'_{k,y,x}$, which requires that $\theta'_{k,y,x}$ must increase in y for $k = 1, 2$ and all x . In the case where $\bar{\pi}_{k,y,x}$ must increase in x , however, one must impose two sets of monotonicity constraints, one in x and the other in y , which limits tractability for adaptive dose-finding. Sensitivity analyses using different marginal distributions showed that, in terms of performance of the dose-finding method, no one model among those described above is uniformly better than the others. However, our simulations (Table 6) showed that, in terms of both dose selection and choosing desirable doses for patients during the trial, the worst performance of the model with saturated marginals was better than the worst performance of the model with proportional odds marginals.

4. Decision Criteria and Trial Design

4.1 Utilities

Denote the elicited utility of outcome $\mathbf{y}$ by $U(\mathbf{y})$. The *mean utility of dose x given θ* is

$$u(x, \theta) = E_{\mathbf{Y}}\{U(\mathbf{Y}) | x, \theta\} = \sum_{y_1=0}^{m_1} \sum_{y_2=0}^{m_2} U(\mathbf{y}) \pi(\mathbf{y} | x, \theta).$$

The *posterior mean utility* is

$$\phi(x, \mathbf{D}_n) = E_{\theta}\{u(x, \theta) | \mathbf{D}_n\} = \sum_{y_1=0}^{m_1} \sum_{y_2=0}^{m_2} U(\mathbf{y}) \int_{\theta} \pi(\mathbf{y} | x, \theta) p(\theta | \mathbf{D}_n) d\theta. \quad (5)$$

Note that (5) reflects the physicians' utilities and the observed data by averaging over the posterior. We denote the dose that maximizes $\phi(x, \mathbf{D}_n)$ by x_n^{opt} . Because $U(\mathbf{y})$ is a patient utility, maximizing $\phi(x, \mathbf{D}_n)$ is very different from the more common Bayesian approach of choosing x to optimize some posterior function of the Fisher information matrix (cf. Haines, Perevozshaya and Rosenberger, 2003; Dragalin and Federov, 2006). In the context of phase I trials, Bartorff and Lai (2010) address the problem of "individual versus collective ethics" by considering an objective function including components for both current and future

patients. In our setting, always choosing based on x_n^{opt} based on $\phi(x, D_n)$ alone is an example of a “greedy algorithm,” which in general is a sequentially adaptive decision rule that always takes the locally optimal action at each stage. Motivated by both ethical and practical considerations, we next introduce additional dose acceptability criteria, and an AR procedure, to improve this greedy algorithm.

4.2 Dose Acceptability and Adaptive Randomization

Simply maximizing $\phi(x, D_n)$ ignores the undesirable but important possibility that all doses are too toxic. To control the risk of toxicity, we elicit the level y^* of toxicity considered to be unacceptable by the physicians, and an accompanying fixed limit π_1^* on the probability $\bar{\pi}_{1,y^*,x}$ of toxicity at or above y^* . We say that a dose x is *unacceptably toxic* if

$$\Pr(\bar{\pi}_{1,y^*,x} > \pi_1^* | D_n) > p_U, \quad (6)$$

where p_U is an upper probability cut-off, usually in the range .80 to .95. In the simplest case where Y_1 is a binary indicator of toxicity at or above some specified severity level, so that $y = 0$ or 1 denote the absence or presence of toxicity defined in this way, $y^* = 1$ by default. The inequality (6) says that x has a high posterior probability of producing an unacceptably high level of toxicity. We will limit all dose assignments to *the set of acceptably safe doses*, A_n^{safe} , defined to be all $x \in \mathcal{X}$ for which (6) is not the case. This is similar in spirit to the safety requirements used by Thall and Cook (2004) and Braun et al. (2007), and “escalation with overdose control” proposed by Babb, Rogatko and Zacks (1998) for phase I trials.

A very important practical problem when using $\phi(x, D_n)$ to select doses adaptively is that, in some cases, little or no information may be obtained for the dose that actually has the highest true mean utility, $u(x, \theta^{true})$. This occurs when the algorithm that chooses x by maximizing $\phi(x, D_n)$ repeatedly assigns a dose that actually is suboptimal, and does not escalate to higher levels that include the true optimal dose. The general phenomenon of a greedy sequential decision procedure becoming stuck at a suboptimal treatment is well-known. A common solution for this problem is to randomly assign some patients to suboptimal treatments. This distributes patients more evenly among treatments and consequently more is learned about the design space, often with a resulting improvement in the method’s reliability. To do this in an ethical way for adaptive dose-finding, we use the following constrained AR procedure.

Let $\{\delta_n, n = 1, \dots, N_{max}\}$ be a non-increasing sequence of differences in the utility domain. We define *the set of δ_n -optimal doses* to be

$$A_n^\delta = \{x \in \mathcal{X} : |\phi(x, D_n) - \phi(x_n^{opt}, D_n)| \leq \delta_n\}. \quad (7)$$

This is the set of doses having posterior mean utility within δ_n of the maximum value. The sequence $\{\delta_n\}$ quantifies what is meant by posterior mean utilities being “close” in expression (7); and we require it to be non-increasing with n to accommodate the decreasing variability in the posteriors of the $u(x, \theta)$ ’s as n increases. Restriction of the AR to doses in A_n^δ is motivated by both ethical considerations and the practical fact that the posteriors of the utilities $\{u(x, \theta), x \in A_n^{safe}\}$ may be quite disperse, especially for the small values of n in a dose-finding trial. Moreover, $\phi(x, D_n)$ may be nearly flat around its maximum, with the numerical superiority of $\phi(x_n^{opt}, D_n)$ over $\phi(x, D_n)$ for one or more $x \neq x_n^{opt}$ quite small. Our simulations, summarized below in Section 5, show that it sometimes is more ethical to treat some patients at suboptimal doses having $\phi(x, D_n)$ near $\phi(x_n^{opt}, D_n)$ because on average, in

many scenarios, this leads to more patients in the trial being treated at doses having higher utilities.

A third acceptability criterion that may be used is to require that a dose should not be unlikely to have the highest utility. We say that a dose x is *unlikely to be best* if

$$\Pr[u(x, \theta) = \max_{z \in \mathcal{X}} \{u(z, \theta)\} | \mathbf{D}_n] < p_L \quad (8)$$

for a small lower probability cut-off p_L . We denote the set of doses that do not have this property by $\mathbf{A}_n^{\min}$. A dose in $\mathbf{A}_n^{\min}$ is admissible in the sense that it satisfies the minimality requirement that it has at least a non-trivial probability of having the highest utility. Combining the three criteria (6), (7), and (8), we define *the set of acceptable doses* to be

$$\mathbf{A}_n = \mathbf{A}_n^{\text{safe}} \cap \mathbf{A}_n^{\delta} \cap \mathbf{A}_n^{\min}. \quad (9)$$

Thus, a dose is acceptable if it (i) has acceptable toxicity, (ii) has posterior mean utility that is δ_n -close to the maximum, and (iii) is not unlikely to have the highest posterior utility.

Our design randomizes patients adaptively among the doses in $\mathbf{A}_n$. This may be done using many different criteria. An ethically attractive approach is to define AR probabilities in terms of a set of “good” outcomes, defined as $G = \{y : U(y) \geq \underline{U}\}$, where the lower limit $\underline{U}$ is elicited from the physicians who provided the utilities. Given G , the *probability of a good outcome for a patient treated with dose x* is $\Pr(Y \in G | x, \theta)$. Denoting the posterior means $\mu_G(x, \mathbf{D}_n) = E\{\Pr(Y \in G | x, \theta) | \mathbf{D}_n\}$, we randomize a patient to dose $x \in \mathbf{A}_n$ with probability

$$r(x, \mathbf{D}_n) = \frac{\mu_G(x, \mathbf{D}_n)}{\sum_{z \in \mathbf{A}_n} \mu_G(z, \mathbf{D}_n)}. \quad (10)$$

Thus, the dose-finding method includes the requirements that it only assigns doses that are safe and δ_n -optimal, as defined by (6) and (7), and that satisfy the minimality requirement (8), and it uses the good outcome set G to determine the AR probabilities (10).

4.3 Establishing a Prior

Denote the normal distribution with mean μ and variance σ^2 by $N(\mu, \sigma^2)$, and denote the normal distribution truncated below at 0 by $N_0(\mu, \sigma^2)$. For the prior $p(\theta | \tilde{\theta})$, we assume $\mu_{k,y} \sim N(\tilde{\mu}_{k,y}, \tilde{\sigma}_{\mu,k,y}^2)$ and $\gamma_{k,y,x} \sim N_0(\tilde{\gamma}_{k,y,x}, \tilde{\sigma}_{\gamma,k,y,x}^2)$. The numerical values of the hyperparameters $\tilde{\theta} = (\tilde{\mu}, \tilde{\gamma}, \tilde{\sigma}_{\mu}^2, \tilde{\sigma}_{\gamma}^2)$ may be established from elicited probabilities using the following algorithm, similar to that given in Thall, et al. (2011).

Step 1. Assume a non-informative pseudo-prior on θ with all entries $N(0, \sigma_o^2)$ for large σ_o^2 .

Step 2. Use the elicited prior probabilities to simulate a large pseudo-sample of size N balanced equally among the doses.

Step 3. Compute the pseudo-posterior from the pseudo-prior and pseudo-sample, and record the pseudo-posterior mean.

Step 4. Repeat steps 2 and 3 M times, and set the prior mean $(\tilde{\mu}, \tilde{\gamma})$ of θ equal to the mean of the M pseudo-posterior means.

Step 5. Using the effective sample size (ESS) as a criterion, calibrate the values of $(\tilde{\sigma}_\mu^2, \tilde{\sigma}_\gamma^2)$ to obtain ESS values of the $\theta_{k,y,x}$'s in the range 0.20 to 1.0.

As a practical guideline, this algorithm may be applied effectively with $N = 100J$, i.e. 100 observations per dose, $M = 1000$ pseudo-samples, and pseudo-prior variances σ_o^2 in the range 10^2 to 100^2 , with the particular value of σ_o^2 chosen to ensure that the pseudo-posteriors are insensitive to the pseudo-prior. In practice, one or two numerical values may be used for the entries of $(\tilde{\sigma}_\mu^2, \tilde{\sigma}_\gamma^2)$. If desired, an elaboration of Step 5 is to simulate the trial for each of several values of the hyper variances to ensure that the trial will have good operating characteristics, as well as the prior being non-informative in terms of ESS. One overall ESS of $p(\theta | \tilde{\theta})$ may be computed using the formal method of Morita, Thall, and Mueller (2008) or, alternatively, one may approximate the prior of each $\pi_{k,y,x}$ as a beta(a, b), set its ESS equal to $a + b$, and average these values to obtain a single summary ESS. For the association parameter ρ in the Gaussian copula, a uniform prior on $(-1, +1)$ may be assumed.

4.4 Trial Conduct

The trial is conducted as follows. An initial cohort is treated at a starting dose chosen by the physicians. For all subsequent cohorts, once the posterior is updated based on the observed outcomes of previous patients, if A_n is empty then the trial is stopped and no dose is chosen; otherwise, patients are randomized among the doses in A_n using the updated AR probabilities given by (10). A rule superseding the above is that no untried dose may be skipped when escalating. At the end of the trial, if A_n is not empty, $x_{N_{max}}^{opt}$ is selected.

To summarize, the design requires the following quantities to be elicited from the physicians: the prior means of all $\pi_{k,y,x}$'s, the utilities, the cut-off $\underline{U}$ that determines G , the safety parameters y^* and π_1^* , a starting dose, and the number of patients treated at the starting dose. Given this information, prior parameters must be determined from the elicited values, e.g. by using the simulation-based approach described earlier, and the additional design parameters $p_U, p_L, \{\delta_n\}, N_{max}$, and cohort size must be specified. These parameters, along with prior variances, should be calibrated by simulating the trial to obtain good operating characteristics.

5. Application to the Radiation Therapy Trial

5.1 Trial Design

For the RT trial, since $m_1 = m_2 = 3$, the distribution $\{\pi_k(y | x, \theta), y_1, y_2 = 0, 1, 2, 3\}$ is determined by 15 probabilities for each x . Since $J = 3$, there are $J(m_1 m_2 + m_1 + m_2) = 45$ bivariate probabilities. For each k and x , the marginals (3) are given by

$$\begin{aligned}\pi_{k,0,x} &= 1 - \lambda_{k,1,x} \\ \pi_{k,1,x} &= \lambda_{k,1,x} - \lambda_{k,1,x} \lambda_{k,2,x} \\ \pi_{k,2,x} &= \lambda_{k,1,x} \lambda_{k,2,x} - \lambda_{k,1,x} \lambda_{k,2,x} \lambda_{k,3,x} \\ \pi_{k,3,x} &= \lambda_{k,1,x} \lambda_{k,2,x} \lambda_{k,3,x}.\end{aligned}$$

We assume a logit link, so $\lambda_{k,y,x} = e^{\theta_{k,y,x}} / (1 + e^{\theta_{k,y,x}})$. For each $k = 1, 2$ and $y = 1, 2, 3$, $\theta_{k,y} = (\mu_{k,y}, \gamma_{k,y,2}, \gamma_{k,y,3})$, so $\dim(\theta_k) = 9$ and $p = \dim(\theta) = 19$.

The prior was obtained from the elicited prior mean probabilities in Table 1 using the algorithm described in Section 4.3. We simulated 1000 pseudo samples, each of size 300 with 100 patients in each pseudo sample assigned to each dose. For each simulated data set,

a pseudo posterior was computed starting with a pseudo-prior on θ with each $\theta_{k,y,x} \sim N(0, 60^2)$, and the mean of the 1000 pseudo posterior means was used as the prior means $\tilde{\mu}_{k,y}$ and $\tilde{\gamma}_{k,y,x}$. The standard deviations of the $\theta_{k,y,x}$'s were calibrated by approximating the prior of each $\pi_{k,x,y}$ with a beta(a, b) and using ESS = $a + b$. Setting all $\tilde{\sigma}_{\mu,k,y} = \tilde{\sigma}_{\gamma,k,y,x} = \tilde{\sigma} = 6$ gave ESS values ranging from 0.31 to 0.70, with mean ESS = 0.42.

For the safety criterion, A_n^{safe} , the physicians specified $y^* = 3$ (severe toxicity) and $\pi_1^* = 0.10$. Using the conservative upper cut-off $p_U = 0.80$, a dose x is unacceptably toxic if $\Pr(\pi_{1,x,3} | D_n) > .80$. The lower cut-off $p_L = 0.10$ was used to define A_n^{min} , after studying effects of the values $p_L = 0.05, 0.10, 0.15$ in preliminary simulations. An initial cohort of 3 patients will be treated at $x = 1$ (BED = 40.00), with subsequent cohorts of size 1 and the AR started at the 4th patient. A maximum of $N_{max} = 30$ patients will be treated, chosen in part because an accrual rate of 6 to 10 patients per year is anticipated, so it will require 3 to 5 years to complete the trial. The physicians specified the lower utility $\underline{U} = 25$ to determine the good outcome set, which is used to define the AR probabilities. The nine outcomes considered to be good by this criterion are shown by the gray shaded values in Table 3. After a preliminary sensitivity analysis examining the design's behavior by simulation for various sequences $\{\delta_n\}$, including 10 δ_n 30 and several non-decreasing functions, it was decided to use the step function $\delta_n = 20$ for 4 n 15 and $\delta_n = 15$ for 16 n 30 to define A_n^δ . All posterior quantities were computed using MCMC with Gibbs sampling (Robert and Cassella, 1999). All simulations are based on 5000 iterations of each case studied.

5.2. Simulation Results

Using this design, we simulated the trial under eight scenarios, each determined by a set of assumed true probabilities, $\{\pi_{k,y,x}^{true}\}$, given in Figure 1. Scenario 1 is based on the elicited prior means, which give equal utility 64.6 to $x = 1$ and 2 and utility 57.0 to $x = 3$. Scenario 2 has steeply increasing utility with $x = 3$ best. Scenario 3 has steeply decreasing utility with $x = 1$ best, the middle dose $x = 2$ is best in Scenario 4, and the utility is V-shaped in Scenario 5 with $x = 3$ best. No dose is acceptably safe in Scenario 6. Scenario 7 is similar to Scenario 4 in terms of utilities, but has slightly higher toxicity and efficacy probabilities. The utility is V-shaped in Scenario 8 with $x = 1$ best. In each simulation scenario, the true marginal probabilities $\{\pi_{k,y,x}^{true}\}$ do not depend on the assumed model, although we use the copula to determine correlations. The Gaussian copula correlation parameter $\rho^{true} = 0.10$ was used throughout. The true utilities $u^{true}(x)$ are determined by the assumed $\{\pi_{k,y,x}^{true}\}$, ρ^{true} , and the elicited utilities $U(y)$ in Table 1.

The results of simulating the RT trial design under Scenarios 1 – 8 are summarized in Table 3. In Scenario 6, where no dose is acceptably safe and the true severe toxicity probabilities are 0.25, 0.28, 0.30 at the three doses, the method correctly stops early and chooses no dose 91% of the time, and treats on average 14.6 patients. Modifying Scenario 6 so that the true severe toxicity probabilities at the three doses are the slightly higher values .30, .35, .40, the stopping probability is 98% and the mean sample size drops to 11.6 patients. In all of the other 7 scenarios, the selection probabilities reflect the utilities quite closely. The algorithm's selection reliability is quite striking in Scenarios 5 and 8, which have V-shaped utilities with the middle dose $x = 2$ least desirable. The sample size distributions are biased toward lower doses in all scenarios, reflecting the initial cohort of 3 patients at $x = 1$, the do-not-skip rule, and the fact that the prior was biased toward the lower doses, which favors AR to these doses early in the trial.

While Table 3 shows that our method has very desirable properties for the RT trial design, since it relies on three particular admissibility criteria and AR, it is important to assess each

design component's effect on the method's performance. We thus simulated four alternative versions of the procedure. Method 1 is our proposed procedure with A_n given by (9).

Method 2 drops the minimality requirement A_n^{min} from A_n . Method 3 drops both the minimality requirement A_n^{min} and the requirement A_n^δ in the AR that a dose must have $\phi(x, D_n)$ that is δ_n -close to the optimum. Thus, Methods 1, 2, and 3 all use AR based on the posterior probabilities $r(x, D_n)$ of a good outcome, but with different definitions of an acceptable dose. Method 4 drops the AR entirely, and uses the greedy algorithm that simply chooses x to maximize $\phi(x, D_n)$, subject only to the safety constraint on toxicity. Thus, $A_n = A_n^{safe}$ for both Methods 3 and 4, but Method 4 does not use AR. The results of this additional comparative simulation are given in Table 4. To present the results in a more compact way than giving four versions of Table 2, one for each version of the method, instead we use the following two summary statistics to quantify each method's performance. For each scenario, let $u^{true}(x_{select})$ denote the true utility of the final selected dose, and u_{max} and u_{min} the largest and smallest possible true utilities among all $x \in \mathcal{X}$. A statistic that quantifies how well the method selects a final dose for future patients is

$$R_{select} = \frac{u^{true}(x_{select}) - u_{min}}{u_{max} - u_{min}},$$

the proportion of the difference between the utilities of the best and worst possible choices that is achieved by x_{select} . A similar statistic that quantifies how well the method assigns doses to patients throughout the trial is

$$R_{treat} = \frac{\frac{1}{N} \sum_{i=1}^N u^{true}(x_{[i]}) - u_{min}}{u_{max} - u_{min}},$$

where $u^{true}(x_{[i]})$ is the true utility of the dose given to the i^{th} patient, and N is the achieved final sample size. Both statistics have range $[0, 1]$, with a larger value corresponding to better overall performance. Thus, R_{select} quantifies future patient benefit while R_{treat} quantifies the benefit to the patients in the trial.

To interpret the results in Table 4, we first note that all four methods are safe in that they stop and select no dose between 90% and 93% of the time in the toxic Scenario 6. This is as expected, since all of the methods include the safety requirement (6). For the other seven scenarios, Table 4 shows that Method 4, the greedy algorithm with the safety requirement but no AR, may perform either extremely well (in Scenarios 1, 3, and 8), moderately well (in Scenarios 4 and 5), or extremely poorly (in Scenarios 2 and 7). That is, Method 4 is inconsistent and sometimes yields disastrous results. In particular, Method 4 gives the extremely small values $R_{treat} = 0.10$ in Scenario 2 and $R_{treat} = 0.20$ in Scenario 7, compared to much larger values for Methods 1 – 3. Thus, depending on the actual state of nature, the greedy algorithm without AR may perform either very well or very poorly in terms of benefit to the patients in the trial. In contrast, across all scenarios considered Methods 1 – 3 are all much more reliable, with no scenarios where any of these methods have extremely poor performance. Thus, the addition of AR acts like an insurance policy against disaster, but with the price being reduced performance in terms of R_{select} and R_{treat} in some cases. Comparing Methods 1, 2, and 3, the greatest effect of adding the admissibility constraints A_n^δ and A_n^{min} is that they both increase R_{treat} substantially in almost all scenarios. Thus, these two dose admissibility constraints have the effect of making the dose-finding algorithm more ethically attractive for the patients in the trial. In contrast, in terms of future patient benefit, Table 4 shows that most of the advantage over Method 4 in Scenarios 2, 4, 5, and 7

is due to including AR, since the R_{select} values for Methods 1, 2, and 3 about the same in these scenarios.

Our motivating application is somewhat atypical in that most phase I/II trials have more than 3 dose levels. To more fully illustrate the methodology, we also include a simulation study with $J = 5$ doses. The model, which now has $p = 31$ parameters, and the design used for this simulation are very similar to those of the RT trial, but we assume $N_{max} = 40$. With regard to choice of N_{max} in practice, it should be kept in mind that a phase I/II trial replaces the more conventional approach of conducting phase I based on toxicity only and phase II based on efficacy only. Thus, $N_{max} = 40$ is quite reasonable for a phase I/II trial. All numerical values in the prior and model corresponding to the new doses $x' \in \{1, 2, 3, 4, 5\}$ were obtained by matching values for $x' = 1, 3, 5$ to $x = 1, 2, 3$, respectively, and interpolating to obtain values for $x' = 2$ and $x' = 4$. This simulation is summarized in Table 5, which shows that the qualitative behavior of the design for $J = 5$ dose levels is very similar to what is shown for $J = 3$ in Table 3.

A natural question is how well the method works using a more conventional model with fewer parameters. To answer this, we also simulated the trial assuming a bivariate model with PO marginals defined by $\text{logit}(\bar{\pi}_{k,y,x}) = \alpha_{k,y} + \beta_k x$ with $\alpha_{k,1} > \alpha_{k,2} > \alpha_{k,3}$ so $\theta = (\alpha_{1,1}, \alpha_{1,2}, \alpha_{1,3}, \alpha_{2,1}, \alpha_{2,2}, \alpha_{2,3}, \beta_1, \beta_2, \rho)$ and $p = 9$. The prior for this model was obtained similarly to that of the 19-parameter model, using the method described in Section 4.3. The results are summarized in Table 6, which shows that the simpler model gives a design with greatly inferior performance in Scenarios 1, 4, 6, and 7, superior performance in Scenarios 2 and 5, and the comparative results are mixed in Scenarios 3 and 8. This illustrates the point that, in the present setting, no model is uniformly best. For the model with saturated marginals, excluding Scenario 6 where no doses are acceptable, R_{select} has range [0.64, 0.87] and R_{treat} has range [0.45, 0.84], whereas the proportional odds model has corresponding ranges [0.32, 0.92] and [0.30, 0.82]. Thus, the saturated model performs much more consistently across a diverse set of scenarios, with far better worst-case results than the simple model.

6. Discussion

In our proposed design, we have modified the approach of maximizing the posterior mean utility by imposing several dose acceptability criteria and using AR among near-optimal doses. As shown in Table 4, each modification improved the method's performance, in terms of both reliability and ethical desirability. While it may seem counterintuitive, in some scenarios randomizing patients to interimly sub-optimal doses actually increases the average utility of the doses assigned to patients throughout the trial. We also found that using a model with saturated marginals gives a design with more consistent overall performance compared to a model with conventional proportional odds marginals.

Several questions remain, and are areas for future investigation. In principle, the utility function might be generalized to include information, monetary costs, or a future patient horizon. Since these utilities have qualitatively different ranges, combining them to actually conduct a trial is a difficult practical and ethical problem (cf. Dragalin and Fedorov, 2006; Bartroff and Lai, 2010). Other issues include the use of bivariate event times as outcomes (cf. Yuan and Yin, 2009), and addressing the more general problems of optimizing a two-agent combination or both dose and schedule (Braun, et al. 2007).

Acknowledgments

The authors thank H. Fontanilla and A. Mahajan for presenting the problem of designing the radiation therapy trial, and for providing the elicited values, and Pat Fox for preparing Figure 1. This research was supported by NCI grant RO1-CA-83932.

References

- Babb J, Rogatko A, Zacks S. Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine*. 1998; 17:1103–1120. [PubMed: 9618772]
- Bartroff J, Lai TL. Approximate dynamic programming and its applications to the design of phase I cancer trials. *Statistical Science*. 2010; 25:255–257.
- 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. *J American Statistical Assoc*. 2004; 99:26–35.
- Berger, James O. *Statistical Decision Theory and Bayesian Analysis*. 2. New York: Springer-Verlag; 1985.
- Braun TM. 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]
- Dragalin V, Fedorov V. Adaptive designs for dose-finding based on efficacy–toxicity response. *Journal of Statistical Planning and Inference*. 2006; 136:1800–1823.
- Fowler J. The Linear-quadratic formula and progress in fractionated radiotherapy. *Br J Radiology*. 1989; 62:679–694.
- Haines LM, Perevozskaya I, Rosenberger WF. Bayesian optimal design for phase I clinical trials. *Biometrics*. 2003; 59:591–600. [PubMed: 14601760]
- Houede N, Thall PF, Nguyen H, Paoletti X, Kramar A. Utility-based optimization of combination therapy using ordinal toxicity and efficacy in phase I/II trials. *Biometrics*. 2010; 66:532–540. [PubMed: 19673865]
- McCullagh P. Regression models for ordinal data (with discussion). *J Royal Statistical Society, Series B*. 1980; 42:109–142.
- McCullagh, P.; Nelder, JA. *Generalized Linear Models*. 2. New York: Chapman and Hall; 1989.
- Nelsen, RB. *An Introduction to Copulas Lecture Notes in Statistics*. Vol. 139. New York: Springer-Verlag; 1999.
- 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.
- Sutton, RS.; Barto, AG. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press; 1998.
- Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*. 2004; 60:684–693. [PubMed: 15339291]
- Thall PF, Nguyen HQ, 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]
- Van Meter EM, Garrett-Mayer E, Bandyopadhyay D. Proportional odds model for dose-finding clinical trial designs with ordinal toxicity grading. *Statistics in Medicine*. 2011; 30:2070–2080. [PubMed: 21344472]
- 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]

- Yuan Y, Yin G. Bayesian dose-finding by jointly modeling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society, Series C*. 2009; 58:954–968.
- Zhang W, Sargent DJ, Mandrekar S. An adaptive dose-finding design incorporating both toxicity and efficacy. *Statistics in Medicine*. 2005; 25:2365–2383. [PubMed: 16220478]

Figure 1.
Fixed probabilities of all outcomes for each of the eight scenarios in the simulation study.

Elicited consensus utilities $U(y_1, y_2)$ for all 16 possible patient outcomes $(y_1, y_2) = (\text{Toxicity severity, Efficacy score})$ in the radiation therapy trial. The gray-shaded utilities identify the set of “good” outcomes, defined as having utility ≥ 25 .

Table 1

		Toxicity Severity			
		Low	Moderate	High	Severe
Efficacy Score	0	50	25	10	0
	1	85	50	15	5
	2	92	60	20	7
	3	100	75	25	10

Table 2

Elicited prior mean marginal toxicity severity and efficacy score probabilities for the model used in the radiation therapy trial.

x	BED	$Y_1 = \text{Toxicity Severity}$			$Y_2 = \text{Efficacy Score}$				
		Low	Moderate	High	Severe	0	1	2	3
1	40.00	0.65	0.20	0.12	0.03	0.20	0.40	0.35	0.05
2	45.76	0.55	0.25	0.15	0.05	0.10	0.30	0.45	0.15
3	53.39	0.40	0.30	0.23	0.07	0.10	0.20	0.50	0.20

Table 3

Simulation results. The assumed true marginal probabilities $\{\pi_{k,y,x}^{true}\}$ defining each scenario are given in Figure 1. Under each scenario, each tabled dose selection % and sample size is the mean from 5000 simulated trials. “None” means that no dose was selected.

Dose	$\mu^{true}(x)$	Scenario 1				Scenario 2			
		% Sel	# Pats	$\mu^{true}(x)$	% Sel	% Sel	$\mu^{true}(x)$	% Sel	# Pats
$x = 1$	64.6	50	15.7	51.1	13	11.3			
$x = 2$	64.6	32	8.5	58.6	22	8.8			
$x = 3$	57.0	17	5.6	63.7	61	9.5			
None		1			4				

Scenario 3				Scenario 4			
$x = 1$	65.2	63	17.2	61.7	31	13.5	
$x = 2$	60.8	20	7.2	66.2	51	10.4	
$x = 3$	54.6	14	5.0	53.1	15	5.7	
None		4			3		

Scenario 5				Scenario 6			
$x = 1$	60.7	35	14.7	37.6	6	10.3	
$x = 2$	54.3	9	6.6	37.6	2	2.9	
$x = 3$	65.7	54	8.4	33.4	1	1.4	
None		3			91		

Scenario 7				Scenario 8			
$x = 1$	52.4	14	11.5	65.4	64	17.4	
$x = 2$	64.4	51	10.7	55.1	9	6.2	
$x = 3$	59.7	32	7.4	59.9	25	6.2	
None		4			2		

Simulations comparing four alternative dose-finding methods. Method 1 is the proposed procedure. Methods 2, 3, and 4 successively drop the minimality requirement $x \in A_n^{min}$, the δ_H -close requirement $x \in A_n^\delta$, and the use of adaptive randomization (AR). The numbers in parentheses after R_{select} are the percentages of the trial being stopped with no dose selected.

Table 4

Method	1		2		3		4	
AR	Yes		Yes		Yes		No	
A_n	$A_n^{safe} \cap A_n^\delta \cap A_n^{min}$		$A_n^{safe} \cap A_n^\delta$		A_n^{safe}		A_n^{safe}	
Scenario	R_{select}	R_{treat}	R_{select}	R_{treat}	R_{select}	R_{treat}	R_{select}	R_{treat}
1	0.83 (1)	0.81	0.82 (1)	0.82	0.81 (1)	0.77	1.00 (2)	1.00
2	0.77 (4)	0.50	0.77 (2)	0.44	0.79 (1)	0.42	0.12 (5)	0.10
3	0.77 (4)	0.73	0.76 (4)	0.71	0.76 (4)	0.65	0.99 (5)	0.99
4	0.74 (3)	0.65	0.73 (1)	0.65	0.72 (1)	0.62	0.68 (2)	0.67
5	0.75 (3)	0.56	0.74 (2)	0.50	0.76 (1)	0.50	0.54 (3)	0.53
6	0.92 (91)	0.90	0.91 (90)	0.90	0.89 (90)	0.90	1.00 (93)	1.00
7	0.73 (4)	0.51	0.74 (2)	0.49	0.73 (1)	0.45	0.25 (5)	0.20
8	0.77 (2)	0.68	0.75 (2)	0.64	0.75 (1)	0.58	0.98 (2)	0.97

Table 5

Simulation results for the 5-dose version of the radiation therapy trial.

Dose	$\mu^{true}(x)$	Scenario 1			Scenario 2		
		% Sel	# Pats	$\mu^{true}(x)$	% Sel	# Pats	# Pats
$x = 1$	64.6	44	15.4	51.1	10	12.0	
$x = 2$	64.6	21	8.8	54.9	10	7.8	
$x = 3$	64.6	16	6.7	58.6	15	7.3	
$x = 4$	60.8	10	5.1	61.2	24	6.8	
$x = 5$	57.0	7	3.9	63.7	39	5.8	
None		1			2		

Dose	$\mu^{true}(x)$	Scenario 3			Scenario 4		
		% Sel	# Pats	$\mu^{true}(x)$	% Sel	# Pats	# Pats
$x = 1$	65.2	58	16.9	61.7	27	13.9	
$x = 2$	63.0	18	8.5	64.0	22	8.9	
$x = 3$	60.8	10	6.1	66.2	30	7.7	
$x = 4$	57.7	7	4.5	59.8	15	5.4	
$x = 5$	54.6	5	3.4	53.1	6	3.8	
None		3			1		

Dose	$\mu^{true}(x)$	Scenario 5			Scenario 6		
		% Sel	# Pats	$\mu^{true}(x)$	% Sel	# Pats	# Pats
$x = 1$	60.7	36	15.3	37.6	4	10.5	
$x = 2$	57.5	11	7.8	37.7	1	3.2	
$x = 3$	54.3	5	5.7	37.6	1	1.6	
$x = 4$	60.0	13	5.8	35.5	0	0.8	
$x = 5$	65.7	34	5.2	33.4	0	0.4	
None		1			94		

Dose	$\mu^{true}(x)$	Scenario 7			Scenario 8		
		% Sel	# Pats	$\mu^{true}(x)$	% Sel	# Pats	# Pats
$x = 1$	52.4	11	11.9	65.4	60	17.8	
$x = 2$	58.5	16	8.4	60.3	13	8.1	
$x = 3$	64.4	33	8.2	55.1	5	5.4	
$x = 4$	62.1	21	6.2	57.5	8	4.7	
$x = 5$	59.7	17	4.9	59.9	14	3.8	

Dose	$u^{true}(x)$	% Sel	# Pats	$u^{true}(x)$	% Sel	# Pats
None		1				
		2				

Table 6

Comparison of the performance under the simplified 9-parameter bivariate proportional odds model ($p=9$) versus the model with saturated marginals ($p=31$) in the 5-dose trial. The numbers in parentheses after R_{select} are the percentages of the trial being stopped with no dose selected.

Scenario	<u>Prop. odds marginals</u>		<u>Saturated marginals</u>	
	R_{select}	R_{treat}	R_{select}	R_{treat}
1	0.55 (3)	0.79	0.87 (1)	0.84
2	0.92 (8)	0.55	0.71 (2)	0.45
3	0.67 (8)	0.82	0.82 (3)	0.72
4	0.32 (9)	0.47	0.74 (1)	0.68
5	0.84 (4)	0.66	0.64 (1)	0.47
6	0.53 (99)	0.92	0.94 (94)	0.94
7	0.50 (11)	0.30	0.70 (2)	0.52
8	0.66 (5)	0.78	0.75 (1)	0.62