Published in final edited form as:
J R Stat Soc Ser C Appl Stat. 2014 January 1; 63(1): 159–173. doi:10.1111/rssc.12039.

A Bayesian Dose-finding Design for Oncology Clinical Trials of Combinational Biological Agents

Chunyan Cai,
Biostatistics/Epidemiology/Research Design Core, Center for Clinical and Translational Sciences, The University of Texas Health Science Center at Houston, Houston, TX 77030, USA

Ying Yuan, and
Department of Biostatistics, The University of Texas M. D. Anderson Cancer Center, Houston, TX 77030, USA

Yuan Ji
Cancer Research Informatics, Center for Clinical and Research and Informatics, NorthShore University HealthSystem, Evanston, IL 60201, USA

Abstract

Treating patients with novel biological agents is becoming a leading trend in oncology. Unlike cytotoxic agents, for which efficacy and toxicity monotonically increase with dose, biological agents may exhibit non-monotonic patterns in their dose-response relationships. Using a trial with two biological agents as an example, we propose a dose-finding design to identify the biologically optimal dose combination (BODC), which is defined as the dose combination of the two agents with the highest efficacy and tolerable toxicity. A change-point model is used to reflect the fact that the dose-toxicity surface of the combinational agents may plateau at higher dose levels, and a flexible logistic model is proposed to accommodate the possible non-monotonic pattern for the dose-efficacy relationship. During the trial, we continuously update the posterior estimates of toxicity and efficacy and assign patients to the most appropriate dose combination. We propose a novel dose-finding algorithm to encourage sufficient exploration of untried dose combinations in the two-dimensional space. Extensive simulation studies show that the proposed design has desirable operating characteristics in identifying the BODC under various patterns of dose-toxicity and dose-efficacy relationships.

Keywords

Biologically optimal dose combination; Non-monotonic pattern; Drug combination; Dose finding; Change-point model; Adaptive design

1. Introduction

The paradigm of oncology drug development is expanding from traditional cytotoxic agents to biological agents (Hoff and Ellis, 2007; Mandrekar and Sargent, 2009). Examples of biological agents include biospecimens targeting a specific tumor pathway, gene products aiming for DNA repair, and immunotherapies stimulating the immune system to attack a tumor. These novel agents differ from traditional cytotoxic agents in a variety of ways. For example, whereas toxicity and efficacy typically monotonically increase with the dose level

Address for correspondence: Ying Yuan, Department of Biostatistics, The University of Texas M. D. Anderson Cancer Center, Houston, TX 77030, USA. yyuan@mdanderson.org.

for cytotoxic agents, such monotonic relationships may not be true for biological agents. Specifically, toxicity may increase at low dose levels and then approximately plateau at higher dose levels when the biological agent has reached the saturation level in the body. In addition, the dose-efficacy curves for the biological agents may follow a non-monotonic pattern, and efficacy may even decrease at higher dose levels (Hoff and Ellis, 2007). Therefore, traditional dose-finding designs with a focus on finding the maximum tolerated dose (MTD) are not suitable for trials of biological agents. Novel designs that consider both the toxicity and efficacy of these agents are imperative.

Numerous designs have been proposed to find the MTD for the drug combination trials with cytotoxic agents. Kramar et al. (1999) proposed monotonically ordering of a selected subset of drug combinations so that standard dose finding method can be used. Thall et al. (2003) developed a six-parameter logistic regression model of the toxicity probability to identify an entire toxicity “contour” of drug combinations. Conaway et al. (2004) examined the simple and partial orders for drug combinations based on the pool adjacent violators algorithm. Wang and Ivanova (2005) proposed a two-stage design to identify the MTD combinations based on a loglinear model for toxicity probabilities. Yuan and Yin (2008) proposed a sequential dose-finding design that allows single-agent dose-finding methods to be used in multiple-agent combination trials. Yin and Yuan (2009) proposed a Bayesian dose-finding design based on a copula-type regression model. Braun and Wang (2010) proposed a hierarchical model that explicitly accounts for patient heterogeneity. Recently, Wages et al. (2011) extended the continual reassessment method (CRM) to two-dimensional dose finding by converting a partially ordered two-dimensional dose space into a series of fully ordered dose sequences. In contrast to this rich body of literature for cytotoxic agents, published research on drug combination trial designs for biological agents has been very limited. Mandrekar et al. (2007) proposed a dose-finding design for trials evaluating com-binational biological agents based on a continuation ratio model, but that design requires collapsing binary toxicity and efficacy outcomes into a trinary outcome.

Our research is motivated by a drug combination trial at The University of Texas MD Anderson Cancer Center for patients diagnosed with relapsed lymphoma. The trial combined two novel biological agents, A and B (their names are masked to maintain confidentiality), that target two different components in the PI3K/AKT/mTOR signaling pathway. This pathway has been associated with several genetic aberrations related to the promotion of cancer (Ihle and Powis, 2009). Agent A is a PI3K kinase inhibitor and agent B is a downstream inhibitor of mTOR kinase within that pathway. Research has suggested that some types of lymphomas are promoted and maintained by the activation of the PI3K/AKT/mTOR pathway, making the pathway an important target for drug development (Smith et al., 2010). Both agents A and B have individually demonstrated a partial inhibition of the pathway and some therapeutic activity. By combining these two agents, the investigators expect to obtain a more complete inhibition of the PI3k/AKT/mTOR pathway, and thereby to achieve better treatment responses. The trial investigates the combinations of 4 dose levels of agent A with 4 dose levels of agent B, which results in 16 dose combinations. The goal is to find the *biologically optimal dose combination* (BODC), defined as the dose combination with the highest efficacy and tolerable toxicity (e.g., with a toxicity probability < 0.3). We note that, depending on the clinical setting, other defi-nitions of BODC (e.g., based on a toxicity-efficacy tradeoff) may be more appropriate for different clinical trials. For this trial, the physicians expect the toxicity of the combinations to increase at low doses and become (approximately) flat at high doses, and they consider the possibility that the dose-efficacy curve of the combinations may be non-monotonic (i.e., the dose with the highest efficacy is not necessarily the highest dose).

We introduce a dose-finding design to identify the BODC for oncology trials of combinational biological agents. The proposed design explicitly accounts for the unique properties of biological agents. We propose a change-point model to reflect the property that the dose-toxicity surface of the combinational agents may plateau at higher dose levels, and use a general logistic model with quadratic terms to accommodate the possible non-monotonic pattern of the dose-efficacy relationship. Our design is conducted in two stages: in stage I, we escalate doses along the diagonal of the dose combination matrix as a fast exploration of the dosing space; in stage II, based on the observed toxicity and efficacy data from stages I and II, we continuously update the posterior estimates of toxicity and efficacy and assign patients to the most appropriate dose combination. We propose a novel dose-finding algorithm to encourage sufficient exploration of the two-dimensional dose space, which facilitates the identification of the BODC. Extensive simulation studies show that the proposed design has desirable operating characteristics in identifying the BODC under various patterns of dose-toxicity and dose-efficacy relationships.

The remainder of this paper is organized as follows. In Sections 2 and 3, we introduce the probability models and the dose-finding design for finding the BODC. In Section 4 we apply our design to the lymphoma clinical trial and examine the design's operating characteristics through extensive simulation studies and sensitivity analysis. We conclude with a brief discussion in Section 5.

2. Methods

2.1. Modeling Toxicity and Efficacy

Consider a trial combining J doses of biological agent A, denoted by $a_1 < a_2 < \cdots < a_J$, with K doses of biological agent B, denoted by $b_1 < b_2 < \cdots < b_K$. Without loss of generality, we assume $J \geq K$ and that the dose values of the a_j's and b_k's have been standardized to have mean 0 and standard deviation of 0.5. This standardization is used to anticipate the prior elicitation in Section 2.2. Let (a_j, b_k) denote the combination of dose a_j and dose b_k, and let p_{jk} and q_{jk} denote the toxicity and efficacy probabilities of (a_j, b_k), respectively, for $j = 1, 2, \ldots, J$, and $k = 1, 2, \ldots, K$. Here, toxicity and efficacy are two binary events that reflect the side effects (toxicity) and therapeutic effects (efficacy) of the biological agents. Therefore, p_{jk} and q_{jk} are simply the probabilities of the toxicity event and efficacy event, respectively, at dose combination (a_j, b_k). The goal of the trial is to identify the BODC in the $J \times K$ dose combination matrix.

2.1.1. Dose-toxicity Model—Previous research has shown that, for the purpose of dose finding, as data are observed only at the discrete doses prespecified in the trial, the choice of the dose-toxicity model is not critical as long as the model is (i) adequately flexible to capture the basic feature of the dose-response curve and (ii) reasonably parsimonious to accommodate small sample sizes of dose-finding trials (O'Quigley et al., 1990; Paoletti and Kramar, 2009). When modeling the dose-toxicity relationship for biological agents, the basic feature that needs to be taken into account is that the dose-toxicity curve may initially increase at low doses and then plateau at high doses. In this article, we consider two candidate dose-toxicity models that can capture this feature of biological agents. As we show later, these two models work equally well and yield very similar operating characteristics. The first model is the change-point model of the form

$$\text{logit}(p_{jk}) = (\beta_0+\beta_1 a_j+\beta_2 b_k)\, I\,(\beta_0+\beta_1 a_j+\beta_2 b_k \leq\ \omega) +\omega I\,(\beta_0+\beta_1 a_j+\beta_2 b_k > \omega), \quad (1)$$

where $I(\cdot)$ is the indicator function and $\beta_0, \beta_1, \beta_2$ and ω are unknown parameters. Under this model, the shape of the dose-toxicity surface initially is monotone with the dose level but

changes to flat once it passes the threshold defined by $\beta_0 + \beta_1 a_j + \beta_2 b_k = \omega$ (see Fig. 1). We assume that $\beta_1 > 0$ and $\beta_2 > 0$ such that the toxicity probability initially increases with the doses of A and B before it plateaus, at which time the toxicity probability is given by $e^{\omega}/(1 + e^{\omega})$.

The second candidate model we considered for the dose-toxicity curve is the scaled-logistic model given by

$$\text{logit}\left(\frac{p_{jk}}{\rho}\right) = \beta_0 + \beta_1 a_j + \beta_2 b_k, \quad (2)$$

where $\beta_1 > 0$, $\beta_2 > 0$ and β_0 are regression parameters, and $0 < \rho \leq 1$ is the scale parameter. Under this model, the toxicity probability first increases with the dose levels of agents A and B and then eventually plateaus at the toxicity probability ρ (see Fig. 1).

The change-point model and the scaled-logistic model are parsimonious. They both have 4 unknown parameters, just one parameter more than the standard logistic model (with two covariates). The extra parameter (i.e., ω or ρ) is used to capture the potential plateau behavior of the dose-toxicity curve. In these two models, we do not include an interactive effect for the two agents (e.g., an interaction term $\beta_3 a_j b_k$) because the reliable estimation of such an interaction term requires a large sample size (e.g., a few hundreds), which is typically not available in phase I trials. For the purpose of dose finding, we do not seek to model the entire dose-toxicity surface but aim to obtain an adequate local fit to facilitate dose escalation and de-escalation. A model may provide a poor global fit to the entire dose-toxicity surface; however, as long as the model provides a good local fit around the current dose, it will lead to correct decisions of dose escalation and selection. In the context of drug combination trials, Wang and Ivanova (2005) found that a model without interaction performed as well as one with interaction for dose finding.

2.1.2. Dose-efficacy Model—For biological agents, the dose-efficacy curve often follows a non-monotonic pattern. For example, in immunotherapy trials, the dose-efficacy relationship could be bell-shaped. That is, the most effective dose may be a dose in the middle of the therapeutic dose ranges, and when a dose level is lower or higher than the most effective dose, efficacy decreases. To incorporate such a non-monotonic pattern for the dose-efficacy relationship, we assume that the efficacy probability of (a_j, b_k), that is, q_{jk}, follows a logistic model of the form

$$\text{logit}(q_{jk}) = \gamma_0 + \gamma_1 a_j + \gamma_2 b_k + \gamma_3 a_j^2 + \gamma_4 b_k^2, \quad (3)$$

where $\gamma_0, \ldots, \gamma_4$ are unknown parameters. The quadratic terms render the model adequate flexibility to capture the non-monotonic shape of the dose-efficacy surface. In this dose-efficacy model, we exclude the interaction effect $a_j b_k$ for the same reason described previously.

2.2. Likelihood and Prior Specification

Suppose that at a certain stage of the trial, among n_{jk} patients treated at the paired dose (a_j, b_k), x_{jk} and y_{jk} patients have experienced dose-limiting toxicity and efficacy, respectively, where $j = 1, \cdots, J$ and $k = 1, \cdots, K$. Let $= \{\beta_0, \beta_1, \beta_2\}$ and $\gamma = \{\gamma_0, \gamma_1, \gamma_2, \gamma_3, \gamma_4\}$ denote the regression coefficients in models (1) and (3). The likelihood function of the observed data $D = \{x_{jk}, y_{jk}\}$ can be expressed as

$$L\ (\mathscr{D}|\omega,\ \beta,\ \gamma) \propto \prod_{j=1}^{J}\prod_{k=1}^{K} p_{jk}^{x_{jk}}\ (1-p_{jk})^{n_{jk}-x_{jk}} \times q_{jk}^{y_{jk}}\ (1-q_{jk})^{n_{jk}-y_{jk}}.$$

Let $f(\omega)$, $f(\beta)$, and $f(\gamma)$ denote the prior distributions for ω, β, and γ, respectively. Assuming prior independence among ω, β, and γ, we write the joint posterior distribution as

$$f\,(\omega,\beta\,,\gamma\ \,|\mathscr{D}) \propto L\,(\mathscr{D}|\,\omega,\ \beta,\gamma\,)\,f\,(\omega)\,f\,(\beta)\,f\,(\gamma)\,,$$

from which the full conditional distributions can be obtained. The Gibbs sampler will be used to obtain posterior draws of unknown parameters for statistical inferences.

For the prior specification of the efficacy model, we assign γ a weakly informative default prior recommended by Gelman et al. (2008). That is, $\gamma_0 \sim$ *Cauchy*(0, 10), and $\gamma_1, \cdots, \gamma_4 \sim$ *Cauchy*(0, 2.5), where *Cauchy*(*c*, *d*) denotes a Cauchy distribution with the center parameter *c* and the scale parameter *d*, assuming that doses a_j and b_k have been standardized with mean 0 and standard deviation 0.5. The advantages of using these priors include that they are weakly informative and also appropriately regularized such that a dramatic change in efficacy probability (e.g., from 0.01 to 0.5) is unlikely when dose changes by one level. Consequently, using the weakly informative default priors improves the estimation stability while still being vague enough to ensure that the data are able to dominate the priors (Gelman et al., 2008).

For the toxicity models (1) and (2), we use the weakly informative default prior *Cauchy*(0, 10) for intercept β_0. We assign β_1 and β_2 independent gamma prior distributions with the shape parameter of 0.5 and the rate parameter of 0.5 to ensure the monotonicity before the dose-toxicity surface reaches the plateau. To specify a prior for ω in model (1), we assume that the toxicity probability at the plateau is between 0.2 and 0.8, which corresponds to a value of ω ranging from -1.39 to 1.39. Thus, we assign ω a normal prior $N(0, 4)$, which provides sufficient coverage for all plausible toxicity probabilities at the plateau. We specify a uniform prior on (0, 1] for the scale parameter ρ in the scaled-logistic model (2).

3. Trial design

The proposed dose-finding design consists of two stages. Stage I is a run-in period, in which the goal is to explore the dose-combination space quickly and collect preliminary data so that the proposed probability models can be reliably estimated in stage II for systematic dose finding. We start stage I of the design by treating the first cohort of patients at the lowest dose combination (a_1, b_1), and then escalate the dose along the diagonal of the dose combination matrix until we encounter a dose combination that violates the safety requirement

$$Pr\,(p_{jk} < \phi|\mathscr{D}) > \delta, \quad (4)$$

where φ denotes the target toxicity upper limit and δ is a prespecified safety cutoff. If the dose matrix is not square (i.e., $J > K$), we first escalate the dose along the diagonal from (a_1, b_1) to (a_2, b_2) and so on until we reach (a_K, b_K); thereafter we escalate the dose by holding the dose level of B at K and increasing the dose level of A from (a_K, b_K) to (a_{K+1}, b_K) and so on until we reach the highest dose combination (a_J, b_K). In stage I, only a small fraction of

patients are enrolled into the trial and the observed data are sparse. Therefore, in this stage, we evaluate the safety requirement based on a simple beta-binomial model rather than the proposed dose-toxicity models. Specifically, we assume that the number of toxicities x_{jk} follows a binomial distribution $Bi(n_{jk}, p_{jk})$, and that the toxicity probability p_{jk} follows a beta distribution $Beta(\zeta, \xi)$ with two shape parameters ζ and ξ. To ensure that the data dominate the posterior distribution, we set ζ=0.1 and ξ=0.2. Under the beta-binomial model, $Pr(p_{jk} < \varphi|\mathscr{D}) = \mathscr{B}(\varphi|\zeta + x_{jk}, \xi + n_{jk} - x_{jk})$, where $\mathscr{B}(\cdot)$ is the cumulative density function for a beta distribution. In stage I we also collect efficacy data; however, these data will not be used to determine the dose escalation. Whenever a dose combination (a_j, b_k) violates the safety requirement, i.e. $Pr(p_{jk} < \varphi|\mathscr{D}) \leq \delta$, or we reach the highest dose combination (a_J, b_K), stage I is complete and the trial moves on to stage II. Under this conservative stage I completion rule, the maximum number of patients allocated to the dose that violates the safety requirement will be not more than one cohort, which is comparable to the conventional 3+3 design. In the conventional 3+3 design, we treat patients in cohort of size 3 based on the following dose transition rule: escalate the dose if 0 out of 3 patients experiences toxicity, deescalate the dose if 2 or more out of 3 patients experience toxicity, and treat 3 more patients at the current dose if 1 out of 3 patients experiences toxicity. Therefore, stage I does not pose particular safety concerns although we escalate the dose along the diagonal.

In stage II of the trial, we invoke the toxicity and efficacy models described in Section 2 for systematic dose finding. Stage II dose finding is highlighted by two features. First, the proposed algorithm encourages the exploration of untried dose combinations to avoid the problem of trapping in suboptimal doses, which is of particular concern for combinations of biological agents. Because of complex drug-drug interactions and non-monotonic dose-response patterns, the assumed (simple) dose-response model is not expected to estimate the dose-response surface well, especially at the beginning of the trial when only a few observations are available. Consequently, the resulting estimates of efficacy and toxicity may substantially deviate from the truth, and the "optimal" dose identified based on these estimates may actually be a suboptimal dose. In other words, the dose with the highest estimate of efficacy is not necessarily the one actually having the highest efficacy. By intentionally visiting untried dose combinations, the proposed method increases the chance of finding better combinations and avoids trapping in suboptimal doses. Second, we introduce a concept of *g-degree neighbor* and *g-degree admissible neighbor* to restrict the dose escalation/deescalation within the neighbors of the current dose, which avoids dramatic dose changes and improves the reliability of the dose finding.

Specifically, given the current dose combination (a_j, b_k), we define *g-degree neighbors* of (a_j, b_k), denoted by $\mathscr{N}_g$, as dose combinations $\{(a_{j'}, b_{k'})\}$ whose dose levels are different from (a_j, b_k) no more than g levels, i.e., $\mathscr{N}_g = \{(a_{j'}, b_{k'}) : |j' - j| \leq g \text{ and } |k' - k| \leq g\}$. Note that the dose set of $\mathscr{N}_g$ includes the current dose combination itself. We further define a g-degree admissible dose set $\mathscr{A}_g = \{(a_{j'}, b_{k'}) : (a_{j'}, b_{k'}) \in \mathscr{N}_g, Pr(p_{j'k'} < \varphi_T|\mathscr{D}) > \delta\}$, which is a subset of the g-degree neighbors $\mathscr{N}_g$ satisfying the pre-specified safety requirement $Pr(p_{j'k'} < \varphi_T|D) > \delta$. That is, $\mathscr{A}_g$ contains the safe g-degree neighbors of the dose combination (a_j, b_k).

Let N denote the prespecified maximum sample size, N_1 denote the number of patients in stage I, and $N_2 = N - N_1$ be the total number of patients available for stage II. Then the proposed dose-finding algorithm for stage II is described as follows:

a. Based on the accumulated trial data, we determine the dose set $\mathscr{A}_{g^*}$, where $g^* = \min\{g: \mathscr{A}_g \neq \emptyset, g \geq 1\}$. That is, $\mathscr{A}_{g^*}$ is the nonempty admissible set with the smallest degree g^*. If $\mathscr{A}_{g^*}$ does not exist, i.e., all investigational doses violate the safety requirement, we terminate the trial.

b. In $\mathscr{A}_{g*}$, we identify the combination (a_{j*}, b_{k*}) that has the highest posterior mean of efficacy rate $\widehat{q}_{j*k*}$ under the safety constraint $(j^* - j) + (k^* - k) \leq 1$ (i.e., the total dose escalation for two agents cannot more than one level).

c. If combination (a_{j*}, b_{k*}) has not been used to treat any patient thus far, or all doses in $\mathscr{A}_{g*}$ have been used to treat patients, we assign the next cohort of patients to (a_{j*}, b_{k*}). However, if (a_{j*}, b_{k*}) has been used to treat patients and there are some untried doses in $\mathscr{A}_{g*}$, we assign the next cohort of patients to (a_{j*}, b_{k*}) only if

$$\widehat{q}_{j*k*} > \left(\frac{N_2 - n_2}{N_2}\right)^{\alpha}, \quad (5)$$

where n_2 is the total number of patients that have been treated in stage II and a is a known tuning parameter controlling how stringent the threshold is; otherwise, (a_{j*}, b_{k*}) will be excluded from the admissible set $\mathscr{A}_{g*}$ and we return to step 2.

d. We continue the above steps until exhaustion of the sample size, and select as the BODC the dose combination with the highest value of $\hat{q}_{jk}$ and satisfying the safety requirement $Pr(p_{jk} < \varphi | \mathscr{D}) > \delta$.

Remark 1

The dose assignment rule (5) plays a key role in adaptively encouraging the exploration of untried doses and avoiding the problem of trapping in suboptimal doses during dose finding. At early phase of stage II, n_2 is small and thus the value of $\{(N_2 - n_2)/N_2\}^{\alpha}$ is close to 1. Consequently, rule (5) strongly encourages the exploration of untried doses. This is a sensible action because at the beginning of stage II the efficacy estimate $\widehat{q}_{jk}$ is of large variability, and we should give high priority to using new doses rather than putting too much faith in the point estimate $\widehat{q}_{jk}$ Toward the end of the trial (i.e., $n_2 \approx N_2$), we have more precise estimates of $\widehat{q}_{jk}$ based on the accumulated data. As $\{(N_2 - n_2)/N_2\}^{\alpha}$ approaches 0, we essentially assign incoming patients to the dose combination with the highest value of $\hat{q}_{jk}$ because rule (5) is almost always satisfied. In rule (5), the tuning parameter a controls how fast $\{(N_2 - n_2)/N_2\}^{\alpha}$ decays from 1 to 0. The value of a can be calibrated to obtain desirable operating characteristics.

The rule (5) provides a pragmatic strategy to balance two competing goals of the trial: to achieve high statistical power and to maintain strong ethic for patients. On one hand, to achieve high power of identifying the BODC, we should distribute patients evenly among dose combinations to learn the efficacy and toxicity profile of the entire dosing surface; and on the other hand, to maintain strong ethic, we should concentrate patients at the dose with the highest efficacy and lowest toxicity. The proposed adaptive rule balances the two goals. It renders the design to gain power rapidly at the beginning of the trial, and then, after reaching adequate power (and more reliable estimates), to focus on allocating patients to the best dose. A formal approach would be to specify a utility function as the tradeoff between the statistical power and ethic, and then optimize the dose assignment according to the value of the utility. Nevertheless, the specification of such utility could be also subjective and varying by case.

We summarize both stages of the proposed design in Box 1.

Box 1

The proposed algorithm for finding BODC

The trial starts with the treatment of the first cohort of patients at the lowest dose (a_1, b_1). Suppose that patients are being treated at dose (a_j, b_k). A dose is safe if $Pr(p_{jk} < \phi|\mathcal{D}) > \delta$; otherwise, the dose is deemed toxic.

Stage I Run-in Period

I1 If dose (a_j, b_k) is safe, escalate the dose and treat the next cohort at (a_{j+1}, b_{k+1}). If $j = k = K$, escalate the dose to (a_{j+1}, b_K). If (a_1, b_1) is deemed toxic, terminate the trial.

I2 Stage I is complete when either dose (a_j, b_k) is deemed toxic or the highest dose combination (a_J, b_K) is reached. Stage II then starts.

Stage II Systematic Dose Finding

II1 Based on the observed data, identify $\mathscr{A}_{g^*}$ as the nonempty set of safe neighbors of (a_j, b_k) with minimum degree g^*. If $\mathscr{A}_{g^*}$ does not exist (i.e., all experimental doses are deemed toxic), terminate the trial.

II2 Among the doses in $\mathscr{A}_{g^*}$, identify the dose (a_{j^*}, b_{k^*}) with the highest posterior mean of efficacy $\widehat{q}_{j^*k^*}$ under the safety constraint $(j^* - j) + (k^* - k) \le 1$.

II3

(i) If $n_j{}^*k^* = 0$ or $n_{rs} \neq 0$ for all $(a_r, b_s) \in \mathscr{A}_{g^*}$, treat the next cohort at dose (a_{j^*}, b_{k^*}).

(ii)

$$\text{Otherwise,} \begin{cases} If\ \widehat{q}_{j^*k^*} > \left(\frac{N_2 - n_2}{N_2}\right)^{\alpha} & \text{treat the next cohort at } (a_{j^*},\ b_{k^*}), \\ If\ \widehat{q}_{j^*k^*} \le \left(\frac{N_2 - n_2}{N_2}\right)^{\alpha} & \text{remove dose } (a_{j^*},\ b_{k^*}) \text{ from } \mathscr{A}_{g^*} \text{and go to step II2.} \end{cases}$$

II4 Repeat steps II2-4 until exhaustion of the sample size. Select as the BODC the dose combination with the highest $\widehat{q}_{jk}$ among all safe doses.

4. Numerical Studies

4.1. Operating Characteristics

We conducted extensive simulations to evaluate the operating characteristics of the proposed dose-finding design. We examined three major components of the design, including the dose-toxicity model, the dose-assignment rule and the start-up rule, by comparing the proposed design (with the change-point dose-toxicity model) to three alternative designs. These alternative designs were obtained by replacing each of the components with an alternative approach. Specifically, we considered (i) a scaled-logistic design, in which we used the scaled-logistic model (2) rather than the change-point model (1) to model the dose-toxicity relationship. This comparison evaluates the sensitivity of the proposed design to the alternative model specification; (ii) a "greedy" design, in which we replaced the proposed adaptive dose assignment rule (described in Step II3 in our design) with a "greedy" dose-assignment rule that always assigns patients to the dose with the highest estimate of efficacy. Technically, this means that the greedy design replaces the rule (5) with $\widehat{q}_{j^*k^*} > 0$, so that the dose with the highest efficacy among admissible dose set $\mathscr{A}_{g^{**}}$ is always selected; (iii) a

zone-based design, in which we replaced our start-up rule (i.e., dose escalation along diagonal) with the zone-based start-up rule (Huang et al., 2007; Wages et al., 2011). The goal of this comparison is to evaluate the performance of the proposed start-up rule. As shown in Fig. 2, the basic idea of the zone-based start-up rule is that, based on the partial order of the dose-toxicity relationship for the drug combinations, we can divide the dose combination matrix into zones and conduct dose escalation among zones (i.e., if the doses in one zone are deemed safe, we escalate to the next higher zone). For doses within zones, as their toxicity probabilities are not ordered, patients are equally allocated to these doses. The details of the zone-based start-up rule can be found in Huang et al. (2007) and Wages et al. (2011).

We considered trials combining two biological agents, A and B, each with 4 dose levels. The maximum sample size was 45 and patients were treated in cohorts of size 3. We set the toxicity upper limit $\varphi = 0.3$. In the proposed design, we set the safety cutoff $\delta = 0.4$ and the tuning parameter $\alpha = 2$, and used 2,000 posterior samples of unknown parameters ω, β, and γ to make inference after 1,000 burn-in iterations. We investigated 6 different dose-toxicity and dose-efficacy scenarios, as shown in scenarios 1-6 in Table 1 (scenarios 7-12 were prepared for the sensitivity analysis). These scenarios were arbitrary constructed (not generated from the proposed dose-toxicity and dose-efficacy models) to simulate various shapes of the dose-toxicity and dose-efficacy relationships. In scenario 1, the dose-toxicity surface initially increases with the dose levels of agents A and B and then plateaus at some higher doses of the dose combination matrix with a toxicity probability of 0.25; while the dose-efficacy relationship is non-monotonic, characterized by efficacy monotonically increasing with agent A but not with agent B. Scenario 2 shares the same dose-toxicity profile with scenario 1 but possesses a different shape of the dose-efficacy surface. Scenario 3 shares the same dose-efficacy profile with scenario 1 but assumes monotonic dose-toxicity relationship for both agents A and B. In scenario 4, the dose-toxicity surface plateaus at the right upper corner of the dose combination matrix with a toxicity probability of 0.25 and efficacy monotonically increases with dose levels for both two agents. In scenario 5, the dose-toxicity surface becomes flat at high dose levels, which are overly toxic. In scenario 6, we consider the case in which all the doses at the right upper corner of the matrix are toxic. Under each scenario, we carried out 2,000 simulated trials for each of the designs. We used C++ to implement the proposed design; the simulation code is available for access via the journal web site.

The simulation results for scenarios 1-6 are summarized in Table 2, including the selection percentage of the BODC, the percentage of patients allocated to the BODC, the average efficacy rate, the number of patients assigned to over-toxic doses, and the total numbers of patients assigned in stage I and stage II of the trial. The detailed results including the selection percentage of each dose combination and the percentage of patients assigned to each dose combination are provided in the supplementary materials.

In scenario 1, the target BODC is (a_4, b_2). Among the four designs, the proposed design and scaled-logistic design performed best with the comparable selection percentages of the BODC (36.2% versus 37.0%), the same percentage of patients (17.3%) allocated to the target dose combination, and similar average efficacy rates (33.3% versus 32.9%). These results suggest that the choice of the dose-toxicity model is not critical as long as the model is able to capture the basic feature of the dose-toxicity relationship. The greedy design performed poorly. Compared to the proposed design, its selection percentage of the BODC was 12.1% lower, with 7.1% less patients allocated to the BODC. The poor performance of the greedy design was due to the use of the greedy dose assignment rule, which often caused the design to be trapped at the suboptimal doses (such as (a_4, b_4), see Table 1 in the supplementary materials). This result demonstrates that the proposed dose assignment rule

(which adaptively encourages new dose exploration) is useful to avoid trapping in suboptimal doses and improve the overall performance of the trial design. The zone-based design performed well in selecting the BODC, but poorly in terms of allocating patients to the BODC and the average efficacy rate. The percentage of patients allocated to the BODC under the zone-based design was less than one half of that of the proposed design. The poor performance of the zone-based design was caused by the use of the zone-based start-up rule, which requires assigning patients to all dose combinations in a zone before escalating to the next higher zone. As a result, too many patients were assigned to futile doses. This is confirmed by the results that almost 14 cohorts of patients were assigned in stage I (i.e., the start-up phase) and less than two cohorts of patients were assigned in stage II when the zone-based design was used.

Similar results are observed in the remaining five scenarios. That is, the proposed design and the scaled-logistic design were comparable and yielded the best performance. The greedy design was often trapped in suboptimal doses and led to low selection percentages of the BODC and low percentages of patients allocated to the BODC. In some scenarios, such as scenarios 4 and 6, the greedy design performed very well, but that is because the target dose combination happens to be on the dose-escalation path of the start-up phase (i.e., the diagonal of the dose matrix), under which the exploration of the dose space is not actually needed. Zone-based design yielded high selection percentages of the BODC comparable to the proposed design, but performed poorly in assigning patients to the BODC and efficacious doses. It tended to use too many patients at the start-up phase and allocate patients to futile doses.

4.2. Sensitivity Analysis

To further evaluate the robustness of the proposed design, we considered 6 additional dose-toxicity and dose-efficacy scenarios (i.e., scenarios 7-12 in Table 1), which were systematically generated based on probit models. Specifically, we generated the true toxicity probability from $\Phi^{-1}(p_{jk}/\rho) = \beta_0 + \beta_1 a_j + \beta_2 b_k$, where $\Phi^{-1}(\cdot)$ is the inverse cumulative distribution function of the standard normal distribution; and the true efficacy probability from $\Phi^{-1}(q_{jk}) = \gamma_0 + \gamma_1 a_j + \gamma_2 b_k + \gamma_3 a_j^2 + \gamma_4 b_k^2$ The simulation results (Table 3) show that the proposed design consistently outperformed the greedy and zone-based designs with the highest selection percentage of the BODC and the highest percentage of patients allocated to the BODC. Again, the scaled-logistic design exhibited similar operating characteristics as the proposed design, which further confirms that the proposed change-point model and scaled-logistic model worked equally well for the purpose of dose finding.

We model toxicity and efficacy outcomes marginally without accounting for the correlation between these two endpoints. To assess the robustness of our design to dependence between toxicity and efficacy, we conducted a sensitivity analysis with correlated toxicity-efficacy data. We generated correlated toxicity and efficacy data (x_{jk}, y_{jk}) based on a bivariate Gumbel model (Murtaugh and Fisher, 1990)

$$\Pr(x_{jk}=a,\ y_{jk}=b) = p_{jk}^{a}(1-p_{jk})^{1-a} q_{jk}^{b}(1-q_{jk})^{1-b} + (-1)^{a+b} p_{jk}(1-p_{jk})\, q_{jk}(1-q_{jk})\, \frac{e^{\psi}-1}{e^{\psi}+1}$$

where $a, b \in \{0, 1\}$ and ψ is a real-valued parameter controlling the correlation between toxicity and efficacy; and then applied our methods (modeling toxicity and efficacy outcomes independently) to the simulated correlated data. In the Gumbel model, we set $\psi = 3$ to induce a high correlation between the toxicity and efficacy endpoints, and for ease of

comparison we matched the marginal toxicity and efficacy probabilities, p_{jk} and q_{jk}, to those in scenarios 1 through 6 listed in Table 1. The results with the correlated toxicity and efficacy data (see Table 3 in the supplementary materials) were very similar to their counterparts with independent data (see Table 1 in the supplementary materials). The differences in the selection probabilities of the BODC were mostly less than 3% in all scenarios. These results suggest that the independence assumption between toxicity and efficacy has negligible effects on the performance of our design.

5. Conclusions

We proposed a new Bayesian dose-finding design for trials that evaluate combinational biological agents. The proposed design explicitly accounts for the unique properties of biological agents. A change-point model is used to capture the feature that the dose-toxicity surface of biological agents may plateau at high dose levels, and a second-order logistic model is employed to accommodate non-monotonic patterns for the dose-efficacy relationship. We proposed a novel dose-finding algorithm that adaptively encourages the exploration of two-dimensional dose-toxicity and dose-efficacy surfaces during dose finding. In the early stage of the trial, the algorithm gives higher priority to trying new doses, and toward the end of the trial it assigns patients to the most effective dose that is safe. Extensive simulations show that the proposed design has good operating characteristics with a high probability of selecting the BODC.

The proposed design is suitable for trials with endpoints that can be quickly evaluated. In the case that toxicity or/and efficacy take a relatively long time to be assessed, one could consider treating patients in cohorts (of size 3, for example) so that the accrual does not need to be halted as often to wait patients' outcomes to become known. Alternatively, we can accommodate such delayed outcomes in a more systematic way by taking the missing data approach (Yuan and Yin, 2011). In addition, as other adaptive designs, our method requires continuously updating the posterior estimates whenever a new response is observed. However, this is not a concern in practical deployment of our design because fitting the proposed models (based on Markov Chain Monte Carlo) only takes less than one minute. In the proposed design, we are interested in finding the dose with highest efficacy and tolerable toxicity as the target BODC. Our design can be easily extended to the case that the target BODC is defined by a certain toxicity-efficacy trade-off function. In that case, the main exercise is to elicit a reasonable toxicity-efficacy trade-off (or utility) function from clinicians (Thall and Cook, 2004). Once the trade-off is defined, our design can be directly applied by replacing efficacy with the trade-off as the criteria of dose escalation and selection.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors would like to thank the editor, associate editor and two reviewers for very insightful and constructive comments that substantially improved the paper. Yuan and Cai's research was supported by Award Number R01 CA154591 from the National Cancer Institute. Yuan's research was also supported by SPORE for Uterine Cancers (P50 CA098258-08). Cai's research was also supported by the National Institutes of Health's Clinical and Translational Science Award grant (UL1 TR000371), awarded to the University of Texas Health Science Center at Houston in 2012 by the National Center for Clinical and Translational Sciences. Article content is solely the responsibility of the authors and does not necessarily represent the official views of the National Cancer Institute or the National Institutes of Health. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper. URL: http://www.tacc.utexas.edu.

References

Braun TM, Wang S. A hierarchical Bayesian design for phase I trials of novel combinations of cancer therapeutic agents. Biometrics. 2010; 66:805–812. [PubMed: 19995354]

Conaway MR, Dunbar S, Peddada SD. Designs for single- or multiple-agent phase i trials. Biometrics. 2004; 60:661–669. [PubMed: 15339288]

Gelman A, Jakulin A, Pittau MG, Su YS. A weakly informative default prior distribution for logistic and other regression models. The Annals of Applied Statistics. 2008; 2:1360–1383.

Hoff PM, Ellis LM. Targeted therapy trials: approval strategies, target validation, or helping patients? Journal of Clinical Oncology. 2007; 25:1639–1641. [PubMed: 17470854]

Huang X, Biswas S, Oki Y, Issa JP, Berry DA. A parallel phase I/II clinical trial design for combination therapies. Biometrics. 2007; 63:429–436. [PubMed: 17688495]

Ihle NT, Powis G. Take your PIK: phosphatidylinositol 3-kinase inhibitors race through the clinic and toward cancer therapy. Molecular Cancer Therapeutics. 2009; 8:1–9. [PubMed: 19139107]

Kramar A, Lebecq A, Candalh E. Continual reassessment methods in phase I trials of the combination of two drugs in oncology. Statistics in Medicine. 1999; 18:1849–1864. [PubMed: 10407256]

Mandrekar SJ, Cui Y, Sargent DJ. An adaptive phase I design for identifying a biologically optimal dose for dual agent drug combinations. Statistics in Medicine. 2007; 26:2317–2330. [PubMed: 17016867]

Mandrekar SJ, Sargent DJ. Clinical trial designs for predictive biomarker validation: Theoretical considerations and practical challenges. Journal of Clinical Oncology. 2009; 27:4027–4034. [PubMed: 19597023]

Murtaugh PA, Fisher LD. Bivariate binary models of efficacy and toxicity in dose-ranging trials. Communications in Statistics, Part A - Theory and Methods. 1990; 19:2003–2020.

O'Quigley J, Pepe M, Fisher L. Continual reassessment method: A practical design for phase 1 clinical trials in cancer. Biometrics. 1990; 46:33–48. [PubMed: 2350571]

Paoletti X, Kramar A. A comparison of model choices for the continual reassessment method in phase I cancer trials. Statistics in Medicine. 2009; 28:3012–3028. [PubMed: 19672839]

Smith SM, Besien K, Karrison T, Dancey J, Mclaughlin P, Younes A, Smith S, Stiff P, Lester E, Modi S, Doyle LA, Vokes EE, Pro B. Temsirolimus has activity in nonmantle cell non-Hodgkins lymphoma subtypes: The University of Chicago phase II consortium. Journal of Clinical Oncology. 2010; 28:4740–4746. [PubMed: 20837940]

Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. Biometrics. 2004; 60:684–693. [PubMed: 15339291]

Thall PF, Millikan RE, Mueller P, Lee SJ. Dose-finding with two agents in phase I oncology trials. Biometrics. 2003; 59:487496.

Wages NA, Conaway MR, O'Quigley J. Continual reassessment method for partial ordering. Biometrics. 2011; 67:1555–1563. [PubMed: 21361888]

Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. Biometrics. 2005; 61:217–222. [PubMed: 15737096]

Yin G, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. Journal of the Royal Statistical Society: Series C (Applied Statistics). 2009; 58:211–224.

Yuan Y, Yin G. Sequential continual reassessment method for two-dimensional dose finding. Statistics in Medicine. 2008; 27:5664–5678. [PubMed: 18618901]

Yuan Y, Yin G. Robust EM continual reassessment method in oncology dose finding. Journal of the American Statistical Association. 2011; 106:818–831. [PubMed: 22375092]

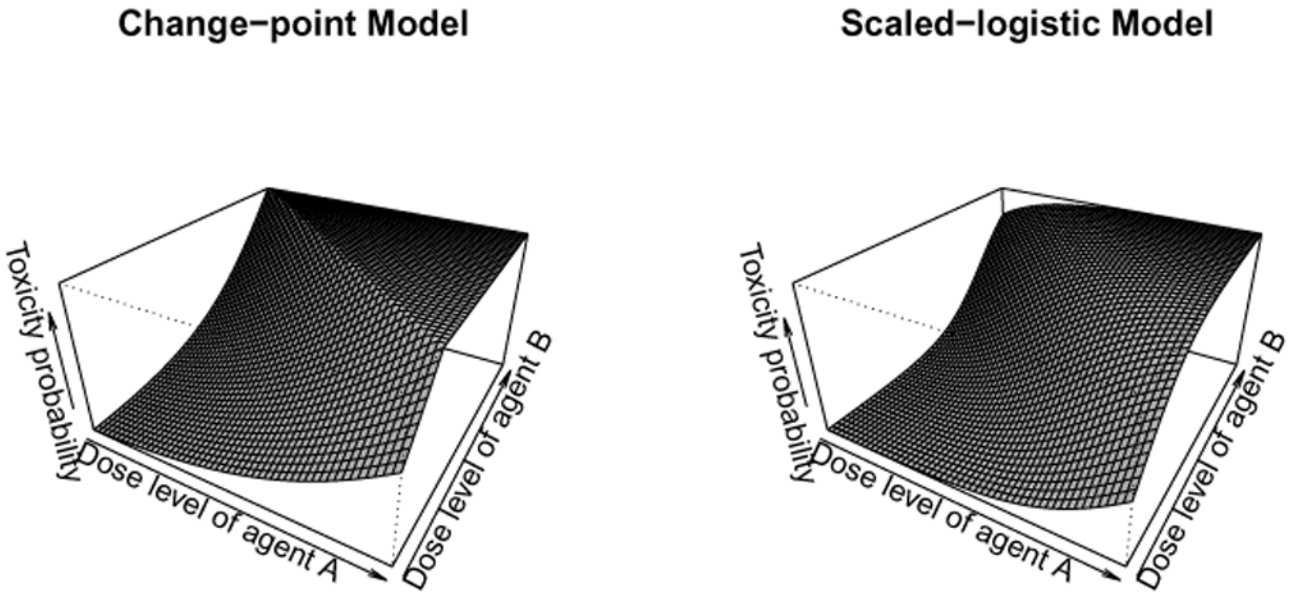

Fig. 1.
Surface of the toxicity probabilities for combinational agents using the proposed change-point model and scaled-logistic model. Toxicity initially increases with dose level and then plateaus.

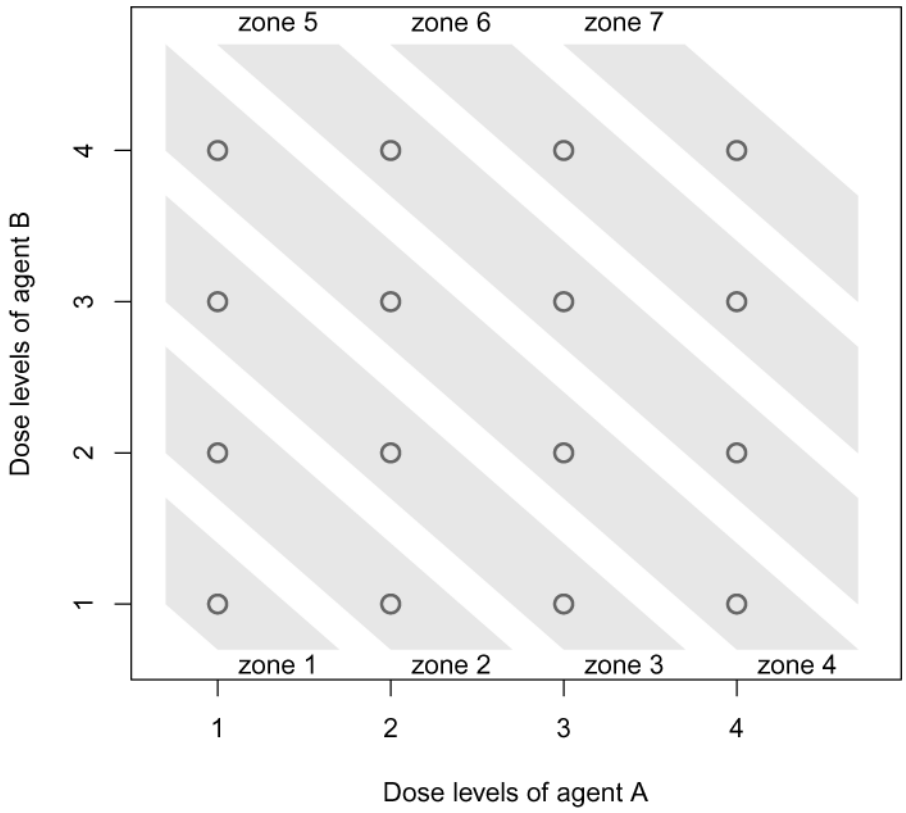

Fig. 2.
An illustration of zones based on the partial order of the dose-toxicity relationship for combinational agents.

Table 1

Twelve dose-toxicity and dose-efficacy scenarios for the simulation studies. The target BODCs are bolded.

	Agent B								Agent B							
	Toxicity probability				Efficacy probability				Toxicity probability				Efficacy probability			
Agent A	**1**	**2**	**3**	**4**	**1**	**2**	**3**	**4**	**1**	**2**	**3**	**4**	**1**	**2**	**3**	**4**
	Scenario 1								Scenario 7							
4	.14	.25	.25	.25	.42	.60	.38	.32	.24	.25	.25	.25	.12	.15	.24	.40
3	.09	.15	.18	.25	.19	.44	.20	.18	.22	.24	**.25**	.25	.26	.31	.43	**.61**
2	.04	.08	.13	.18	.12	.29	.15	.10	.17	.22	.24	.25	.13	.16	.25	.41
1	.02	.04	.07	.12	.05	.22	.10	.08	.09	.17	.22	.24	.00	.01	.01	.04
	Scenario 2								Scenario 8							
4	.14	.25	.25	.25	.10	.10	.18	.24	.24	.25	**.25**	.25	.06	.44	**.62**	.42
3	.09	.15	.18	.25	.14	.14	.24	.43	.22	.24	.25	.25	.01	.23	.39	.22
2	.04	.08	.13	.18	.23	.28	.42	**.60**	.17	.22	.24	.25	.01	.14	.26	.13
1	.02	.04	.07	.12	.08	.10	.29	.42	.09	.17	.22	.24	.00	.11	.23	.11
	Scenario 3								Scenario 9							
4	.15	**.18**	.21	.25	.42	**.60**	.38	.32	.23	**.24**	.24	.24	.06	.44	**.62**	.42
3	.10	.15	.19	.23	.19	.44	.20	.18	.19	.20	.22	.22	.01	.23	.39	.22
2	.05	.12	.15	.20	.12	.29	.15	.10	.12	.14	.16	.17	.01	.14	.26	.13
1	.01	.07	.12	.18	.05	.22	.10	.08	.05	.06	.08	.10	.00	.11	.23	.11
	Scenario 4								Scenario 10							
4	.25	.25	.25	**.25**	.20	.30	.48	**.60**	**.12**	.23	.38	.54	**.63**	.46	.23	.06
3	.15	.25	.25	.25	.12	.25	.39	.43	.05	.12	.24	.40	.40	.25	.10	.02
2	.10	.25	.25	.25	.08	.18	.26	.28	.02	.06	.13	.25	.25	.13	.04	.01
1	.05	.10	.15	.25	.05	.07	.15	.19	.01	.02	.06	.14	.17	.08	.02	.00
	Scenario 5								Scenario 11							
4	.23	.42	.43	.44	.05	.08	.15	.10	.10	.25	.42	.54	.49	.36	.10	.00
3	.15	.23	.42	.43	.10	.14	.24	.18	**.03**	.10	.24	.41	**.66**	.53	.20	.01
2	.10	.15	.23	.42	.20	.28	.37	.26	.00	.02	.09	.23	.48	.35	.10	.00
1	.05	.10	**.18**	.25	.30	.40	**.60**	.37	.00	.00	.02	.09	.08	.04	.00	.00
	Scenario 6								Scenario 12							

	Agent B								Agent B							
	Toxicity probability				Efficacy probability				Toxicity probability				Efficacy probability			
Agent A	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
4	.39	.45	.50	.55	.15	.10	.08	.05	.21	.38	.44	.45	.00	.00	.01	.00
3	.25	.42	.46	.51	.24	.18	.14	.10	.06	.22	.38	.44	.00	.08	.18	.08
2	.15	.26	.40	.49	.37	.26	.20	.13	.01	.07	.23	.39	.03	.32	.50	.32
1	**.05**	.18	.26	.38	**.60**	.37	.30	.24	.00	.01	**.07**	.24	.06	.45	**.63**	.45

Table 2

Simulation results for scenarios 1-6 under the proposed, scaled-logistic, greedy and zone-based designs.

	Scenario					
Design	**1**	**2**	**3**	**4**	**5**	**6**
Selection percentage of the BODC						
Proposed	36.2	38.7	37.8	44.8	36.1	74.0
Scaled-logistic	37.0	37.7	38.4	41.1	37.1	74.0
Greedy	24.1	29.4	23.9	56.6	29.8	79.7
Zone-based	34.4	35.3	35.8	47.2	37.5	69.1
Percentage of patients allocated to the BODC						
Proposed	17.3	17.9	17.6	21.9	16.4	42.2
Scaled-logistic	17.3	17.6	17.8	19.6	16.4	41.6
Greedy	10.2	12.2	9.6	46.4	11.4	64.3
Zone-based	7.8	7.8	8.2	7.7	10.9	21.2
Average efficacy rate						
Proposed	33.3	33.6	33.5	34.8	32.8	39.6
Scaled-logistic	32.9	34.1	33.2	33.8	32.4	39.8
Greedy	30.7	30.2	30.4	40.7	30.6	45.6
Zone-based	24.3	26.1	24.5	23.9	27.7	33.0
Number of patients assigned to over-toxic doses						
Proposed	0	0	0	0	8.9	9.9
Scaled-logistic	0	0	0	0	9.0	9.4
Greedy	0	0	0	0	10.0	6.8
Zone-based	0	0	0	0	10.1	14.5
Number of patients assigned in stage I						
Proposed	11.7	11.7	11.5	10.6	10.5	9.8
Scaled-logistic	11.7	11.7	11.5	10.6	10.5	9.8
Greedy	11.6	11.6	11.5	10.6	10.6	9.9
Zone-based	41.1	41.1	39.8	33.8	34.2	27.4
Number of patients assigned in stage II						
Proposed	33.3	33.3	33.4	32.7	33.0	32.5
Scaled-logistic	33.3	33.3	33.4	33.0	33.7	33.2

Design	Scenario					
	1	2	3	4	5	6
Greedy	33.3	33.3	33.4	32.8	33.1	32.6
Zone-based	3.9	3.9	5.2	10.8	10.4	16.0

Table 3

Sensitivity analysis for scenarios 7-12 under the proposed, scaled-logistic, greedy and zone-based designs.

	Scenario					
Design	**7**	**8**	**9**	**10**	**11**	**12**
Selection percentage of the BODC						
Proposed	38.9	40.6	47.8	82.7	49.8	51.3
Scaled-logistic	38.7	38.5	47.0	82.1	48.9	51.5
Greedy	24.5	21.9	25.1	69.0	30.4	40.6
Zone-based	31.6	34.0	39.6	91.9	36.3	42.4
Percentage of patients allocated to the BODC						
Proposed	18.3	18.8	21.4	34.7	22.2	23.3
Scaled-logistic	17.9	18.2	21.2	33.9	22.7	23.6
Greedy	12.1	11.5	12.4	27.1	12.2	16.7
Zone-based	7.8	8.1	7.6	8.2	6.9	8.9
Average efficacy rate						
Proposed	32.1	31.8	33.4	32.7	36.8	34.6
Scaled-logistic	31.4	31.2	34.0	32.4	36.7	34.5
Greedy	34.3	34.0	36.1	28.4	34.0	33.6
Zone-based	21.6	21.1	20.4	19.6	23.9	24.3
Number of patients assigned to over-toxic doses						
Proposed	0	0	0	6.1	6.0	10.0
Scaled-logistic	0	0	0	6.1	5.9	9.5
Greedy	0	0	0	6.0	4.0	10.6
Zone-based	0	0	0	5.2	5.1	11.8
Number of patients assigned in stage I						
Proposed	10.7	10.7	11.2	11.5	11.5	10.9
Scaled-logistic	10.6	10.6	11.2	11.5	11.5	10.9
Greedy	10.7	10.7	11.3	11.5	11.5	10.9
Zone-based	31.1	31.1	37.0	41.7	42.6	38.2
Number of patients assigned in stage II						
Proposed	31.4	31.3	32.9	33.5	33.5	34.1
Scaled-logistic	32.0	32.0	33.1	33.5	33.5	34.1

Design	Scenario						
	7	8	9	10	11	12	
Greedy	31.3	31.3	33.0	33.4	33.5	34.1	
Zone-based	12.3	12.2	7.4	3.3	2.4	6.8	