



Published in final edited form as:

J Biopharm Stat. 2015 ; 25(5): 903–920. doi:10.1080/10543406.2014.920873.

Seamless Phase I/II Adaptive Design For Oncology Trials of Molecularly Targeted Agents

Nolan A. Wages¹ and Christopher Tait²

¹Division of Translational Research & Applied Statistics, Department of Public Health Sciences University of Virginia, Charlottesville, Virginia 22908, U.S.A

²Department of Statistics, University of Virginia, Charlottesville, Virginia 22904-4135, U.S.A

Abstract

In dose-finding trials of chemotherapeutic agents, the goal of identifying the maximum tolerated dose is usually determined by considering information on toxicity only, with the assumption that the highest safe dose also provides the most promising outlook for efficacy. Trials of molecularly targeted agents challenge accepted dose-finding methods because minimal toxicity may arise over all doses under consideration and higher doses may not result in greater response. In this article, we propose a new early-phase method for trials investigating targeted agents. We provide simulation results illustrating the operating characteristics of our design.

Keywords

Continual reassessment method; Dose finding; optimal biological dose; Molecularly targeted agent

1 Introduction

1.1 Background

In general, the primary objective of Phase I clinical trials is to identify the maximum tolerated dose (MTD) of the agent or agents being investigated. In a subsequent Phase II trial, the agent is evaluated for efficacy at the recommended dose (MTD). In oncology trials of chemotherapeutic agents, identification of the MTD is usually determined by considering dose-limiting toxicity (DLT) information only, with the assumption that the MTD is the highest dose that satisfies some safety requirement, so that it provides the most promising outlook for efficacy. Numerous Phase I designs have been proposed for identifying the MTD from a discrete set of doses in which toxicity is described as a binary outcome (Storer, 1989; O'Quigley, Pepe, and Fisher, 1990; Ji, *et al.*, 2010; among many others) or an ordinal outcome (Yuan, Chappell, Bailey, 2007; Van Meter, Garrett-Mayer, and Bandyopadhyay, 2011; Iasonos, Zohar, and O'Quigley, 2011; among others). Usually, Phase I and Phase II trials are performed independently, without formally sharing information across the separate phases. There has been a recent shift in the paradigm of drug development in oncology to integrate Phase I and Phase II trials so that drug development process may be accelerated, while potentially reducing costs (Yin, 2012). To this end, several published Phase I/II methods have extended dose-finding methodology to allow for the modeling of both toxicity

and efficacy (O'Quigley, Hughes, and Fenton, 2001; Braun, 2002; Thall and Cook, 2004; Yin, Li and Ji, 2006; among others). Each of these designs, as well as the one we propose in this current work, is referred to by Federov and Leonov (2013) as a "best intention" (BI) design, meaning that it aims to allocate patients in the current trial to the "best" treatment based on the current knowledge available. We distinguish these designs from "optimal" designs in the area of dose-finding that seek to maximize global information about dose-response curves in order to optimize doses for future patients. In addition to the detailed discussion provided by Federov and Leonov (2013), we refer any reader interested in this type of design to Dragalin and Federov (2006) and Whitehead *et al.* (2006), among many others.

In general, the design of Phase I trials is driven by the assumption of monotone increasing dose-toxicity and dose-efficacy relationships. By contrast, many molecularly targeted agents (MTA's) are essentially considered safe and higher doses do not necessarily produce greater efficacious response. Dose-efficacy relationships may exhibit non-monotone patterns, such as increasing at low doses and either decreasing or plateauing at higher levels, as displayed in Figure 1. Therefore, the goal of the trial shifts to identifying the optimal biological dose (OBD), which is defined as the dose with acceptable toxicity that maximizes efficacious response. Despite the emergence of molecularly targeted agents in oncology drug development, there are relatively few statistical methods for designing Phase I trials of these agents (Mandrekar, Qin, and Sargent, 2010). Assuming minimal toxicity over the therapeutic dose range, Hunsberger *et al.* (2005) described two practical Phase I designs for identifying the OBD for molecularly targeted agents. Zhang, Sargent and Mandrekar (2006) introduced a trinomial continual reassessment method (TriCRM) that utilizes a continuation ratio to model toxicity and efficacy simultaneously. Mandrekar, Cui, and Sargent (2007) extended this method to account for combinations for two agents. Polley and Cheung (2008) outlined a two-stage design for finding the OBD that implements a futility interim analysis. Recently, Hoering, LeBlanc, and Crowley (2011) and Hoering *et al.* (2013) proposed methods for early phase trials of targeted agents that uses a traditional dose-finding method to find the MTD in Phase I, and then subsequently allocates patients in Phase II to the MTD (Hoering *et al.*, 2011) or a small set of doses at and around the MTD (Hoering *et al.*, 2013) based on a randomization scheme.

In this article, we outline a method for identifying the OBD that combines features of the continual reassessment method (CRM; O'Quigley *et al.*, 1990) and order restricted inference (Robertson, Wright, and Dykstra, 1988). The existing methods in this area take a fixed set of admissible doses from the Phase I portion into the Phase II portion of the study in order to assess efficacy. For instance, the design of Hoering *et al.* (2013) uses a 3+3 design in the Phase I portion to locate the MTD. However, if the OBD is not close to the MTD, a design that targets the MTD may fail to accurately locate the most efficacious dose. Moreover, the statistical limitations of the 3+3 design have been discussed at length in the literature (Garrett-Mayer, 2006; Iasonos *et al.*, 2008; Ji and Wang, 2013; among others). We contrast this approach with the one we propose here in that, throughout the duration of the trial, we continuously monitor safety data in order to adaptively update our set of acceptable doses with which to make allocation decisions based on efficacy. The rest of the paper is

organized as follows. We finish this section by outlining our overall strategy and providing a motivating example for this work. In Section 2, we outline the statistical models and inference used in the proposed design. In Section 3, we describe the dose-finding algorithm and in Section 4 we report simulation findings illustrating the operating characteristics of the design. Finally, we conclude with some discussion.

1.2 Locating the OBD

As previously noted, the assumption of monotonicity with regards to dose-toxicity relationships is usually appropriate for MTA's, yet this assumption for dose-efficacy curves may fail. For these types of therapies, dose-efficacy curves could follow a non-monotone pattern, such as a unimodal or plateau relationship. In general, suppose a trial is investigating a set, $\mathcal{D} = \{d_1, \dots, d_I\}$, of I doses and that the probability of efficacious response at dose d_i is denoted $\pi_E(d_i)$. In the presence of unimodal (including monotone increasing and monotone decreasing) or plateau dose-efficacy relationships, the primary objective of the trial is to find the optimal dose, $d_\nu \in \mathcal{D}$, defined such that

$$\pi_E(d_1) \leq \dots \leq \pi_E(d_\nu) \geq \dots \geq \pi_E(d_I). \quad (1)$$

In other words, $\pi_E(d_\nu)$ corresponds to the dose d_ν where the peak of the unimodal relationship occurs, or where the dose-efficacy curve begins to plateau. Using the terminology of Hwang and Peddada (1994), $\pi_E(d_\nu)$ is said to be a *nodal* parameter (or a node) in that it is known that $\pi_E(d_\nu) \geq \pi_E(d_i)$ for all $i \neq \nu$. Order restriction (1) is referred to as an *umbrella* ordering with the node of the umbrella occurring at d_ν . For instance, if the node occurs at the highest dose, we have a monotone increasing dose-efficacy relationship. In actuality, we do not know where in the dose range this node occurs and we must account for this uncertainty. The node could occur at any of the I available dose levels, with non-decreasing efficacy probabilities for doses before the node and non-increasing efficacy probabilities for doses after the node. The strategy we propose is to use this information to formulate a set of possible dose-efficacy relationships and make progress by appealing to Bayesian model choice, in order to account for the uncertainty surrounding the shape of the true dose-efficacy curve.

In locating the OBD, accepted early-phase designs are challenged by the use of non-traditional endpoints that accompany cytostatic agents. The potential growth of targeted therapies will be driven by the need to define suitable measures of biologic effect and finding ways to incorporate them into early-phase designs. Although the relationship between clinical outcome and biologic activity may not be clear, it is generally assumed that the absence of targeted effect will accompany a lack of clinical efficacy. Before a targeted agent can be taken into large-scale trials to test for clinical efficacy, early-phase trials are needed to establish that the therapy can produce a biologic activity with the *potential* to translate to clinical benefit. Protocol-specific endpoints give the investigator a measure of targeted effect that serves as a driving factor in early-phase trial design.

Parulekar and Eisenhauer (2004) and Korn (2004) provide a discussion of the challenges presented by non-toxicity endpoints in Phase I trial design of targeted agents. Unlike

cytotoxic agents, the effects of these agents will likely be cytostatic, which will require measures of anti-tumor activity other than those used as traditional endpoints, such as tumor shrinkage. Surrogate markers of efficacy, such as measurement of target inhibition or pharmacokinetic endpoints, will be necessary. Incorporation of these endpoints can be challenging due to the fact that (1) it may be difficult to define the desired biological effect, and (2) once defined, these endpoints may lack reliable, validated assays, making them practically difficult to measure. For the methods proposed in this article, we make the assumption that a surrogate efficacy endpoint is defined and able to be measured in order to find the OBD.

1.3 A motivating example

As an example, we describe the Phase I trial of Gerber *et al.* (2011) that assessed the safety and pharmacokinetics of a vascular targeting agent, bavituximab, in patients with advanced solid tumors. The trial investigated four dose levels (0.1, 0.3, 1, 3 mg/kg) and the goal of the study was to dose escalate to the predicted biologically effective dose, rather than the MTD, because for monoclonal antibodies, the MTD may not correspond to optimal efficacy. Any of the four dose levels could be the node of the *unimodal or plateau* dose-efficacy curve, corresponding to the optimal dose. These four possible dose-efficacy relationships are provided in Table 1. There are other possibilities, in addition to those in Table 1, that could be considered for inclusion into this subset of possible relationships. For example, for a unimodal curve with node at d_3 , another possibility is given by $\pi_E(d_1) \pi_E(d_4) \pi_E(d_2) \pi_E(d_3)$. This notion of formulating a subset of possible curves has been undertaken by other authors (Conaway, Dunbar, and Peddada, 2004; Wages, Conaway, and O’Quigley, 2011) within the context dose-finding in drug combinations. These papers assessed the possible loss of information in only using a subset of relationships and they found that, in general, using a subset did not drastically reduce the performance of their methods. An important idea to this approach is that we do not need to identify the “correct” relationship at every dose level. We just need one which is “close” to correct, in that it is correct in specifying the node (optimal dose) of the umbrella ordering.

2 Methods

Consider a trial investigating a set of I doses, $\mathcal{D} = \{d_1, \dots, d_I\}$. Toxicity and efficacy are modeled as binary endpoints so that, for each subject j , we measure

$$Y_j = \begin{cases} 0 & \text{if no toxicity,} \\ 1 & \text{if toxicity} \end{cases} \quad Z_j = \begin{cases} 0 & \text{if no response,} \\ 1 & \text{if response.} \end{cases}$$

The dose for the j th entered patient, $X_j, j = 1, \dots, n$ can be thought of as random, taking values $x_j \in \mathcal{D}$.

2.1 Modeling toxicity with CRM

The CRM was introduced as an alternative to the traditional escalation schemes reviewed by Storer (1989). The details of CRM are provided by O’Quigley *et al.* (1990), so we only

briefly recall them here. In its original form the CRM is a Bayesian method that uses sequential updating of dose-toxicity information to estimate the dose level at which to treat the next available patient. The method begins by assuming a functional dose-toxicity curve, $F(d_i, \beta)$, that is monotonic in both dose level, d_i , and the parameter, β , in order to model the DLT probabilities, $\pi_T(d_i)$, at each dose i ; $i = 1, \dots, I$. For instance, the power model,

$F(d_i, \beta) = p_i^{\exp(\beta)}$, is common to CRM literature, where $0 < p_1 < \dots < p_I < 1$ are the standardized units (skeleton) representing the discrete dose levels d_i . The p_i are also referred to as a working model. After having included j subjects, we have toxicity data in the form $\Omega_j = \{(x_1, y_1), \dots, (x_j, y_j)\}$, and we can generate an estimate, $\hat{\beta}_j$, based on the posterior mean of β . Using $\hat{\beta}_j$, we can obtain an estimate of the DLT probability at each dose level via $\hat{\pi}_T(d_i) = F(d_i, \hat{\beta}_j)$, $i = 1, \dots, I$, from which the dose recommended for the $(j + 1)$ th patient is determined. Specifically, for target DLT rate, φ_T , this dose is $d_i = \arg \min_i |F(d_i, \hat{\beta}_j) - \varphi_T|$. The MTD is the recommended dose after the inclusion of some predetermined sample size of N patients is exhausted or a stopping rule takes effect. In the bivariate extension of the CRM that we describe in this article, instead of using the DLT probability estimates, $\hat{\pi}_T(d_i)$, to determine the recommended dose for the next patient, we will use them in defining a set of acceptable doses with regards to safety. To this end, φ_T will be redefined from a target DLT probability to a maximum acceptable toxicity rate that will guide the definition of safe doses.

2.2 Models and inference for efficacy

Suppose now that, instead of utilizing a single working model as we did with toxicity, we make use of some class of working models and model selection techniques in order to allow for more flexibility in modeling the dose-efficacy relationship. This class of skeletons would correspond to various possible dose-toxicity relationships (Table 1). The models and inference for efficacy are based on specifying both a unimodal and plateau skeleton for each of the possible dose-efficacy relationships. For instance, if the node is located at d_2 , we identify a skeleton that (1) peaks at d_2 and decreases after d_2 , and another that (2) begins to plateau at d_2 and remains constant after d_2 . Using Table 1 as a guide, in general, we need to specify $K = 2 \times I - 1$ working models; I unimodal skeletons, with nodes at each of the I doses, and $I - 1$ plateau skeletons, with nodes at each of the first $I - 1$ doses. The dose-efficacy relationship does not plateau at dose level d_I , so only one working model with node at d_I is needed. For a particular skeleton, k ; $k = 1, \dots, K$, we model $\pi_E(d_i)$, the true probability of efficacious response at d_i by

$$\pi_E(d_i) = \Pr(Z_j = 1 | d_i) \approx G_k(d_i, \theta) = q_{ik}^{\exp(\theta)} \quad k = 1, \dots, K,$$

for a class of working dose-efficacy models, $G_k(d_i, \theta)$ and $\theta \in \Theta$. Like toxicity, for efficacy we appeal to the simple power model, which has shown itself to work well in practice. Here, $\Theta = (-\infty, \infty)$ and $0 < q_{1k} < \dots < q_{Ik} < 1$ is the skeleton of model k . Further, we may wish to take account of any prior information concerning the plausibility of each model and so introduce $\alpha(k) = \{\alpha(1), \dots, \alpha(K)\}$, where $\alpha(k) \geq 0$ and where $\sum_k \alpha(k) = 1$. Even when there is no prior information available on the possible skeletons we can formally proceed in the

same way by specifying a discrete uniform for $\tau(k)$. After inclusion of the first j patients into the trial, we have efficacy data in the form of $\mathcal{D}_j = \{(x_1, z_1), \dots, (x_j, z_j)\}$. Under working model k , we obtain an estimate, $\hat{\theta}_{jk}$, for the parameter θ . In the Bayesian framework, we assign a prior probability distribution $h(\theta)$ for the parameter θ of each model, and a prior probability $\tau(k)$ to each possible skeleton. For each model, we utilize a Normal prior with mean 0 and variance 1.34 on θ so that $h(\theta) = \mathcal{N}(0, 1.34)$ (O'Quigley and Shen, 1996). In order to establish running estimates of the probability of efficacy at the available doses, we need an expression for the likelihood for the parameter θ . After inclusion of the first j patients into the study, the likelihood under model k is given by

$$L_k(\theta|\mathcal{D}_j) = \prod_{\ell=1}^j \{G_k(x_\ell, \theta)\}^{z_\ell} \{1 - G_k(x_\ell, \theta)\}^{(1-z_\ell)} \quad (2)$$

which, for each skeleton, can be used in order to generate a summary value, $\hat{\theta}_{jk}$, for θ . Given the set \mathcal{D}_j and the likelihood, the posterior density for θ is given by

$$P_k(\theta|\mathcal{D}_j) = \frac{L_k(\theta|\mathcal{D}_j)h(\theta)}{\int_{\Theta} L_k(\theta|\mathcal{D}_j)h(\theta)d\theta}.$$

This information can be used to establish posterior model probabilities given the data as

$$w(k|\mathcal{D}_j) = \frac{\tau^{(k)} \int_{\Theta} L_k(\theta|\mathcal{D}_j)h(\theta)d\theta}{\sum_{k=1}^K \tau^{(k)} \int_{\Theta} L_k(\theta|\mathcal{D}_j)h(\theta)d\theta}.$$

The prior model probabilities, $\tau(k)$, are updated by the efficacy response data, \mathcal{D}_j . It is expected that the more the data support model k , the greater its posterior probability will be. Thus, we appeal to sequential Bayesian model choice in order to guide allocation decisions. Each time a new patient is to be enrolled, we choose a single skeleton, k^* , with the largest posterior probability such that

$$k^* = \arg \max_k w(k|\mathcal{D}_j), \quad k=1, \dots, K.$$

We then utilize $G_{k^*}(d_i, \theta)$ to generate efficacy probability estimates at each dose. Beginning with the prior for θ and having included the j th subject, we can compute the posterior probability of a response for d_i so that

$$\hat{\pi}_E(d_i) = G_{k^*}(d_i, \hat{\theta}_{jk^*}); \quad \hat{\theta}_{jk^*} = \int_{\Theta} \theta P_k(\theta|\mathcal{D}_j)d\theta.$$

The efficacy probability estimates, $\hat{\pi}_E(d_i)$, are used to make decisions regarding dose allocation as described in Section 3.

3 Dose-finding Algorithm

Overall, we are going to allocate each entered patient to the dose estimated to be the most efficacious, among those with acceptable toxicity. After obtaining the DLT probability estimates, $\hat{\pi}_T(d_i)$, for each dose, we are going to first restrict our attention to those doses with estimated probabilities less than a maximum acceptable toxicity rate, ϕ_T . In general, after j entered patients, we define the set of “acceptable” doses as

$$\mathcal{A}_j = \{d_i; \hat{\pi}_T(d_i) \leq \phi_T; i=1, \dots, I\}.$$

We then allocate the next entered patient to a dose in \mathcal{A}_j based upon estimated efficacy probabilities, $\hat{\pi}_E(d_i)$. By incorporating an acceptable set, \mathcal{A}_j , we exclude overly toxic doses. After each cohort inclusion the acceptable set is redefined based on the current DLT probability estimates, so it is possible, once more data has been observed, for \mathcal{A}_j to include a dose that was previously excluded when a limited amount of data existed. The allocation algorithm depends upon the amount of data that has been observed so far in the trial. If a limited amount of data exists, we will rely on an adaptive randomization phase to allocate future patients to acceptable doses. In the latter portion of the trial, when a sufficient amount of data has been observed, we will utilize a maximization phase in which we allocate according to the most efficacious treatment among the set of acceptable doses.

3.1 Adaptive randomization (AR) phase

Early in the trial, there may not be enough data to rely entirely on maximization of estimated efficacy probabilities within \mathcal{A}_j to accurately assign patients to the most efficacious dose with acceptable toxicity. There may be doses in \mathcal{A}_j that have never been tried and information on these can only ever be obtained through experimentation. Added randomization allows for information to be obtained on competing doses and prevents the method from getting “stuck” on a dose that has been tried early in the trial. Therefore, we do not rely entirely on the maximization of estimated efficacy probabilities for guidance as to the most appropriate treatment but rather implement adaptive randomization (AR) to obtain information more broadly. Based on the estimated efficacy probabilities, $\hat{\pi}_E(d_i)$, for doses in \mathcal{A}_j , calculate a randomization probability R_i ,

$$R_i = \frac{\hat{\pi}_E(d_i)}{\sum_{d_i \in \mathcal{A}_j} \hat{\pi}_E(d_i)}$$

and randomize the next patient or cohort of patients to dose d_i with probability R_i . We are going to rely on this randomization algorithm for a subset of n_R patients in order to allow information to accumulate on untried doses, before switching to a phase in which we simply allocate according to the maximum estimated efficacy probability among the acceptable

doses. The number of patients on which to implement the AR phase may vary from trial to trial. In some cases, clinician/statistician teams may choose to set n_R according to their personal preference. Some may favor randomization for the entirety of the study, while others may choose not to randomize any patients and treat according to which dose the data indicates is the most efficacious. We found that the method exhibits good performance, in terms of choosing desirable doses, when at least some randomization is employed. It is difficult to make a definitive recommendation for n_R , since it is not feasible to determine at what exact value that performance is optimized across a broad range of situations. Optimal performance may depend upon many factors, including the number of dose levels, the true dose-toxicity/efficacy curves, and the maximum sample size N . In the simulations below, we explore performance of our method under various values of n_R . The results will demonstrate that our method, across a wide range of scenarios, is fairly robust to this choice. The size of the AR phase can be expanded or contracted according to the desire of the clinician/statistician team and operating characteristics will be relatively unaffected. As a general rule of thumb, adaptively randomizing for half of the maximum sample size seems to work well in a wide variety of scenarios.

3.2 Maximization phase

Upon completion of the AR phase, the trial design switches to a maximization phase in which maximized efficacy probability estimates guide allocation. Among the doses contained in \mathcal{A}_j , we allocate the $(j + 1)$ th patient cohort to the dose x_{j+1} according to the estimated efficacy probabilities, $\hat{\pi}_E(d_i)$, such that

$$x_{j+1} = \arg \max_{d_i \in \mathcal{A}_j} \hat{\pi}_E(d_i)$$

Continuing in this way, the OBD is the recommended dose $d_i = x_{N+1}$ for the hypothetical $(N + 1)$ th patient after the inclusion of the maximum sample size of N patients or some stopping rule takes effect.

3.3 Starting the trial

In order to get the trial underway, we will choose the efficacy skeleton with the largest prior probability, $\pi(k)$, among the working models being considered. If several, or all, of the working models have the same maximum prior probability, then we will choose at random from these skeletons. Given the starting skeleton, k^* , for efficacy, the starting dose, $x_1 \in \mathcal{D}$, is then chosen. Specifically, based on the toxicity skeleton, p_i , we define the acceptable set, a priori, to be

$$\mathcal{A}_j = \{d_i : p_i \leq \phi_T; i=1, \dots, I\}.$$

Based on the efficacy skeleton, q_{ik}^* , for doses in \mathcal{A}_j , we calculate the randomization probability R_i ,

$$R_i = \frac{q_{ik^*}}{\sum_{d_i \in \mathcal{A}_j} q_{ik^*}}$$

and randomize the first patient or cohort of patients to dose $x_1 = d_i$ with probability R_i .

3.4 Ending the trial

Safety—Investigators will want some measure by which to stop the trial in the presence of undesirable toxicity. At any point in the trial, we can calculate an exact binomial confidence interval for toxicity $(\pi_T^-(d_1), \pi_T^+(d_1))$ at the lowest dose, d_1 . The value of $\pi_T^-(d_1)$ provides a lower bound for the probability of toxicity at d_1 , above which we are 95% confident that the true toxicity probability for d_1 falls. We want to compare this lower bound to our maximum acceptable toxicity rate, ϕ_T . If our lower bound exceeds this threshold, we can be confident that the new treatment has too high of a DLT rate to warrant continuing the trial. Therefore, if $\pi_T^-(d_1) > \phi_T$, stop the trial for safety and no treatment is identified as the OBD.

Futility—If the new treatment is no better than the current standard of care, we would also like for the trial to be terminated. After j inclusions, we calculate an exact binomial confidence interval for efficacy $(\pi_E^-(x_j), \pi_E^+(x_j))$ at the current dose, x_j . The value of $\pi_E^+(x_j)$ provides an upper bound for the probability of efficacy at the current treatment, below which we are 95% confident that the true efficacy probability for that dose falls. We want to compare this upper bound to some futility threshold, ϕ_E , that represents the efficacy response rate for some current standard of care. If our upper bound fails to reach this threshold, we can be confident that the new treatment does not have a sufficient enough response rate to warrant continuing the trial. Therefore, if $\pi_E^+(x_j) < \phi_E$, stop the trial for futility and no treatment is identified as the OBD. This stopping rule will only be assessed in the maximization phase, not the AR phase, so that here $j \leq n_R$. The reason for this is that the upper bound will be calculated on the current treatment, which is based on the maximum estimated efficacy probability. We can be confident that if the maximum of the estimated probabilities did not reach the futility threshold, then none of the other acceptable toxicity doses would reach it, since they are presumed to have lower efficacy than the selected treatment. In the AR phase, when we randomize, there may exist a more efficacious treatment in the set \mathcal{A}_j that was not selected purely because of randomization. Consequently, another treatment may have reached the futility threshold had it been selected, so we would not want to stop the trial in this case.

4 Numerical results

4.1 Simulation setting

For the first set of simulated results contained in this section, we compare the performance of the proposed approach with that of Hoering *et al.* (2013) for identifying the OBD in a trial investigating $I = 6$ dose levels. Throughout the simulation results, we denote the method of Hoering *et al.* (2013) as HMLC, and our proposed method as WT. The comparison assesses

the performance of each method over the twelve true toxicity/efficacy scenarios contained in their paper, which are provided in Figure 2. True toxicity probabilities are denoted T1 through T4, while true efficacy probabilities are denoted R1 through R3. In each set of 1000 simulated trials, the maximum sample size was set to $N = 64$. The Phase II portion of Hoering *et al.* (2013) randomized $n_2 = 48$ patients to three doses levels (arms); (1) the recommended dose (RD) from Phase I, (2) the dose level immediately above the RD (RD +1), and (3) the dose level immediately below the RD (RD-1). For response, the authors test competing hypotheses $H_0 : p = 0.05$ vs. $H_1 : p = 0.30$. The Phase I sample sizes for the Hoering *et al.* (2013) method are reported in Table 1 of Hoering *et al.* (2011), with a minimum sample size of $n_1 = 16$. Therefore, we thought a justifiable total sample size for comparison was $N = 16 + 48 = 64$. For our proposed method, the size of the AR phase was set equal to one quarter the total sample size so that $n_R = 16$. In a subsequent section, we investigated the impact of various sizes, n_R , of the AR phase on operating characteristics.

The proposed design embodies characteristics of the CRM so we can utilize these features in specifying design parameters. The toxicity probabilities were modeled via the power model given in Section 2.1 with skeleton values $p_i = \{0.01, 0.08, 0.15, 0.22, 0.29, 0.36\}$. It has been shown (O'Quigley and Zohar, 2010) that CRM designs are robust and efficient with the implementation of "reasonable" skeletons. Simply defined, a reasonable skeleton is one in which there is adequate spacing between adjacent values. This skeleton also reflects the idea that, for MTA's, it is often assumed that there is minimal toxicity over the dose range. For efficacy, probabilities were modeled via the class of power models in Section 2.2, using $K = 2 \times I - 1 = 11$ skeletons that correspond to the possible dose-efficacy relationship. The six sets of values used for the unimodal relationships were:

- 1 $q_{i1} = \{0.60, 0.50, 0.40, 0.30, 0.20, 0.10\}$
- 2 $q_{i2} = \{0.50, 0.60, 0.50, 0.40, 0.30, 0.20\}$
- 3 $q_{i3} = \{0.40, 0.50, 0.60, 0.50, 0.40, 0.30\}$
- 4 $q_{i4} = \{0.30, 0.40, 0.50, 0.60, 0.50, 0.40\}$
- 5 $q_{i5} = \{0.20, 0.30, 0.40, 0.50, 0.60, 0.50\}$
- 6 $q_{i6} = \{0.10, 0.20, 0.30, 0.40, 0.50, 0.60\}$.

The five sets of values for the plateau relationships were:

- 7 $q_{i7} = \{0.20, 0.30, 0.40, 0.50, 0.60, 0.60\}$
- 8 $q_{i8} = \{0.30, 0.40, 0.50, 0.60, 0.60, 0.60\}$
- 9 $q_{i9} = \{0.40, 0.50, 0.60, 0.60, 0.60, 0.60\}$
- 10 $q_{i,10} = \{0.50, 0.60, 0.60, 0.60, 0.60, 0.60\}$
- 11 $q_{i,11} = \{0.60, 0.60, 0.60, 0.60, 0.60, 0.60\}$

We assumed, a priori, that each of these eleven models was equally likely and set $\tau(k) = 1/11$.

Another design parameter that must be addressed is the association between toxicity and efficacy responses. The models and inference presented in Section 2 estimate toxicity and efficacy probabilities independently, ignoring an association parameter, ψ . In order to provide a justifiable comparison to Hoering *et al.* (2013), as well as to illustrate our method's robustness to the misspecification of an association parameter, we generated correlated binary outcomes using the function `ranBin2` in R package `binarySimCLF` (R Core Team). The log odds ratio specification used to generate the data was set to $\psi = 4.6$ to match that used by Hoering *et al.* In each set of 1000 simulations, in order to define the acceptable set \mathcal{A} , the maximum acceptable toxicity rate was specified to be $\phi_T = 0.33$ and the minimum efficacy threshold to be $\phi_E = 0.05$. Again, these values were chosen in order to present results in a setting mimicking that of Hoering *et al.* (2013). For each scenario, using the definitions of Hoering *et al.*, the **best** dose is defined as the level that maximizes efficacy while assuring safety, and a **good** dose is defined as a level where efficacy is above predefined boundary while maintaining safety. In Figure 2, the **best** doses are indicated with gray shading, while **good** doses are indicated with a \square .

4.2 Operating characteristics

The results of the comparison are summarized in Table 2 and Figure 3, with each bar in the figure representing the proportion of simulated trials that each method selected a **good** dose [Figure 3(a)] and the **best** dose [Figure 3(b)]. Overall, the results indicate a superior performance for the proposed method over that of Hoering *et al.* (2013). The left column of Figure 2 and the first four rows of Table 2 represent cases in which the dose-efficacy curves have monotone increasing probabilities. In the first three of these scenarios (R1T1, R1T2, and R1T3), d_4 is the only good and best dose. WT outperforms HMLC in these scenarios by selection percentages of 40%, 34% and 8%, respectively. In R1T4, the toxicity probabilities plateau at a value of 0.05 across the dose range, so the highest dose, d_6 , is considered the best dose. The HMLC method selects the best dose as the OBD in 63% of simulated trials, compared to 57% by the WT method. The next four rows of Table 2 and the middle column of Figure 2 illustrate situations in which the dose-efficacy curves increase until dose d_4 and remain constant after d_4 (plateau). Again, in terms of selecting both good and best doses, WT improves upon the selection percentages of HMLC in most cases (R2T1, R2T2, R2T4), by as much as 30 to 40% in some scenarios. HMLC performs slightly better in selecting a good dose in R2T3, doing so in 71% of trials compared to 67% by WT. Finally, in the last four rows of Table 2 and the right column of Figure 2, most of the doses are considered to be best/good doses, so, as expected, both methods produce high selection percentages in these cases. The WT method does perform slightly better in these last four scenarios, yielding selection percentages that range from 3% to 10% higher than those produced by HMLC.

We feel that this improved performance by the WT across the 12 scenarios considered here can mostly be attributed to the HMLC's use of a Phase I portion to identify an MTD. This may not be appropriate for MTA's, due to the uncertainty surrounding the shape of the dose-efficacy curve. If the MTD-based Phase I portion does not identify doses around the OBD, then the Phase II portion could possibly be randomizing to incorrect doses. In contrast, throughout the duration of the trial, our method continuously monitors toxicity and efficacy

data to identify a set of safe doses, within which we allocate to a dose with high efficacy. This allows for the adaptive updating of the doses being considered after each patient inclusion, rather than restricting attention to two or three doses after a Phase I component designed to find the MTD. Overall, the strong showing of our method against published work in the area in extensive simulation studies makes us feel confident in recommending it as a viable alternative in Phase I/II trials of MTA's.

4.3 Sensitivity analysis

In this set of simulations, we investigated the sensitivity of our proposed method to the design specifications over a class of five true (toxicity, efficacy) scenarios. The analysis assesses the effect of changing (1) the number of dose levels, (2) the true value of the association parameter ψ , (3) the size of the AR phase n_R , and (4) the maximum sample size N . We also investigated the sensitivity of the proposed methods to the choice of both toxicity and efficacy skeletons, but we exclude these results for the sake of brevity. As expected, since our method leans upon the CRM, it is robust to skeleton choice, so long as "reasonable" skeletons are chosen (O'Quigley and Zohar, 2010). We are happy to share these results with any interested reader.

We ran simulations using outcomes generated from the three fixed values of ψ : $\psi = \{-2, 0, -2\}$. The maximum sample size for each simulated trial was set to $N = 48$, and we varied the size of the AR phase; i.e. $n_R = \{12, 24, 36, 48\}$. The toxicity probabilities were modeled via the power model given in Section 2.1 with skeleton values $p_i = \{0.01, 0.08, 0.15, 0.22, 0.29\}$. For efficacy, probabilities were modeled via the class of power models in Section 2.2, using $K = 2 \times I - 1 = 9$ skeletons. The nine sets of values used were: $q_{i1} = \{0.30, 0.40, 0.50, 0.60, 0.70\}$, $q_{i2} = \{0.40, 0.50, 0.60, 0.70, 0.60\}$, $q_{i3} = \{0.50, 0.60, 0.70, 0.60, 0.50\}$, $q_{i4} = \{0.60, 0.70, 0.60, 0.50, 0.40\}$, $q_{i5} = \{0.70, 0.60, 0.50, 0.40, 0.30\}$, $q_{i6} = \{0.70, 0.70, 0.70, 0.70, 0.70\}$, $q_{i7} = \{0.60, 0.70, 0.70, 0.70, 0.70\}$, $q_{i8} = \{0.50, 0.60, 0.70, 0.70, 0.70\}$ and $q_{i9} = \{0.40, 0.50, 0.60, 0.70, 0.70\}$. We assumed, a priori, that each of these nine skeletons is equally likely and set $\pi(k) = 1/9$.

The previous set of true scenarios (Figure 2) investigated monotone increasing (R1) and plateau (R2, R3) dose-efficacy curves. The set of true scenarios for the sensitivity analysis assess performance for unimodal relationships. Table 3 reports the percentage of trials in which each dose was selected as the optimal dose at the end of each trial, and the average trial size for each of the Scenarios 1–4. The true (toxicity, efficacy) probabilities for each dose are given next to the number of the scenario. In Scenario 1, all doses are safe and the peak of the unimodal curve occurs at dose d_3 . In Scenario 2, all but the highest dose are safe, and the OBD occurs at d_2 . In Scenario 3, again all but d_6 are safe and the OBD occurs at dose d_5 . Finally, in Scenario 4, all doses are safe and there exists a monotone decreasing dose-efficacy curve with the OBD at d_1 . In each set of 1000 simulations, the maximum acceptable toxicity rate was specified to be $\varphi_T = 0.33$ and the efficacy threshold $\varphi_E = 0.20$.

Examination of Table 3 reveals strong performance for our method in these scenarios, recommending the OBD in a large percentage of trials ($\approx 50 - 70\%$) after $N = 48$ patients. The results in Table 3 also demonstrate the robustness of our method to the association between efficacy and toxicity. The performance of WT is practically unchanged for all

values of ψ , given that n_R remains unchanged. Therefore, even in the presence of correlated binary outcomes, the fact that our method ignores any association between toxicity and efficacy in estimating probabilities does not affect its operating characteristics. In general, the selection percentages remain fairly consistent for $n_R = 12, 24, 36,$ and 48 . All percentages in each scenario are within 10% of one another, so the chosen size of the AR phase does not drastically alter performance. These simulations represent only a small set of situations. More study, under a broader range of scenarios, may provide more insight into the general behavior of n_R .

5 Discussion

In this article, we have outlined a new Phase I/II design that accounts for the both toxicity and efficacy of molecularly targeted agents. The simulation results demonstrated the method's ability to effectively recommend the OBD, defined by acceptable toxicity and high efficacy, in a high percentage of trials with manageable sample sizes. The method we propose in this work is a bivariate extension of the CRM, leaning upon properties of order restricted inference utilized in the partial order CRM (POCRM; Wages *et al.*, 2011). This approach falls into a broader class of extended model-based designs that utilize multiple non-monotonic skeletons, which increases the ability of CRM designs to handle more complex dose-finding problems (O'Quigley and Conaway, 2011). This class of design is developed within the framework of CRM, with the added step of Bayesian model choice to account for the multiple models. Therefore, many of the features associated with CRM design specifications, such as one-parameter models and "reasonable skeletons" (O'Quigley and Zohar, 2010), can be employed in these extended designs.

The proposed design is most appropriate when both binary toxicity and binary efficacy endpoints can be observed in a reasonably similar time-frame. In some practical situations, this may not be possible due to the fact that efficacy may occur much later than toxicity. This would create a situation where we would be estimating DLT probabilities based on more patient observations than efficacy probabilities. Since we are ignoring their association in modeling these responses, we can fit the likelihood for each response based on different amounts of data, and simply utilize the efficacy data we have available, even though it may be less than that of the toxicity data. This idea requires further study and we are exploring modifications to the proposed methodology to handle such practical issues. Along these same lines, incorporating time-to-event outcomes may be an effective extension of the method as a means of handling delayed response.

References

1. 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]
2. Cheung, YK. Dose-finding by the continual reassessment method. New York: Chapman and Hall/CRC Press; 2011.
3. Conaway MR, Dunbar S, Peddada S. Designs for single- or multiple-agent phase I trials. *Biometrics*. 2004; 60:661–669. [PubMed: 15339288]
4. Dragalin V, Federov VV. Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference*. 2006; 136:1800–1823.

5. Federov, VV.; Leonov, SL. Optimal design for nonlinear response models. New York: Chapman and Hall/CRC Press; 2013.
6. Garrett-Mayer E. The continual reassessment method for dose-finding studies: a tutorial. *Clinical Trials*. 2006; 3:57–71. [PubMed: 16539090]
7. Gerber DE, Stopeck AT, Wong L, Rosen LS, Thorpe PE, Shan JS, Ibrahim NK. Phase I safety and pharmacokinetic study of baviximab, a chimeric phosphatidylserine-targeting monoclonal antibody, in patients with advanced solid tumors. *Clinical Cancer Research*. 2011; 17:6888–6896. [PubMed: 21989064]
8. Hoering A, LeBlanc M, Crowley J. Seamless phase I/II trial design for assessing toxicity and efficacy for targeted agents. *Clinical Cancer Research*. 2011; 17:640–46. [PubMed: 21135145]
9. Hoering A, Mitchell A, LeBlanc M, Crowley J. Early phase trial design for assessing several dose levels for toxicity and efficacy for targeted agents. *Clinical Trials*. 2013; 10:422–429. [PubMed: 23529697]
10. Hunsberger S, Rubinstein L, Dancy J, Korn E. Dose escalation trial designs based on a molecularly targeted endpoint. *Statistics in Medicine*. 2005; 24:2171–2181. [PubMed: 15909289]
11. Hwang J, Peddada S. Confidence interval estimation subject to order restrictions. *Annals of Statistics*. 1994; 22:67–93.
12. Iasonos A, Wilton A, Riedel E, Seshan V, Spriggs D. A comprehensive comparison of the continual reassessment method to the standard 3+3 dose escalation scheme in Phase I dose-finding studies. *Clinical Trials*. 2008; 5:465–477. [PubMed: 18827039]
13. Iasonos A, Zohar S, Quigley OJ. Incorporating lower grade toxicity information into dose finding designs. *Clinical Trials*. 2011; 8:370–79. [PubMed: 21835856]
14. Ji Y, Liu P, Li Y, Bekele BN. A modified toxicity probability interval method for dose-finding trials. *Clinical Trials*. 2010; 7:653–663. [PubMed: 20935021]
15. Ji Y, Wang SJ. Modified toxicity probability interval design: a safer and more reliable method than the 3+3 design for practical phase I trials. *Journal of Clinical Oncology*. 2013; 31:1785–1791. [PubMed: 23569307]
16. Korn E. Nontoxicity endpoints in phase I trial designs for targeted, non-cytotoxic agents. *J Natl Cancer Inst*. 2004; 96:977–978. [PubMed: 15240771]
17. Mandrekar S, Cui Y, Sargent D. An adaptive phase I design for identifying a biologically optimal dose for dual agent drug combinations. *Statistics in Medicine*. 2007; 26:2317–2330. [PubMed: 17016867]
18. Mandrekar S, Qin R, Sargent D. Model-based phase I designs incorporating toxicity and efficacy for single and dual agent drug combinations: methods and challenges. *Statistics in Medicine*. 2010; 29:1077–1083. [PubMed: 20419760]
19. 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]
20. O'Quigley J, Shen L. Continual reassessment method: a likelihood approach. *Biometrics*. 1996; 52:673–684. [PubMed: 8672707]
21. O'Quigley J, Hughes MD, Fenton T. Dose-finding designs for HIV studies. *Biometrics*. 2001; 57:1018–1029. [PubMed: 11764240]
22. O'Quigley J, Zohar S. Retrospective robustness of the continual reassessment method. *Journal of Biopharmaceutical Statistics*. 2010; 5:1013–1025. [PubMed: 20721788]
23. O'Quigley J, Conaway M. Extended model-based designs for more complex dose-finding studies. *Statistics in Medicine*. 2011; 30:2062–2069. [PubMed: 21351287]
24. Parulekar WR, Eisenhauer EA. Phase I trial design for solid tumors studies of targeted, non-cytotoxic agents: theory and practice. *J Natl Cancer Inst*. 2004; 96:970–977.
25. Polley MY, Cheung YK. Two-stage designs for dose-finding trials with a biologic endpoint using stepwise tests. *Biometrics*. 2008; 64:232–241. [PubMed: 17573866]
26. R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing; Vienna, Austria: 2013. <http://www.R-project.org>
27. Robertson, T.; Wright, F.; Dykstra, R. Order Restricted Statistical Inference. New York: Wiley; 1988.

28. Storer BE. Design and analysis of Phase I clinical trials. *Biometrics*. 1989; 45:925–937. [PubMed: 2790129]
29. Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*. 2004; 60:684–693. [PubMed: 15339291]
30. 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]
31. Wages NA, Conaway MR, O’Quigley J. Continual Reassessment Method for Partial Ordering. *Biometrics*. 2011; 67:1555–1563. [PubMed: 21361888]
32. Whitehead J, Zhou Y, Stevens J, Blakey G, Price J, Leadbetter J. Bayesian decision procedures for dose-escalation based on evidence of undesirable events and therapeutic benefit. *Statistics in Medicine*. 2006; 25:37–53. [PubMed: 16342336]
33. Yin G, Li Y, Ji Y. Bayesian dose finding in phase I/II clinical trials using toxicity and efficacy odds ratio. *Biometrics*. 2006; 62:777–787. [PubMed: 16984320]
34. Yin, G. *Clinical Trial Design: Bayesian and Frequentist Adaptive Methods*. Hoboken, New Jersey: John Wiley & Sons; 2012.
35. Yuan Z, Chappell R, Bailey H. The continual reassessment method for multiple toxicity grades: A Bayesian quasi-likelihood approach. *Biometrics*. 2007; 63:173–79. [PubMed: 17447942]
36. Zhang W, Sargent D, Mandrekar S. An adaptive dose-finding design incorporating both toxicity and efficacy. *Statistics in Medicine*. 2006; 25:2365–2383. [PubMed: 16220478]

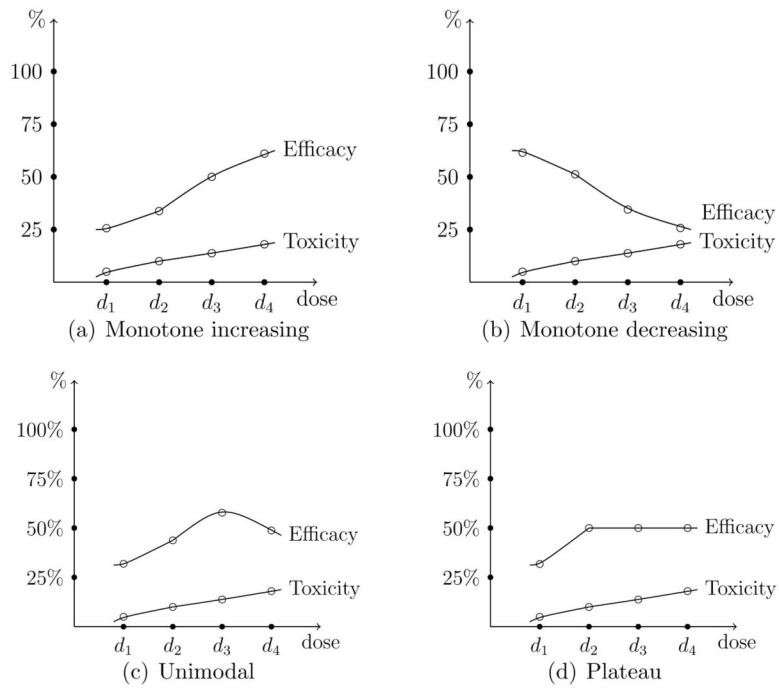


Figure 1. Possible dose-toxicity and dose-efficacy relationships for molecularly targeted agents.

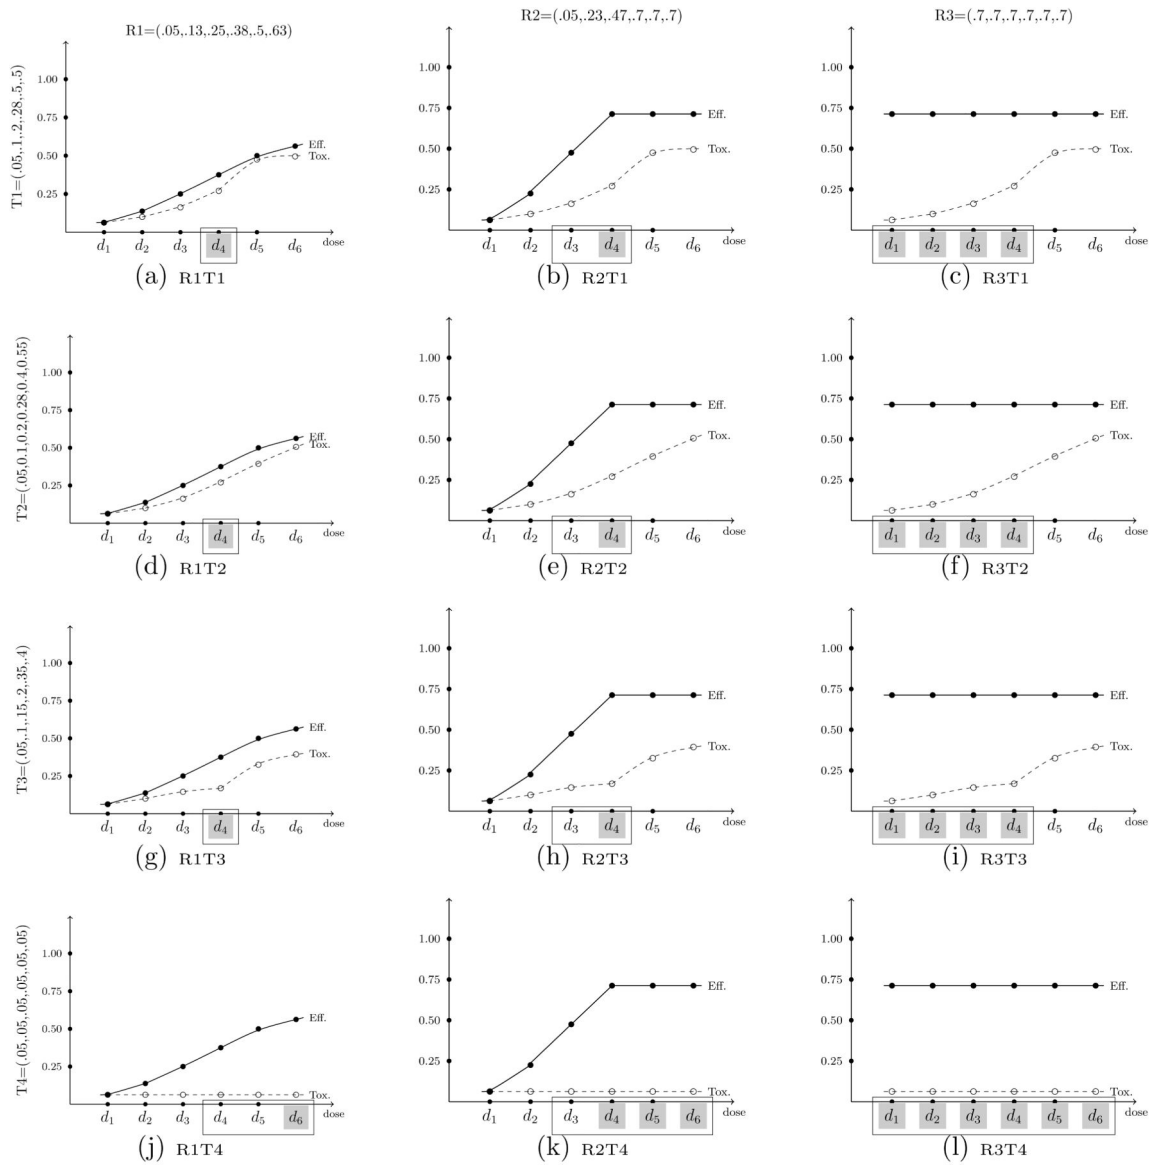


Figure 2. A set of twelve true (toxicity, efficacy) scenarios. True toxicity probabilities are denoted T1 through T4, while true efficacy probabilities are denoted R1 through R3. The **good** doses are indicated with a □, while the **best** doses are shaded in gray.

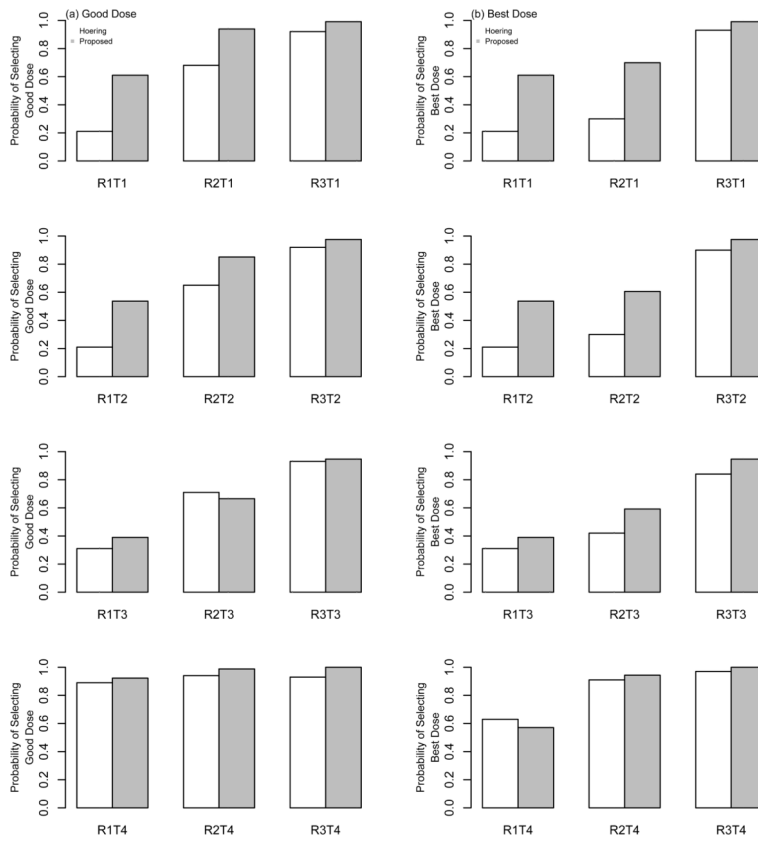


Figure 3. Operating characteristics of the proposed design and Hoering et al. (2013). Each bar represents the proportion of times that each method recommended (a) good doses and (b) best doses as the OBD at the conclusion of a simulated Phase I/II trial with a maximum sample size of $N = 64$ patients.

Table 1

Possible dose-efficacy relationships for $I = 4$ dose levels.

Shape of relationship	Location of node	Possible Dose-efficacy relationship
Unimodal or plateau	d_1	$\pi_E(d_4)$ $\pi_E(d_3)$ $\pi_E(d_2)$ $\pi_E(d_1)$
Unimodal or plateau	d_2	$\pi_E(d_1)$ $\pi_E(d_3)$ $\pi_E(d_4)$ $\pi_E(d_2)$
Unimodal or plateau	d_3	$\pi_E(d_1)$ $\pi_E(d_2)$ $\pi_E(d_4)$ $\pi_E(d_3)$
Unimodal	d_4	$\pi_E(d_1)$ $\pi_E(d_2)$ $\pi_E(d_3)$ $\pi_E(d_4)$

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

Operating characteristics of the proposed design and Hoering et al. (2013). Overall probability of selecting a best or good dose level as the OBD after 1000 simulation runs with a maximum sample size of $N = 64$ patients.

Efficacy Scenario	Toxicity Scenario	Best Dose Level	Good Dose Level	Probability of Picking Best Dose with Our Proposed Design	Probability of Picking Good Dose with Our Proposed Design	Probability of Picking Best Dose with Hoering <i>et al.</i> Design	Probability of Picking Good Dose with Hoering <i>et al.</i> Design
R1	T1	4	4	0.61	0.61	0.21	0.21
R1	T2	4	4	0.54	0.54	0.20	0.20
R1	T3	4	4	0.39	0.39	0.31	0.31
R1	T4	6	4-6	0.57	0.92	0.63	0.89
R2	T1	4	3, 4	0.70	0.94	0.30	0.68
R2	T2	4	3, 4	0.61	0.85	0.30	0.65
R2	T3	4	3, 4	0.59	0.67	0.42	0.71
R2	T4	4-6	3-6	0.94	0.99	0.91	0.94
R3	T1	1-4	1-4	0.99	0.99	0.93	0.93
R3	T2	1-4	1-4	0.98	0.98	0.90	0.90
R3	T3	1-4	1-4	0.95	0.95	0.84	0.85
R3	T4	1-6	1-6	1.00	1.00	0.97	0.97

Selection percentages for each dose under the proposed method for various design specifications. The maximum sample size for each simulated trial was $N = 48$ patients. The OBD in each scenario is indicated in bold-type.

Table 3

Scenario 1	True (toxicity, efficacy) probability					
	(0.01, 0.30)	(0.05, 0.50)	(0.10, 0.60)	(0.15, 0.40)	(0.20, 0.25)	
$\psi = 0; n_R = 12$	0.069	0.321	0.515	0.078	0.017	
$\psi = 0; n_R = 24$	0.049	0.327	0.567	0.053	0.004	
$\psi = 0; n_R = 36$	0.042	0.305	0.593	0.058	0.002	
$\psi = 0; n_R = 48$	0.062	0.285	0.587	0.060	0.006	
$\psi = -2; n_R = 12$	0.082	0.315	0.513	0.077	0.013	
$\psi = 2; n_R = 12$	0.063	0.303	0.512	0.107	0.015	
<hr/>						
Scenario 2	(0.02, 0.38)	(0.06, 0.50)	(0.12, 0.40)	(0.30, 0.30)	(0.40, 0.25)	
$\psi = 0; n_R = 12$	0.253	0.481	0.211	0.050	0.005	
$\psi = 0; n_R = 24$	0.241	0.538	0.182	0.034	0.005	
$\psi = 0; n_R = 36$	0.255	0.542	0.180	0.021	0.002	
$\psi = 0; n_R = 48$	0.265	0.537	0.158	0.035	0.005	
$\psi = -2; n_R = 12$	0.227	0.512	0.205	0.051	0.005	
$\psi = 2; n_R = 12$	0.250	0.468	0.206	0.064	0.012	
<hr/>						
Scenario 3	(0.03, 0.25)	(0.09, 0.35)	(0.16, 0.48)	(0.28, 0.65)	(0.42, 0.52)	
$\psi = 0; n_R = 12$	0.098	0.124	0.278	0.474	0.026	
$\psi = 0; n_R = 24$	0.061	0.095	0.271	0.540	0.033	
$\psi = 0; n_R = 36$	0.048	0.082	0.266	0.561	0.043	
$\psi = 0; n_R = 48$	0.061	0.061	0.215	0.539	0.124	
$\psi = -2; n_R = 12$	0.078	0.130	0.297	0.455	0.040	
$\psi = 2; n_R = 12$	0.073	0.118	0.256	0.526	0.027	
<hr/>						
Scenario 4	(0.02, 0.68)	(0.05, 0.56)	(0.07, 0.49)	(0.09, 0.40)	(0.11, 0.33)	
$\psi = 0; n_R = 12$	0.627	0.211	0.115	0.034	0.013	

Scenario 1	True (toxicity, efficacy) probability				
	(0.01, 0.30)	(0.05, 0.50)	(0.10, 0.60)	(0.15, 0.40)	(0.20, 0.25)
$\psi = 0; n_R = 24$	0.672	0.230	0.084	0.009	0.005
$\psi = 0; n_R = 36$	0.713	0.206	0.069	0.011	0.001
$\psi = 0; n_R = 48$	0.748	0.183	0.055	0.011	0.003
$\psi = -2; n_R = 12$	0.644	0.214	0.094	0.031	0.017
$\psi = 2; n_R = 12$	0.622	0.216	0.099	0.045	0.018