



Published in final edited form as:

Stat Med. 2019 December 10; 38(28): 5299–5316. doi:10.1002/sim.8361.

A Utility-based Bayesian Optimal Interval (U-BOIN) Phase I/II Design to Identify the Optimal Biological Dose for Targeted and Immune Therapies

Yanhong Zhou, MS^{1,2}, J.Jack Lee, PhD^{*,2}, Ying Yuan, PhD^{*,2}

¹Quantitative Science, The University of Texas MD Anderson Cancer Center UTHealth Graduate School of Biomedical Sciences, TX, US

²Department of Biostatistics, The University of Texas MD Anderson Cancer Center, TX, US

Summary

In the era of targeted therapy and immunotherapy, the objective of dose finding is often to identify the optimal biological dose (OBD), rather than the maximum tolerated dose (MTD). We develop a utility-based Bayesian optimal interval (U-BOIN) phase I/II design to find the OBD. We jointly model toxicity and efficacy using a multinomial-Dirichlet model, and employ a utility function to measure dose risk-benefit trade-off. The U-BOIN design consists of two seamlessly connected stages. In stage I, the Bayesian optimal interval (BOIN) design is used to quickly explore the dose space and collect preliminary toxicity and efficacy data. In stage II, in light of accumulating efficacy and toxicity from both stages I and II, we continuously update the posterior estimate of the utility for each dose after each cohort, and use this information to direct the dose assignment and selection. Compared to existing phase I/II designs, one prominent advantage of the U-BOIN design is its simplicity for implementation. Once the trial is designed, it can be easily applied using predetermined decision tables, without complex model fitting and estimation. Our simulation study shows that, despite its simplicity, the U-BOIN design is robust and has high accuracy to identify the OBD. We extend the design to accommodate delayed efficacy by leveraging the short-term endpoint (e.g., immune activity or other biological activity of targeted agents), and using it to predict the delayed efficacy outcome to facilitate real-time decision making. A user-friendly software to implement the U-BOIN is freely available at www.trialdesign.org.

Keywords

Phase I/II; utility-based design; optimal biological dose; targeted therapy; immunotherapy

*Correspondence Ying Yuan, yyuan@mdanderson.org, J.Jack Lee, jjlee@mdanderson.org.

Present Address Ying Yuan and J.Jack Lee. Department of Biostatistics, The University of Texas MD, Anderson Cancer Center, Houston, TX, 77030, US

1 | INTRODUCTION

Immunotherapy and targeted therapies have revolutionized cancer treatment. Unlike conventional chemotherapy, immunotherapy drugs do not target the tumor directly. Instead, they work by reactivating the immune system and, hence, reestablishing its capacity to combat tumors. A major class of immunotherapy drugs are monoclonal antibodies, known as immune checkpoint inhibitors (ICI) (e.g., nivolumab, pembrolizumab, and ipilimumab). Checkpoint proteins are receptors on immune cells that can be activated to block immune response, for example checkpoint proteins on T cells (e.g., PD-1 and CTLA-4). The ICI bind to the checkpoint receptors on T cells and release “brake” such that T cells can kill cancer cells, achieving treatment efficacy. Nivolumab and pembrolizumab are PD-1 inhibitors, and ipilimumab is CTLA-4 inhibitor.

Traditional dose-finding designs developed for chemotherapy aim to find the maximum tolerated dose (MTD). The underlying assumption is that both toxicity and efficacy monotonically increase with the dose and, thus, the MTD presents the most efficacious dose that is safe. This assumption, however, is often questionable for immunotherapy and targeted agents. Although it is reasonable to assume that toxicity increases with the dose, the same is not necessarily true for efficacy. For example, once the checkpoint binding is saturated, further incrementing the ICI dose does not increase treatment efficacy. For some monoclonal antibodies, it is observed that higher doses actually lead to lower efficacy.¹ In addition, while increasing the dose of immunotherapy agents (or targeted biological agents) may improve efficacy, it still may cause substantial toxicity, due to over-activation of the immune system. When the overall benefit is limited, the value of increasing the dose becomes questionable.

To optimize the treatment benefit of immunotherapy and targeted therapy, therefore, it is important to consider toxicity and efficacy simultaneously and their risk-benefit trade-off during dose finding. The objective of dose finding for targeted therapy and immunotherapy is to identify the optimal biological dose (OBD), defined as the dose that has the highest desirability in terms of the risk-benefit trade-off. Our research is motivated by a phase I/II trial to identify the OBD of a novel humanized antiTROP2 monoclonal antibody in patients with an advanced solid tumor. Trophoblast cell-surface antigen 2 (TROP2) is associated with increased tumor growth. It is overexpressed in the majority of human epithelial cancers, including esophageal, breast, and lung cancers.^{2,3} Five doses will be studied in the trial. Toxicity will be graded according to the NCI-CTCAE version 4.03, with a 21-day assessment window using the scale of 5 grades. The dose limiting toxicity (DLT) is defined as toxicity with grade 3 or higher. Tumor response will be evaluated using RECIST version 1.1, scored as complete remission (CR), partial remission (PR), stable disease (SD) and progressive disease (PD). Response is defined as CR/PR. TROP2 expression will be measured on day 8, after dosing, to provide a biomarker for measuring the biological activity of the drug.

Here, we develop a utility-based seamless phase I/II design to find the OBD. We jointly model toxicity and efficacy using a multinomial-Dirichlet model and employ a utility function to measure dose risk-benefit trade-off. We show that the utility approach is flexible and more general, in that it contains the existing marginal toxicity-efficacy trade-off methods

as special cases. The design consists of two seamlessly connected stages. In stage I, the Bayesian optimal interval (BOIN) design⁴ is used to quickly explore the dose space and collect preliminary toxicity and efficacy data. In stage II, in light of accumulating efficacy and toxicity from both stages I and II, we continuously update the posterior estimate of the utility for each dose after each cohort, and use this information to direct the dose assignment and selection. We refer to the resulting design as the utility-based BOIN (U-BOIN) design. To accommodate the delayed efficacy observed in some targeted and immunotherapy trials, we extend the U-BOIN design and use the short-term endpoint (e.g., immune activity or other biological activity of targeted agents) to predict the delayed efficacy outcome to facilitate real-time decision making.

Numerous phase I/II trial designs have been proposed. Thall and Russell⁵ developed a Bayesian phase I/II design that characterizes patient outcomes using a trinary variable that accounts for both toxicity and efficacy. Braun⁶ generalized the continual reassessment method (CRM)⁷ to accommodate toxicity and efficacy simultaneously. Thall and Cook⁸ presented the EffTox design, based on the trade-offs between toxicity and efficacy. Yin et al.⁹ proposed a phase I/II design that uses the odds ratio of the efficacy and toxicity as a measure of desirability. Yuan and Yin¹⁰ described a phase I/II design that jointly models toxicity and efficacy as time-to-event outcomes. Jin et al.¹¹ proposed a phase I/II design that accommodates late-onset toxicity and efficacy. Liu and Johnson¹² proposed a robust Bayesian phase I/II design, based on a flexible Bayesian dynamic model. Guo and Yuan¹³ proposed a personalized Bayesian phase I/II design that accounts for patient characteristics and biomarker information. Zang et al.¹⁴ proposed several practical phase I/II trial designs to find the OBD. Liu et al.¹⁵ proposed a Bayesian phase I/II trial design for immunotherapy that considers the immune response, toxicity, and efficacy. Yuan et al.¹⁶ provided a comprehensive coverage of phase I/II designs.

Compared to existing phase I/II designs, the proposed U-BOIN design has several strengths. First, the U-BOIN is simple and easy to implement. It is a model-assisted design in that it uses a model for efficient decision making, but its decision rule for dose escalation/de-escalation can be tabulated and included in the trial protocol before the trial starts.^{17,18,19} Once the trial is designed, it can be easily implemented using the predetermined decision table. No complicated computation or model estimation is needed. In contrast, most existing phase I/II designs are model-based designs, and require complicated model fitting and estimation after treating each cohort. Second, the U-BOIN is robust, because it models toxicity and efficacy at each dose independently, without imposing any parametric structure on dose-toxicity and dose-efficacy curves. In contrast, most existing designs assume that the dose-toxicity and dose-efficacy curves follow certain parametric forms (e.g., logistic model), and thus are susceptible to the influence of model misspecification.

Third, most existing phase I/II designs use the trade-off in the marginal toxicity and efficacy probabilities of a dose to measure the desirability of the dose, while the U-BOIN design uses the utility, which is not only easier to elicit from physicians, but also more flexible and general. We prove that the trade-off based on marginal toxicity and efficacy probabilities is a special case of the utility approach. Although utility has been used in previous designs,^{20,21,22} to the best of our knowledge, this paper is the first to formally show that the utility

approach contains the trade-off based on marginal toxicity and efficacy probabilities as a special case. This provides a theoretical justification for using the utility approach. Lastly, the U-BOIN is capable of accommodating delayed efficacy, whereas most existing designs assume that the efficacy endpoint is quickly observed.

The remainder of this article proceeds as follows. Section 2 introduces the statistical model, utility function, dose-finding algorithm, and software to implement the U-BOIN design. Section 3 uses simulation to compare the U-BOIN design with an existing phase I/II design, and also assesses the robustness of the design using sensitivity analysis. Section 4 provides a brief summary.

2 | METHODS

2.1 | Efficacy-toxicity model

Consider a phase I/II trial with J doses under investigation. Let $Y_E = 0, \dots, R-1$ denote the categorical efficacy endpoint with R levels, where a higher level represents a more desirable treatment response; and $Y_T = 0, \dots, Q-1$ denote the categorical toxicity endpoint with Q levels, where a higher level represents a more severe toxicity. The bivariate outcomes (Y_E, Y_T) can be equivalently represented by a single variable Y with $K = R \times Q$ levels, where each level of Y maps to a distinct value of (Y_E, Y_T) .

As an example, consider the conventional setting that both Y_E and Y_T are binary. Let $Y_E = 1$ denote response, 0 otherwise. Similarly, let $Y_T = 1$ denote DLT, 0 otherwise. Then, Y has $K = 4$ levels, with $Y = 1$, if $(Y_E, Y_T) = (0, 1)$; $Y = 2$, if $(Y_E, Y_T) = (0, 0)$; $Y = 3$, if $(Y_E, Y_T) = (1, 1)$; and $Y = 4$, if $(Y_E, Y_T) = (1, 0)$. The value of Y ascribed to each possible (Y_E, Y_T) is not critical, as long as we keep track of the mapping. Without loss of generality, we assume that $Y = 1$ denotes the least favorable clinical outcomes, and $Y = K$ denotes the most favorable clinical outcomes.

Define $\pi_{jk} = \Pr(Y = k | d = j)$, $k = 1, \dots, K$ and $j = 1, \dots, J$, with $\sum_{k=1}^K \pi_{jk} = 1$, where d denotes the dose level. We assume that Y follows a Dirichlet-multinomial model as follows:

$$Y = k | d = j \sim \text{Multinomial}(\pi_{j1}, \dots, \pi_{jK}). \quad (1)$$

$$(\pi_{j1}, \dots, \pi_{jK}) \sim \text{Dirichlet}(a_1, \dots, a_K), \quad (2)$$

where $a_1, \dots, a_K > 0$ are hyper-parameters. We set $\sum_{k=1}^K a_k = 1$, such that the prior is vague and equivalent to a prior sample size of 1.

At an interim decision time, assume that n_j patients have been treated at dose $d = j$, among which n_{jk} patients had outcome $Y = k$, where $n_j = \sum_{k=1}^K n_{jk}$. Given the observed interim data $D_j = (n_{j1}, \dots, n_{jK})$, the posterior distribution of $\pi_j = (\pi_{j1}, \dots, \pi_{jK})$ is

$$\pi_j | D_j \sim \text{Dirichlet}(a_1 + n_{j1}, \dots, a_K + n_{jK}). \quad (3)$$

2.2 | Utility

We measure the desirability of the investigational doses using utility. Let ψ_k denote the utility ascribed to outcome $Y = k$, $k = 1, \dots, K$. The utility ψ_k should be elicited from physicians to reflect the risk-benefit trade-off underlying their medical decisions, which can be done using the following procedure:

1. Fix the value of the utility for the least desirable outcome $Y = 1$ as $\psi_1 = 0$, and for the most desirable outcome $Y = K$ as $\psi_K = 100$. For example, for binary Y_E and Y_T , the least desirable outcome is $(Y_E = 0, Y_T = 1)$, i.e., (no response, DLT), and the most desirable outcome is $(Y_E = 1, Y_T = 0)$, i.e., (response, no DLT).
2. Ask the clinician to use these two utilities as a reference to score the utility values $\psi_2, \dots, \psi_{K-1}$ for the other $K - 2$ possible outcomes $Y = 2, \dots, K - 1$ to quantify the risk-benefit trade-off under each outcome.

Table 1 shows three examples of the utility function. The first two examples consider the scenario where both toxicity and efficacy are binary outcomes. Example 1 has utility values $\{\psi_1 = 0, \psi_2 = 30, \psi_3 = 50, \psi_4 = 100\}$ for the outcomes $\{(Y_E = 0, Y_T = 1), (Y_E = 0, Y_T = 0), (Y_E = 1, Y_T = 1), (Y_E = 1, Y_T = 0)\}$. Compared to example 1, example 2 rewards the response (i.e., $Y_E = 1$) more, in the presence of DLT (i.e., $Y_T = 1$), by assigning a larger value to ψ_3 (65 versus 50). This is appropriate for the trial where toxicity can be well managed and response is highly desirable (e.g., leading to long survival). Example 3 shows the case where Y_T has three levels (i.e., minor, moderate, and severe toxicity), and Y_E has also three levels (i.e., PD, SD, and CR/PR).

In our experience, clinicians quickly understand what the utilities mean and provide values for ψ_k 's, since the values reflect the actual clinical practice. After completing this process, simulation should be performed to verify the operating characteristics of the design. In some cases, the simulation results may motivate slight modification of some of the numerical utility values, although such modification typically has little or no effect on the design's operating characteristics. One possible criticism for using the utility values is that they require subjective input. However, we are inclined to view this as a strength, rather than a weakness. The process of specifying the utility requires clinicians to carefully consider the potential risks and benefits of the treatment that underlie their clinical decision making in a more formal way and incorporate that into the trial. In addition, our simulation study and previous studies^{13,15,21} show that the design is generally not sensitive to the numerical values of the utility, as long as it reflects a similar trend.

Given the values of ψ_k , the true mean utility for dose j is given by

$$U_j = \sum_{k=1}^K \psi_k \pi_{jk}, \quad (4)$$

Since the true mean utility U_j depends on π_{jk} , which is unknown, we need to estimate it, based on the observed data. Given the interim data $D = \{D_j\}$, the estimate of mean utility is given by

$$\hat{U}_j = \sum_{k=1}^K \psi_k E(\pi_{jk} | D). \quad (5)$$

For the conventional setting with binary Y_T and Y_E , another common approach to defining the efficacy-toxicity trade-off is directly based on the marginal efficacy probability $\pi_{E,j} = \Pr(Y_E = 1 | d = j)$ and marginal toxicity probability $\pi_{T,j} = \Pr(Y_T = 1 | d = j)$,¹⁶ which can be expressed as

$$U_j^M = \pi_{E,j} - \omega \pi_{T,j}, \quad (6)$$

where ω is a prespecified weight. This trade-off function says that patients are willing to trade an increase of ω in the DLT rate for a unit increase in the efficacy rate. If $\omega = 0$, we obtain the special case that the dose with the highest efficacy is the most desirable. The following theorem shows that this marginal-probability-based approach is a special case of the utility approach described above. The proof is provided in the Appendix.

Theorem 1.—Marginal-probability-based trade-off U_j^M , defined in (6), is a special case of the utility method defined in (4), in the sense that for a pre-specified weight ω , we can find (ψ_2, ψ_3) , such that $U_j = \xi U_j^M$, where ξ is a non-zero constant.

In the Supplementary Material, we provide an example to show how to map U_j with U_j^M by choosing appropriate values of (ψ_2, ψ_3) .

Liu and Johnson¹² proposed another marginal-probability-based trade-off function

$$U_j^{M2} = \pi_{E,j} - \omega_1 \pi_{T,j} - \omega_2 \pi_{T,j} I(\pi_{T,j} > \rho), \quad (7)$$

where ω_1 and ω_2 is a prespecified weight, $I(\cdot)$ is an indicator function, and ρ is a prespecified toxicity threshold deemed of substantial concern. Compared to U_j^M in (6), this trade-off function is more flexible and allows to impose a higher penalty (i.e., $\omega_1 + \omega_2$) when the true DLT rate $\pi_{T,j}$ exceeds the threshold ρ . U_j^{M2} becomes U_j^M when $\omega_2 = 0$. It is not clear if trade-off function U_j^{M2} can be expressed equivalently in the form of U_j as (4). Nevertheless, the proposed U-BOIN design can be used with different types of utility/trade-off, including U_j^{M2} . We plan to incorporate this functionality in our software provided later.

2.3 | Optimal biological dose (OBD)

To define the OBD, we first define the admissible dose to safeguard patients from toxic or futile doses. As the objective here is to rule out toxic or futile doses, it is more natural to start with the definition of an inadmissible dose. Let $\bar{\pi}_T$ denote the maximum tolerable DLT rate, and $\underline{\pi}_E$ the lowest acceptable response rate. We define that dose j is inadmissible, if it meets either one or both of the following two criteria:

$$\text{(Toxic)} \quad \Pr(\pi_{T,j} > \bar{\pi}_T \mid D) > C_T, \tag{8}$$

$$\text{(Futile)} \quad \Pr(\pi_{E,j} < \underline{\pi}_E \mid D) > C_E, \tag{9}$$

where C_E and C_T are probability cutoffs. In general, $C_T = 0.95$ and $C_E = 0.9$ work well, but should be calibrated using simulation to ensure desirable operating characteristics. This can be done easily using the software provided.

The admissible dose is then defined as the dose for which none of the criteria (8) and (9) is satisfied. When Y_T has more than 2 categories (e.g., grade 0–2, grade 3, and grade 4–5), for the purpose of defining the admissible dose, Y_T can be temporarily collapsed into DLT/no DLT (i.e., DLT = grade 3 and higher). This dichotomization simplifies the definition of the admissible dose and is often adequate to safeguard patients from overly toxic doses. While it is not necessary, more complicated safety criteria could be entertained to accommodate multiple categories. Similar dichotomization (i.e., response/no response) is also applied to Y_E , when it has more than two categories. Note that although we dichotomize Y_T and Y_E for defining the admissible dose, our model and utility are still based on their original scales. We define the OBD as the dose that is admissible and has the highest utility value, i.e.,

$$\text{OBD} = \operatorname{argmax}_{j \in \mathcal{A}}(U_j),$$

where \mathcal{A} denotes the set of admissible doses.

2.4 | Phase I/II OBD finding algorithm

The U-BOIN design consists of two seamless, connected stages (Figure 1). The objective of stage I is to quickly explore the dose space to identify a set of admissible doses that are reasonably efficacious and safe for stage II. In stage I, we conduct dose escalation based on only the toxicity outcome, but efficacy data are also collected and will be used for decision making in stage II. Given the exploratory nature of stage I, if Y_T has more than two categories, we dichotomize it as DLT/no-DLT to facilitate the exploration of the dose space. This is in line with the clinical practice and serves well for the purpose of stage I. Note that for the estimation and finding the OBD (i.e., the primary objective of the trial), we retain the original scale of Y_T .

Stage I dose escalation/de-escalation is guided by the BOIN design,⁴ which has been widely used in variety of oncology trials, including solid tumors,^{23,24} liquid tumors,^{25,26} and various treatment agents.^{27,28,29,30} Due to very limited data and large uncertainty, for patient safety, we set the target DLT rate $\phi_T = \bar{\pi}_T - 0.05$, slightly lower than the maximum tolerable DLT rate $\bar{\pi}_T$, to ensure that stage I dose exploration concentrates around up to, but not exceeding, $\bar{\pi}_T$. Let $\hat{\pi}_{j,T}$ denote the empirical (or maximum likelihood) estimate of $\pi_{j,T}$, given by $\hat{\pi}_{j,T} = m_j/n_j$, where m_j is the number of patients who experienced DLT at the dose level j ; and λ_e and λ_d denote the predetermined optimal escalation boundary and de-escalation boundary. Table 2 provides the values of λ_e and λ_d for the commonly used target

DLT rate ϕ_T , see Liu and Yuan⁴ for the derivation and formula to calculate λ_e and λ_d . The dose-finding algorithm in stage I proceeds as follows.

A1 Patients in the first cohort are treated at dose level 1 or a prespecified starting dose.

A2 Suppose j is the current dose, use the following rules to assign a dose to the next cohort of patients.

- Escalate the dose to $j + 1$ if $\hat{\pi}_{j,T} \leq \lambda_e$.
 - De-escalate the dose to $j - 1$ if $\hat{\pi}_{j,T} \geq \lambda_d$.
 - Otherwise, stay at the current dose j .
-

A3 Repeat step A2 until the number of patients treated on one of the doses reaches s_1 , and then move to stage II. We recommend $s_1 = 12$ as the default value, while $s_1 = 9$ to 15 generally yields good operating characteristics.

Stage II proceeds as follows:

B1 Let j^* denote the highest dose level that has been tried. If $\hat{\pi}_{j^*,T} \leq \lambda_e$ and j^* is not the highest dose in the trial, escalate the dose to d_{j^*+1} for treating the next cohort of patients; otherwise, proceed to step B2.

B2 Given the observed interim data D collected in both stages I and II, determine the admissible dose set \mathcal{A} . If no dose is admissible, terminate the trial and no dose should be selected as the OBD. Otherwise, assign the next cohort of patients to the admissible dose (i.e., $\in \mathcal{A}$) that has the largest posterior mean utility, which can be pre-tabulated.

B3 Repeat steps B1 and B2 until reaching the prespecified maximum sample size N or the number of patients treated at one of the doses reaches $s_2 (> s_1)$, and then select the OBD as the admissible dose (i.e., $\in \mathcal{A}$) that has the largest posterior mean utility. For most trials, a value between 18 to 24 is a reasonable choice for s_2 .

In stage I, following the BOIN design, we impose an overdose control rule as follows: if $\Pr(\pi_{T,j} > \bar{\pi}_T \mid m_j, n_j) > 0.95$ and $n_j \geq 3$, dose level j and higher are eliminated from the trial; the trial is terminated if the lowest dose level is eliminated, where $\Pr(\pi_{T,j} > \bar{\pi}_T \mid m_j, n_j) > 0.95$ is evaluated based on a beta-binomial model with the uniform prior. Once the trial move to stage II, this overdose control rule is seamlessly merged as the inadmissible rule (8). In the overdose control rule, we use $\bar{\pi}_T$, rather than ϕ_T , as the DLT rate threshold to ensure that the overdose control rule seamlessly connects with the inadmissible rule.

For stage II step B1, the reason that we perform dose escalation when $\hat{\pi}_{j^*,T} \leq \lambda_e$ is to allow the trial to continue exploring the dose space, given that the highest tried dose is safe, to reduce the risk of being stuck at a local suboptimal dose, due to a large variation caused by a small sample size. Besides the pick-the-winner (PW) approach (i.e., deterministically assigning the next cohort of patients to dose $j \in \mathcal{A}$ that has the largest posterior mean utility), other strategies can also be used to assign patients. For example, we can adaptively randomize the next cohort of patients to dose $j \in \mathcal{A}$, with probability ω_j proportional to its posterior mean utility, i.e.,

$$\omega_j = \frac{U_j}{\sum_{j \in \mathcal{A}} U_j}. \quad (10)$$

The adaptive randomization (AR) approach can reduce the risk of being stuck at a suboptimal dose, but as a tradeoff, it tends to treat fewer patients at the OBD. Another approach is equal randomization (ER), where the next cohort of patients are assigned to the admissible doses with equal probability, i.e.,

$$\omega_j = \frac{1}{\sum_{j \in \mathcal{A}} 1}. \quad (11)$$

We compare the performance of PW, AR, and ER in our simulation study. None of the methods dominates the others. Thus, we generally recommend the PW approach because of its simplicity.

Unlike most existing phase I/II designs, which require complicated model fitting and estimation after each cohort to make the decision of dose assignment, one prominent advantage of the U-BOIN is that its dose assignment rules can be pre-tabulated in decision tables and included in the trial protocol before the trial starts. To conduct the trial, no complicated calculation is needed. The investigator can simply use the decision tables to determine dose assignment, e.g., determine whether escalation/de-escalation is needed (steps A2 and B1) or identify admissible doses (step B2). Because of this feature, the U-BOIN can be classified as a model-assisted design.^{17,18,19} Section 1 in the Supplementary Material provides an example of using the decision tables in a hypothetical phase I/II trial.

2.5 | Delayed efficacy response

In some trials, efficacy endpoint Y_E requires a long time to be ascertained. In our motivating example, it takes three months to evaluate Y_E . The long assessment window causes a major logistics issue for decision making in stage II. For example, given that the accrual rate is three patients per month, and that patients are treated in cohorts of three, on average, six new patients will be accrued while waiting to evaluate the previous three patients' outcomes. This begs the question: how can new patients receive timely treatment, when the previous patients' outcomes are not yet observed?

Statistically, this means that D are not fully observed, as some Y_E 's are unavailable. As a result, the mean utility estimate (5) and the inadmissible criterion (9) cannot be evaluated for making interim decisions. This issue is known as the late-onset or delayed-outcome issue, which has been studied in literature. Cheung and Chappell³¹ proposed a weighting method to handle late-onset toxicity for phase I clinical trials. Yuan and Yin¹⁰ modeled toxicity and efficacy as time-to-event outcomes that accommodate the unobserved as censored events. Liu et al.³² and Jin et al.¹¹ treated unobserved outcomes as missing data, and proposed a Bayesian data augmentation approach to predict unobserved toxicity and/or efficacy outcomes to facilitate decision making. Cai et al.³³ took the multiple imputation approach to handle unobserved efficacy outcomes in phase II trials.

We follow the approach of Cai et al.³³ and use multiple imputation to handle unobserved Y_E . The innovation here is that we use the measure of biological activity, which is routinely recorded in targeted and immunotherapy trials, as an ancillary variable to help predict (or impute) Y_E . Examples of the measure of biological activity include immune response (e.g., CD8+ T cell count) in immunotherapy trials and gene expression related to the pathway targeted by the treatment agent. The measure of biological activity is often quickly observable after drug administration and correlated with the clinical response. Daud et al.³⁴ showed that the abundance of CD8+ T cells predict response to anti-PD1 therapy, and advocated using the immune activity to predict the likelihood of achieving a clinical response to the PD-1 pathway inhibitor. For ease of exposition, hereafter we use the immune response, denoted by Y_I , as the example to illustrate our approach, but Y_I can be any reasonable biological activity measure predictive of treatment efficacy. We assume that Y_I is quickly observable, thus its value is always available at the time of decision making.

Given the observed value of Y_I , we predict the unobserved Y_E based on the following scaled logistic regression model:

$$\text{logit}\left(\frac{\pi_{E,j}}{\lambda}\right) = \beta_0 + \beta_1 Y_I, \quad (12)$$

where $\pi_{E,j}$ is the probability of efficacy for dose j , β_0 and β_1 are regression parameters, and $0 < \lambda < 1$ is a plateau (or scale) parameter used to reflect that the probability of clinical response often levels out after the immune activity reaches a certain level. Under this model, the probability of efficacy $\pi_{E,j}$ increases with Y_I and then plateaus at the value of λ when Y_I is sufficiently large. When $\lambda = 1$, it becomes a standard logistic regression model. We do not include dose level j in the model, because the treatment effect of immunotherapy is mediated by the immune response, and thus it is often reasonable to assume that, conditional on Y_I , $\pi_{E,j}$ is independent of j . If this assumption is not plausible in some situations, one can simply add j as a covariate to the model. In addition, more biomarkers, when available, can be added to model (12) as predictors to improve the prediction accuracy. One attractive property of our imputation approach is that, when the imputation model (12) is misspecified, its impact on the design diminishes over time and eventually goes away when the trial is completed. This is because, as the trial proceeds, more and more patients' Y_E become observed, and accordingly fewer and fewer percentage of Y_E needs prediction. As a result, our method is generally robust to model misspecification, as shown later in the sensitivity analysis.

In terms of prior specifications in model (12), following Gelman et al.,³⁵ we assume that the model parameters $(\lambda, \beta_0, \beta_1)$ are independent and have their own prior distributions. We specify a uniform distribution for the scale parameter $\lambda \sim \text{uni} f(0, 1)$. The magnitude of the coefficients (β_0 and β_1) could be very large or small, depending on different trials. This makes it difficult to specify standard prior distributions for the coefficients. To tackle this problem, we first standardize the variable Y_I . In our application, Y_I is continuous; we standardize it to have a mean of 0 and a standard deviation of 0.5. We regulate the prior distributions to make sure that a typical change in a covariate should not lead to a dramatic change in the efficacy rate.

After standardizing data, a change of 2.5 on the logit scale can move the probability of a favorable response to the therapy from 0.2 to 0.75. We assume that the effect of the immune response is unlikely to be more dramatic than that. This typically is true for immunotherapy trials, as the efficacy rate is rarely outside of that range. We assign a normal distribution with a mean of 0 and a standard deviation 1.25 for β_0 : $\beta_0 \sim N(0, 1.25^2)$. The parameter β_1 is supposed to be positive, due to the positive relationship between immune response and tumor response, and thus we assign a gamma distribution with a shape parameter of 1 and a rate parameter of 1.2: $\beta_1 \sim \text{Gamma}(\text{shape} = 1, \text{rate} = 1.2)$. The priors ensure that a change in the covariate Y_I from one standard deviation below the mean to one standard deviation above the mean will lead to an absolute change that is mostly less than 2.5 on the logit scale.

Let $f(\beta_0, \beta_1, \lambda)$ denote the joint prior distribution of $(\beta_0, \beta_1, \lambda)$. The posterior distribution of $(\beta_0, \beta_1, \lambda)$ is given by

$$f(\beta_0, \beta_1, \lambda | D) \propto f(\beta_0, \beta_1, \lambda) \prod_{j=1}^J \prod_{i=1}^{n_j} \left(\frac{\lambda e^{\beta_0 + \beta_1 y_{I,ji}}}{1 + e^{\beta_0 + \beta_1 y_{I,ji}}} \right)^{y_{E,ji}} \left(\frac{1 + (1 - \lambda)e^{\beta_0 + \beta_1 y_{I,ji}}}{1 + e^{\beta_0 + \beta_1 y_{I,ji}}} \right)^{1 - y_{E,ji}}, \quad (13)$$

where $y_{I,ji}$ is the immune response for the i th patient treated with dose j and $y_{E,ji}$ is the tumor response for this patient.

At an interim decision time, let $Y_{E,obs}$ and $Y_{E,mis}$ denote the observed and missing parts of Y_E . We impute the value of $Y_{E,mis}$ using multiple imputations as follows.

1. Conditional on observed data $(Y_{E,obs}, Y_I)$, sample L draws of $(\lambda, \beta_0, \beta_1)$ from their posterior distribution (13) using the adaptive rejection Metropolis sampling.³⁶ A certain number of burn-in iterations are typically needed before collecting the posterior draws. In our simulation, we set the number of burn-in iterations as $L' = L/2$.
2. Thin the posterior draws by taking a sample after every L'/H draw, resulting in a total of H sets of posterior draws of $(\lambda, \beta_0, \beta_1)$, denoted as $(\lambda^{(1)}, \beta_0^{(1)}, \beta_1^{(1)}), \dots, (\lambda^{(H)}, \beta_0^{(H)}, \beta_1^{(H)})$.
3. Impute $Y_{E,mis}$ for H times, where the h th imputed value is generated by drawing a random sample from $\text{Bernoulli}(q)$, where $h = 1, \dots, H$ and $q = \lambda e^{\beta_0 + \beta_1 Y_I} / (1 + e^{\beta_0 + \beta_1 Y_I})$. We denoted the H sets of imputed values as $Y_{E,imp}^{(1)}, \dots, Y_{E,imp}^{(H)}$.

After filling in $Y_{E,mis}$ with each of the H imputed values, we obtained an H imputed complete dataset, $D^{(1)}, \dots, D^{(H)}$, where $D^{(h)}$ is obtained by filling in $Y_{E,mis}$ with $Y_{E,imp}^{(h)}$, $h = 1, \dots, H$. Then, the estimate of mean utility (5) is given by

$$\hat{U}_j = \frac{1}{H} \sum_{h=1}^H \sum_{k=1}^4 \psi_k E(\pi_{jk} | D^{(h)}), \quad (14)$$

and the left side of the inadmissible criterion (9) is calculated as

$$\Pr(\pi_{E,j} < \underline{\pi}_E | D) = \frac{1}{H} \sum_{h=1}^H \Pr(\pi_{E,j} < \underline{\pi}_E | D^{(h)}). \quad (15)$$

The dose finding follows the same algorithm described in section 2.4. Little and Rubin³⁷ suggested that, for practical use, the number of multiple imputation (H) can be sufficient when $H=5$ or greater. In our simulation, we perform $H=20$ imputations.

To further improve the efficiency of the design, Y_I can also be used to refine the admissible rule. The rationale is that when Y_E takes a long time to be observed, the inadmissible criterion (9) is not effective for screening out ineffective doses, because a high percentage of Y_E may not be observed at the time of making interim decisions. Since Y_I is quickly observable, it can be used, supplementary to Y_E , to improve the power of identifying effective doses or safeguard patients from ineffective doses. Specifically, define $\mu_{I,j} = E(Y_I | d = j)$, and let $\underline{\mu}_I$ denote the lowest acceptable mean immune response. We define dose j inadmissible if it meets at least one of the toxicity and futility criteria in (8) and (9), and

$$\text{(Insufficient activity)} \quad \Pr(\mu_{I,j} < \underline{\mu}_I | D) > C_I, \quad (16)$$

where C_I is a pre-specified probability threshold. In other words, if a dose has little activity to activate immune system, it is deemed unpromising and inadmissible. Again, the admissible dose is defined as the dose that is not inadmissible. The posterior probability $\Pr(\mu_{I,j} < \underline{\mu}_I | D)$ can be evaluated based on the Bayesian normal model:

$$\begin{aligned} Y_I | d = j &\sim N(\mu_{I,j}, \sigma_j^2) \\ f(\mu_{I,j}, \sigma_j^2) &\propto \sigma_j^{-2}. \end{aligned} \quad (17)$$

2.6 | Software and trial implementation

To facilitate the use of the U-BOIN design, we develop an easy-to-use web app, which is freely available at <http://www.trialdesign.org>. Figure 2 shows the graphical user interface of the app. Extensive help files are available by clicking on the yellow question mark at the upper right corner of each tab.

To design a trial, follow the steps:

- Specify the design parameters (e.g., dose information, sample size, cohort size, utility function, admissible criteria, etc.), under the **Trial Setting** tab. After parameters are entered, design diagram and decisions tables will be provided under the tab.

- Generate the operating characteristics of the design by supplying different scenarios (under the **Simulation** tab).
- Download trial protocol under **Trial Protocol** tab. The protocol includes (1) a brief description of the U-BOIN design, (2) decision tables based on input under **Trial Setting**, and (3) operating characteristics for the scenarios entered under **Simulation**.

To conduct a trial, use the decision tables generated under **Trial Setting** or use the **Trial Conduct** tab. The app will return the dose for treating the next cohort of patients if the data provided indicates that the trial is not completed, and it will select the OBD if the trial is completed.

3 | NUMERICAL STUDY

3.1 | Simulation A

Simulation A considers the case that Y_T and Y_E are quickly ascertainable. Following our motivating trial, we consider $J = 5$ doses, and the total sample size $N = 54$ patients with $s_1 = 12$. The minimum acceptable efficacy rate is $\underline{\pi}_E = 0.2$, and the maximum acceptable DLT rate is $\bar{\pi}_T = 0.30$. We set $\pi_T^* = 0.25$ by using $\delta = 0.05$. For the inadmissible criteria (i.e., equations (8) and (9)), probability cutoffs are set as $C_T = 0.95$ and $C_E = 0.9$, based on simulation calibration. The elicited utility is presented in Table 1 as example 1.

We consider 8 representative scenarios that differ in the shape of the dose-toxicity and dose-efficacy curves, and the location of the OBD. The scenarios are shown in Figure 3. In scenarios 1 and 2, the dose-response curve increases with the dose level, and then plateaus. In scenarios 3 and 4, the dose-response curve increases to an optimal point, and then decreases. In scenario 5, the dose-response curve levels off at the first dose level. In scenarios 6, 7, and 8, all dose-response curves monotonically increase. Scenario 6 has the last dose level being the OBD. Scenario 7 has two OBDs, located at dose level 2 and 3. Scenario 8 has no OBD since the doses are either futile or overly toxic. Under each scenario, we generate (Y_T, Y_E) , based on a Gumbel model:

$$\Pr(Y_T = y_T, Y_E = y_E \mid d = j) = (\pi_{E,j})^{y_E} (1 - \pi_{E,j})^{1 - y_E} (\pi_{T,j})^{y_T} (1 - \pi_{T,j})^{1 - y_T} + \pi_{E,j} (1 - \pi_{E,j}) \pi_{T,j} (1 - \pi_{T,j}) (-1)^{y_E + y_T} \left(\frac{e^c - 1}{e^c + 1} \right), \quad (18)$$

where $y_E, y_T \in \{0, 1\}$, and the association parameter c was set as 0.2. The values of $\pi_{E,j}$ and $\pi_{T,j}$ (i.e., the marginal efficacy rate and DLT rate) for each dose level are provided in Table 3. We compare the proposed U-BOIN design with the EffTox design.⁸ We consider the two most important metrics for comparison: (1) percentage of correct selection (PCS), which is the probability of correctly identifying the OBD, and (2) patient allocation, which referred to the average number of patients assigned to each dose. In the EffTox design, the toxicity-efficacy trade-off is constructed based on three equally desirable trade-off target probabilities: $(\pi_E = 0.2, \pi_T = 0)$, $(\pi_E = 1, \pi_T = 0.7)$, and $(\pi_E = 0.45, \pi_T = 0.3)$. Under each scenario, we performed 2,000 simulations.

Table 3 summarizes the operating characteristics for the designs. The U-BOIN design outperforms EffTox as it has a larger PCS and allocates more patients on the OBD. In scenarios 1 and 2, the dose-response curve increases first, and then plateaus, and the OBDs (dose level 2 and 3, respectively) are one dose level lower than the MTD (i.e., dose level with DLT rate closest to the target DLT rate). The U-BOIN has a 23% higher PCS and assigns seven more patients on the OBD than EffTox does in scenario 1. Similarly, the U-BOIN has a larger PCS in scenario 2, for which the two designs have comparable patient allocation. In scenarios 3 and 4 where the dose-response curve increases to an optimal point and then decreases, U-BOIN has 22% and 27% higher PCS than EffTox, respectively. Moreover, U-BOIN assigns nine and ten more patients on the OBD, respectively in the two scenarios. In scenario 5 where the OBD is located on the first dose, and the response rate does not change with dose levels, U-BOIN has a 6% higher PCS and assigns three more patients on the OBD. In scenario 6, where all doses are safe and the OBD is the last dose level, the U-BOIN has a 27% higher PCS and assigns eight more patients on the OBD. This is because EffTox cannot distinguish well the suboptimal doses from the OBD in this scenario, subsequently assigning more patients on dose level 3 and 4 and selecting one of the dose levels as the OBD 55% of the time. U-BOIN has comparable performance to EffTox when there are two OBDs (scenario 7), in terms of PCS and patient allocation. When there is no OBD due to toxicity or futility (scenario 8), U-BOIN has a larger chance to stop the trial early.

3.2 | Simulation B

Simulation B considers the case that Y_T is quickly ascertainable, but Y_E takes a long time to be scored with the assessment window of 3 months. We assume that patient accrual follows a Poisson process, with the rate of 3 patients per month. We simulated Y_T and Y_E based on the same Gumbel model and 8 scenarios, as described in Simulation A. The reason we chose the same simulation scenarios is that, by doing so, the results from Simulation A (i.e., Y_E is always observed) can be used as a benchmark to evaluate the performance of the design in Simulation B (i.e., Y_E is partially observed, due to delayed response). To generate a delayed response, for patients who experience efficacy in the assessment window (i.e., $Y_E = 1$), we simulate their time to efficacy from a truncated Weibull distribution with a support of (0, 3) months. The shape and scale parameter for the Weibull distribution are chosen such that the efficacy rate at the end of assessment time matches those in Table 3, and that 90% of the responses occur in the latter half of the assessment window (i.e., (1.5, 3) months). The immune response Y_I is generated from $\mathcal{N}(\mu_{I,j}, 1)$, where $\mu_{I,j}$ is plotted in Figure 3. We set $C_I = 0.9$. Under each scenario, 2,000 simulations are performed.

Table 4 shows the simulation results. The PCS of the OBD and the number of patients allocated to the OBD are generally comparable to the results in Simulation A (i.e., the optimal benchmark with fully observed data) and U-BOIN still outperforms the EffTox design. The results indicate that U-BOIN efficiently handles the delayed efficacy response. Because the U-BOIN design does not need to suspend accrual to wait Y_E to be fully observed and allows real-time decision making, it has great potential to shorten the trial duration.

To examine the performance of U-BOIN when there are delayed outcomes at a relatively smaller sample size, we conduct the simulation again using $N = 39$. Results show that U-BOIN still maintains its great operating characteristics, even when sample size is relatively small. The simulation results are provided in Table S8 of the Supplementary Material. We also provide operating characteristics for eight additional representative scenarios (Scenarios A1-A8) for both Simulation A (Table S9) and Simulation B (Table S10). The results, again, show that U-BOIN has robust performance in various scenarios in comparison to EffTox.

3.3 | Sensitivity analysis

We conduct sensitivity analyses to assess the robustness of the U-BOIN design by using a different set of utility values $\Psi' = \{\psi_1 = 0, \psi_2 = 20, \psi_3 = 55, \psi_4 = 100\}$, which assigns a lower score for $(Y_E = 0, Y_T = 1)$ and a higher score for $(Y_E = 1, Y_T = 1)$, indicating that patients are willing to tolerate a higher toxicity risk in order to attain a higher efficacy. We also evaluate the performance of U-BOIN using different patient allocation strategies (i.e., step B2 of Stage II of the dose finding algorithm): PW, AR, and ER.

The simulation results (see Figure 4) show that the U-BOIN design performs well with comparable PCS and patient allocations for the two different utilities. This characteristic is important, since two clinicians might have slightly different opinions about the specific utility value assigned to an outcome combination. Our design shows robustness for this type of deviation. Figure 5 shows the results under different patient allocation strategies. The PCS is comparable among the PW, AR, and ER strategies. The difference among the three strategies mainly lies in the number of patients treated on OBD. The PW approach, on average, assigns significantly more patients on the OBD than both the ER and AR approaches, while it has larger variability. Since the variation is more towards the higher end (i.e., more patients are treated on OBD), the PW approach is desirable in this simulation study.

Note that in this Simulation A, we specify $s_2 = N$ for fair comparison with the EffTox design, which does not stop the trial early on the basis of the number of patients treated on a dose. A sensitivity analysis using $s_2 = \{18, 21\}$ is provided in Section 5 of the Supplementary Material. The result in Figure S1 shows that using the recommended values, the change in the percentage of correct selection is negligible, but the saving in sample size is substantial.

For the case with delayed Y_E , we performed three additional sensitivity analyses by (1) assuming that the efficacy assessment window is two months or four months; (2) specifying more vague prior distributions for the coefficients β_0 and β_1 with $\beta_0 \sim N(0, 3.75^2)$ and $\beta_1 \sim \text{Gamma}(shape = 1, rate = 0.4)$, both of which have standard deviations that are three times the previous values in section 2.5; and (3) considering two additional scenarios (scenarios B1 and B2 in Table S11), where Y_I and Y_E are weakly associated with Pearson's correlation coefficients -0.15 and 0.2 , respectively. The results show that the U-BOIN design is also robust to the length of assessment window, prior distribution of the prediction model (Figure S2), and the association between Y_I and Y_E .

4 | SUMMARY

We propose the U-BOIN, a seamless phase I/II model-assisted design, to identify the OBD for targeted and immunotherapy trials. The U-BOIN design accounts for the efficacy-toxicity trade-off using a utility function. Unlike most existing phase I/II designs, which require complicated real-time model fitting and estimation to make dose assignment decisions, the U-BOIN is simple and easy to implement. The dose assignment rules of the U-BOIN can be pre-tabulated in decision tables and included in the trial protocol before onset of the trial. To conduct the trial, no complicated calculation is needed. The investigator can simply use the decision tables to make the decision of dose escalation/de-escalation. Simulation studies show that compared to a more complicated model-based phase I/II design, the U-BOIN has higher accuracy to identify OBD and is more robust. To facilitate the use of the U-BOIN design in clinical trials, we develop a user-friendly software freely available at www.trialdesign.org.

While this article focuses on immunotherapy and targeted trials, U-BOIN also can be used for conventional cytotoxic agent trials. In such cases, both toxicity and efficacy typically increase with the dose, but may do so at different rates. It is likely that increasing the dose causes much higher toxicity, with limited efficacy gain. The idea of risk-benefit tradeoff, and thus finding the optimal dose, is still generally applicable here for use in most medical decisions in practice.

As most model-assisted designs, U-BOIN models efficacy and toxicity at each dose independently, whereas model-based phase I-II designs (e.g., EffTox design) model efficacy and toxicity across all doses, through imposing a parametric dose-efficacy and -toxicity curve model. As a result, one may worry about the potential efficiency loss for U-BOIN. Our numerical study and previous studies show that, for the purpose of dose finding, the efficiency loss caused by using only local data (in model-assisted designs, such as U-BOIN) is minimal or negligible. This can be explained as follows. First, although U-BOIN models only the local data at the current dose, its dose-finding algorithm (e.g., escalate the dose if the current dose is safe, and de-escalate the dose if the current dose is toxic) implicitly uses the dose-toxicity order information across the doses. In addition, in practice, the parametric dose-efficacy and -toxicity model assumed by model-based designs is more likely to be misspecified than correctly specified. Thus, on average, borrowing information across doses through the parametric model leads to rather limited efficiency gain. Furthermore, such limited efficiency gain does not necessarily translate into performance gain. This is because, to make correct decisions of dose assignment and selection, we only need to correctly estimate the rank of utility across the doses. A slightly more variability on the estimate of the utility has no or negligible impact on the performance of the design.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

APPENDIX

A PROOF

Theorem 1.

Marginal-probability-based trade-off U_j^M , defined in (6), is a special case of the utility method defined in (4), in the sense that for a pre-specified weight ω , we can find (ψ_2, ψ_3) , such that $U_j = \xi U_j^M$, where ξ is a non-zero constant.

Proof—If there is a constant ξ , such that $U_j^M = \xi \tilde{U}_j$, then, \tilde{U}_j is a special case of U_j^M .

To fix scale for utility, we propose in this paper that $\psi_1 = 0, \psi_4 = 100$. Then U_j^M is a function of only ψ_2 and ψ_3 .

$$U_j^M = \psi_2 \pi_{j2} + \psi_3 \pi_{j3} + 100 \pi_{j4}$$

It is obvious that $U_j^M > 0$ by definition. Denote $\tilde{U}_j(\omega) = \tilde{U}_j = \pi_{E,j} - \omega \pi_{T,j}$. If $U_j^M = \xi \tilde{U}_j$, then $\xi \neq 0, \tilde{U}_j(\omega) \neq 0$, and $\psi_2 \pi_{j2} + \psi_3 \pi_{j3} + 100 \pi_{j4} = \xi \tilde{U}_j(\omega)$. Since

$0 < \psi_2 \pi_{j2} + \psi_3 \pi_{j3} < 100(\pi_{j2} + \pi_{j3})$, we have $\frac{100 \pi_{j4}}{\tilde{U}_j(\omega)} < \xi < \frac{100(1 - \pi_{j1})}{\tilde{U}_j(\omega)}$, which is nonzero constant.

In summary, for a pre-specified ω and $\tilde{U}_j(\omega) \neq 0$, we can find (ψ_2, ψ_3) and a non-zero constant ξ , such that $U_j^M = \xi \tilde{U}_j$, which says that \tilde{U}_j is a special case of U_j .

References

1. Reynolds AR. Potential relevance of bell-shaped and u-shaped dose-responses for the therapeutic targeting of angiogenesis in cancer. *Dose-response* 2010; 8(3): dose-response.
2. Nakashima K, Shimada H, Ochiai T, et al. Serological identification of TROP2 by recombinant cDNA expression cloning using sera of patients with esophageal squamous cell carcinoma. *International journal of cancer* 2004; 112(6): 1029–1035. [PubMed: 15386348]
3. Ni IBP, Zakaria Z, Muhammad R, et al. Gene expression patterns distinguish breast carcinomas from normal breast tissues: the Malaysian context. *Pathology-Research and Practice* 2010; 206(4): 223–228.
4. Liu S, Yuan Y. Bayesian optimal interval designs for phase I clinical trials. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 2015; 64(3): 507–523.
5. Thall PF, Russell KE. A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* 1998; 251–264. [PubMed: 9544520]
6. Braun TM. The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Controlled clinical trials* 2002; 23(3): 240–256. [PubMed: 12057877]
7. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase 1 clinical trials in cancer. *Biometrics* 1990; 33–48. [PubMed: 2350571]
8. Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* 2004; 60(3): 684–693. [PubMed: 15339291]

9. Yin G, Li Y, Ji Y. Bayesian dose-finding in phase I/II clinical trials using toxicity and efficacy odds ratios. *Biometrics* 2006; 62(3): 777–787. [PubMed: 16984320]
10. Yuan Y, Yin G. Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 2009; 58(5): 719–736.
11. Jin IH, Liu S, Thall PF, Yuan Y. Using data augmentation to facilitate conduct of phase I–II clinical trials with delayed outcomes. *Journal of the American Statistical Association* 2014; 109(506): 525–536. [PubMed: 25382884]
12. Liu S, Johnson VE. A robust Bayesian dose-finding design for phase I/II clinical trials. *Biostatistics* 2016; 17(2): 249–263. [PubMed: 26486139]
13. Guo B, Yuan Y. Bayesian phase I/II biomarker-based dose finding for precision medicine with molecularly targeted agents. *Journal of the American Statistical Association* 2017; 112(518): 508–520. [PubMed: 32863478]
14. Zang Y, Lee JJ, Yuan Y. Adaptive designs for identifying optimal biological dose for molecularly targeted agents. *Clinical Trials* 2014; 11(3): 319–327. [PubMed: 24844841]
15. Liu S, Guo B, Yuan Y. A Bayesian phase I/II trial design for immunotherapy. *Journal of the American Statistical Association* 2018: 1–12. [PubMed: 30034060]
16. Yuan Y, Nguyen HQ, Thall PF. *Bayesian designs for phase I–II clinical trials*. Chapman and Hall/CRC. 2016.
17. Yan F, Mandrekar SJ, Yuan Y. Keyboard: a novel Bayesian toxicity probability interval design for phase I clinical trials. *Clinical Cancer Research* 2017; 23(15): 3994–4003. [PubMed: 28546227]
18. Zhou H, Yuan Y, Nie L. Accuracy, safety, and reliability of novel phase I trial designs. *Clinical Cancer Research* 2018; 24(18): 4357–4364. [PubMed: 29661774]
19. Yuan Y, Lee JJ, Hilsenbeck SG. Model-Assisted Designs for Early Phase Clinical Trials: Simplicity Meets Superiority (under revision). *Journal of Clinical Oncology*.
20. Houede N, Thall PF, Nguyen H, Paoletti X, Kramar A. Utility-based optimization of combination therapy using ordinal toxicity and efficacy in phase I/II trials. *Biometrics* 2010; 66(2): 532–540. [PubMed: 19673865]
21. Murray TA, Yuan Y, Thall PF, Elizondo JH, Hofstetter WL. A utility-based design for randomized comparative trials with ordinal outcomes and prognostic subgroups. *Biometrics* 2018.
22. Murray TA, Thall PF, Yuan Y, McAvoy S, Gomez DR. Robust treatment comparison based on utilities of semi-competing risks in non-small-cell lung cancer. *Journal of the American Statistical Association* 2017; 112(517): 11–23. [PubMed: 28943681]
23. Lim B A Phase II Study of Triple Combination of Atezolizumab + Cobimetinib + Eribulin (ACE) in Patients With Recurrent/Metastatic Inflammatory Breast Cancer. <https://clinicaltrials.gov/ct2/show/NCT03202316>. Published February 8, 2019. Accessed March 18, 2019..
24. Wu J TG02 Plus Dose-Dense or Metronomic Temozolomide Followed by Randomized Phase II Trial of TG02 Plus Temozolomide Versus Temozolomide Alone in Adults With Recurrent Anaplastic Astrocytoma and Glioblastoma. <https://clinicaltrials.gov/ct2/show/NCT02942264>. Published January 4, 2019. Accessed March 18, 2019..
25. Al-Atrash G. Involumab and Ipilimumab After Donor Stem Cell Transplant in Treating Participants With High Risk Refractory or Relapsed Acute Myeloid Leukemia. <https://clinicaltrials.gov/ct2/show/NCT03600155>. Published January 30, 2019. Accessed March 18, 2019..
26. Leonard AJ, Rutherford S. Study of Venetoclax Plus DA-EPOCH-R for the Treatment of Aggressive B-Cell Lymphomas (V+DA-EPOCH-R). <https://clinicaltrials.gov/ct2/show/NCT03036904>. Published May 21, 2018. Accessed March 18, 2019.
27. Li J The Safety, Efficacy of Anti-EGFR Humanized Monoclonal Antibody Combined With Chemotherapy in Advanced Solid Tumors (HLX07Ib/II). <https://clinicaltrials.gov/ct2/show/NCT03577704>. September 3, 2018. Accessed March 18, 2019..
28. Phan J Trial of Stereotactic Hypofractionated Radioablative (HYDRA) Treatment of Laryngeal Cancer. <https://clinicaltrials.gov/ct2/show/NCT03114462>. Published November 29, 2018. Accessed March 18, 2019..

29. Loskog A Phase I/IIa Trial Evaluating Safety of LOAd703, an Armed Oncolytic Adenovirus for Pancreatic Cancer. <https://clinicaltrials.gov/ct2/show/NCT02705196>. 18, 2019. Accessed March 18, 2019..
30. Jazaeri YC. T Cell Immunotherapy for Advanced Ovarian Cancer. <https://clinicaltrials.gov/ct2/show/NCT03318900>. January 29, 2019. Accessed March 18, 2019..
31. Cheung YK, Chappell R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* 2000; 56(4): 1177–1182. [PubMed: 11129476]
32. Liu S, Yin G, Yuan Y. Bayesian data augmentation dose finding with continual reassessment method and delayed toxicity. *The annals of applied statistics* 2013; 7(4): 1837. [PubMed: 24707327]
33. Cai C, Liu S, Yuan Y. A Bayesian design for phase II clinical trials with delayed responses based on multiple imputation. *Statistics in medicine* 2014; 33(23): 4017–4028. [PubMed: 24817556]
34. Daud AI, Loo K, Pauli ML, et al. Tumor immune profiling predicts response to anti-PD-1 therapy in human melanoma. *The Journal of clinical investigation* 2016; 126(9): 3447–3452. [PubMed: 27525433]
35. Gelman A, Jakulin A, Pittau MG, Su YS, others. A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics* 2008; 2(4): 1360–1383.
36. Gilks WR, Best N, Tan K. Adaptive rejection Metropolis sampling within Gibbs sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 1995; 44(4): 455–472.
37. Little RJ, Rubin DB. *Statistical analysis with missing data*. 333. John Wiley & Sons. 2014.

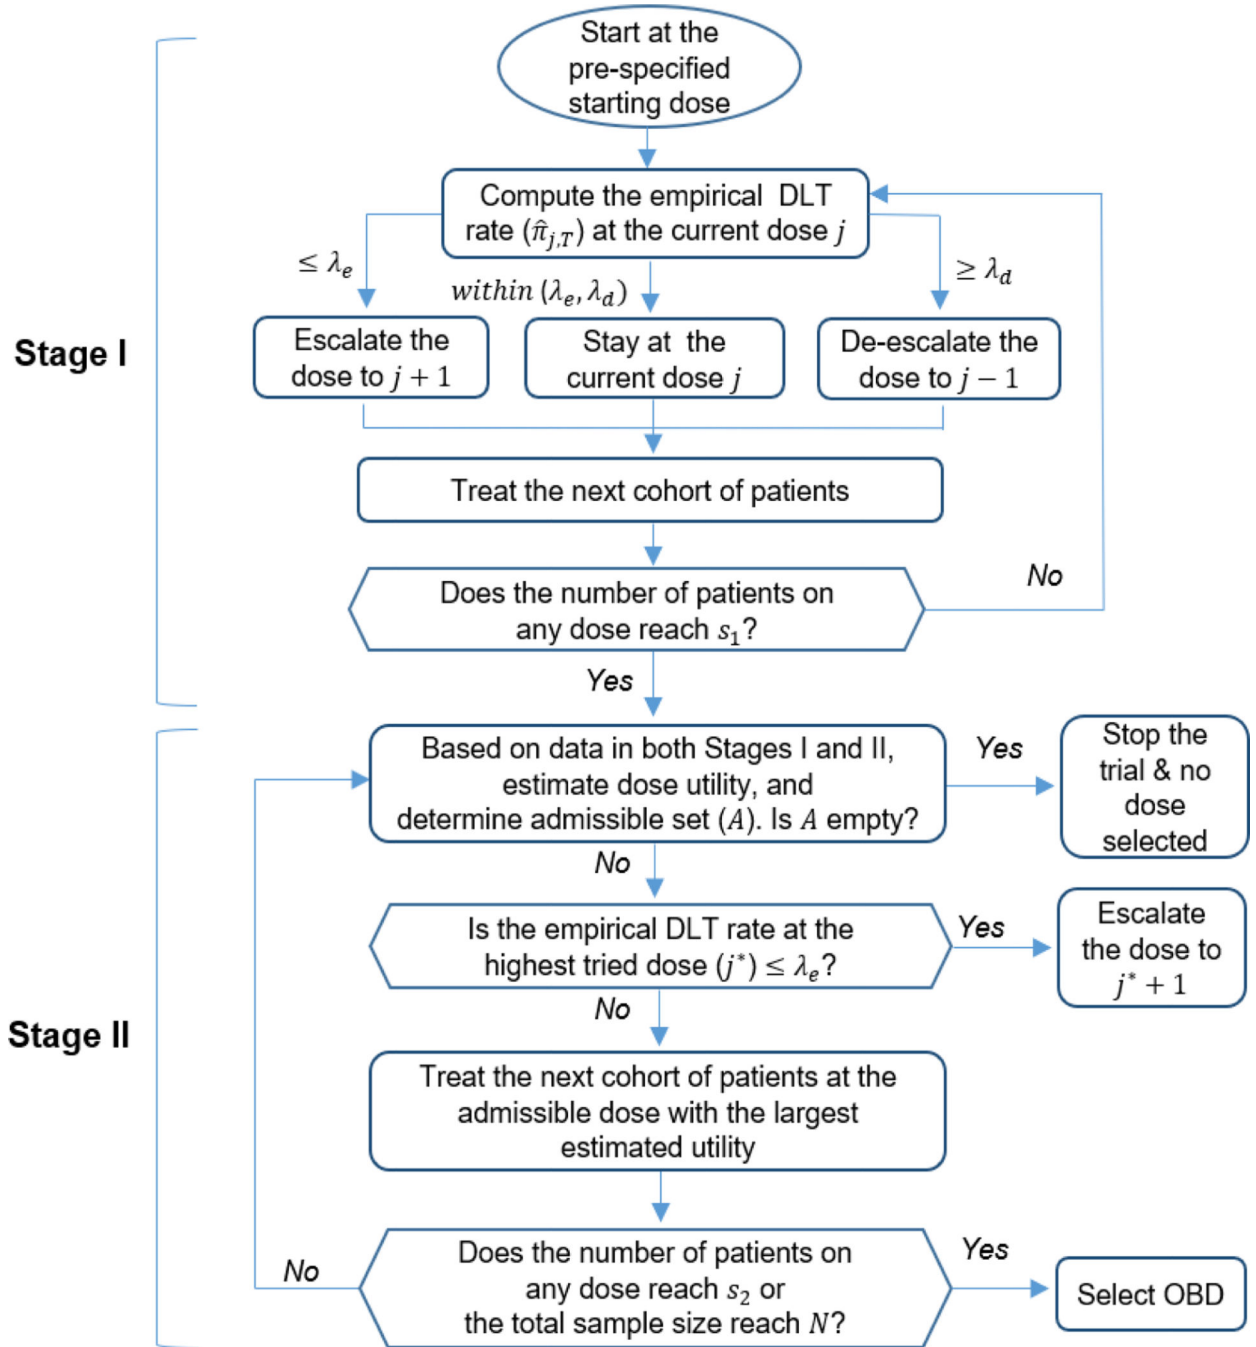


FIGURE 1.
Diagram of the U-BOIN design

Trial Setting
Simulation
Trial Protocol
Trial Conduct

Dose & Sample Size ?

Number of doses:

Starting dose level:

Cohort size :

Number of cohorts :

Stage I is completed if the number of patients treated at any dose reaches s_1 , where $s_1 =$

Stage II is completed if the number of patients treated at any dose reaches s_2 , where $s_2 =$

Utility for Risk-benefit Tradeoff ?

Number of toxicity level:

Number of efficacy level:

Toxicity

Efficacy	Yes	No
No	0	30
Yes	50	100

Admissible Criteria ?

Maximum tolerable toxicity rate ($\bar{\pi}_T$):

Minimum acceptable efficacy rate ($\underline{\pi}_E$):

A dose is deemed admissible if none of the following conditions for the true toxicity rate (π_T) and true efficacy rate (π_E) hold.

(Criterion for toxicity)

$\Pr(\pi_T > \bar{\pi}_T | data) > C_T$, where $C_T =$

(Criterion for futility)

$\Pr(\pi_E < \underline{\pi}_E | data) > C_E$, where $C_E =$

FIGURE 2.
The user interface of the U-BOIN software

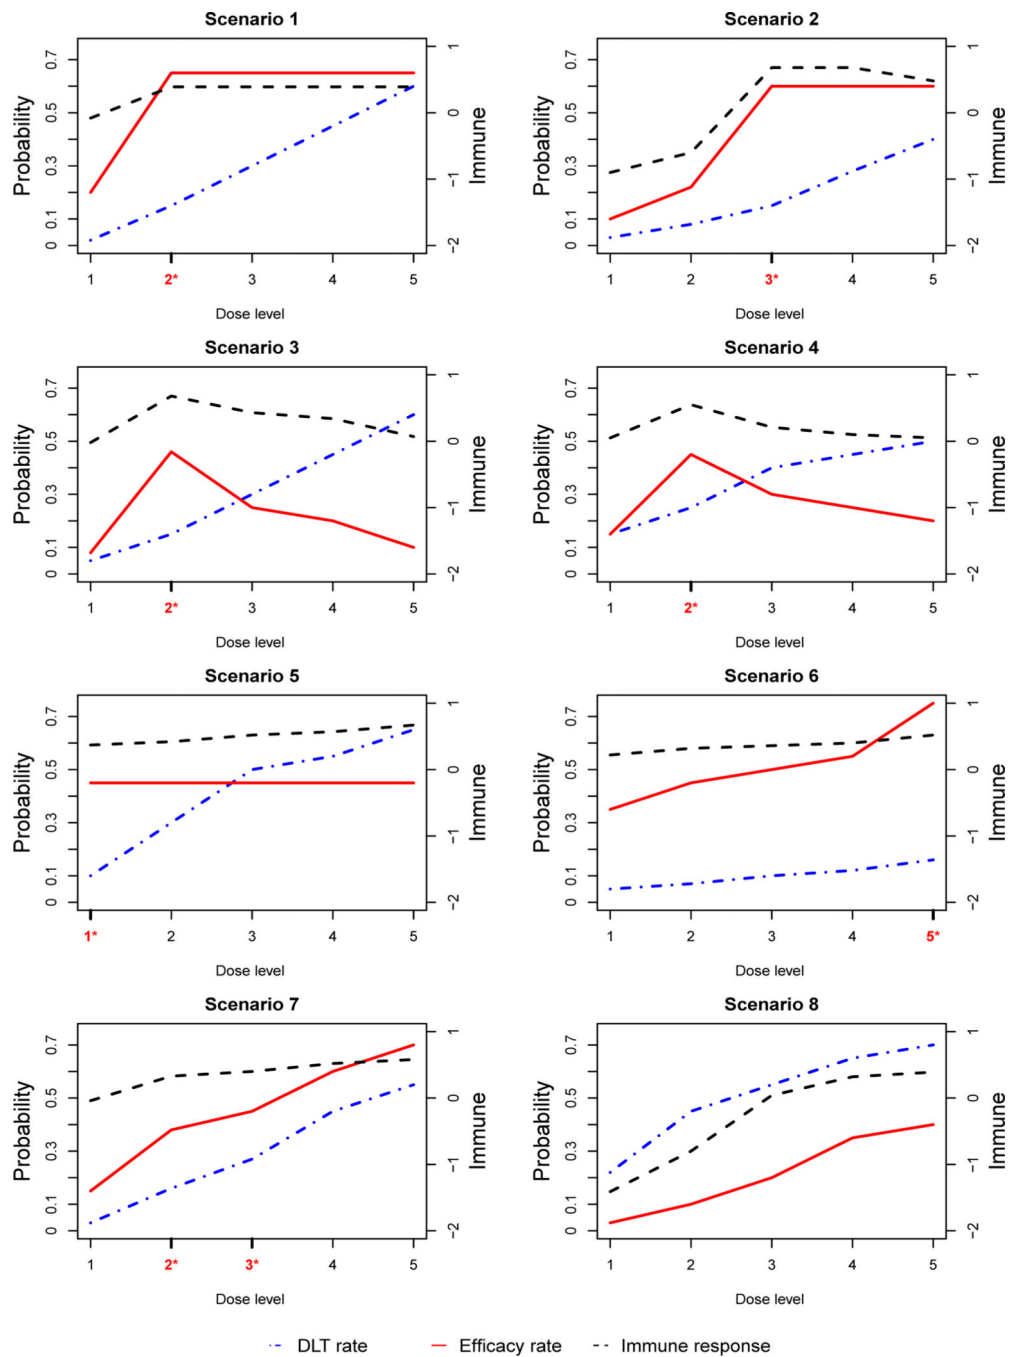


FIGURE 3. Simulation scenarios. The dash-dotted line (blue) is the dose-toxicity curve, the solid line (red) is the dose-efficacy curve, and the dashed line (black) is the dose-immune response curve. The OBD is highlighted by red asterisk in the x-axis. Simulation A considers only the efficacy and toxicity curves, while simulation B considers efficacy, toxicity, and immune response.

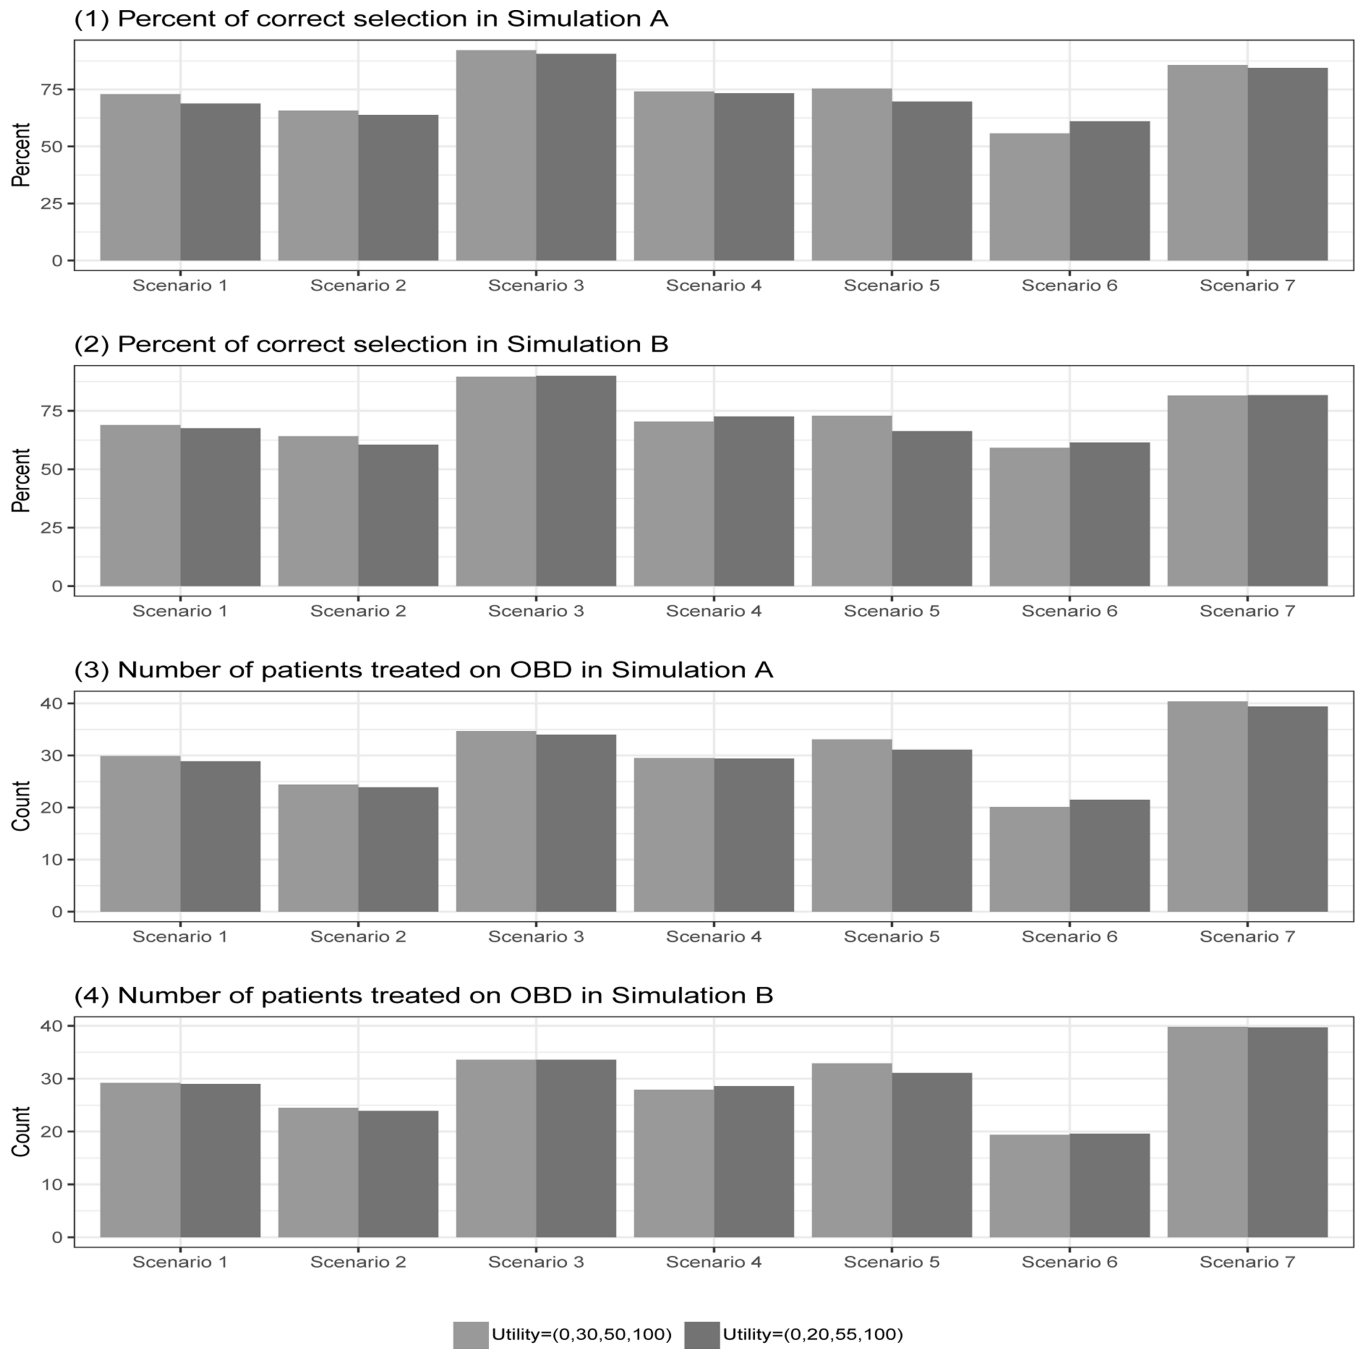


FIGURE 4. Results of sensitivity analysis for different utilities. Scenario 8 is not included, as the OBD does not exist in that scenario.

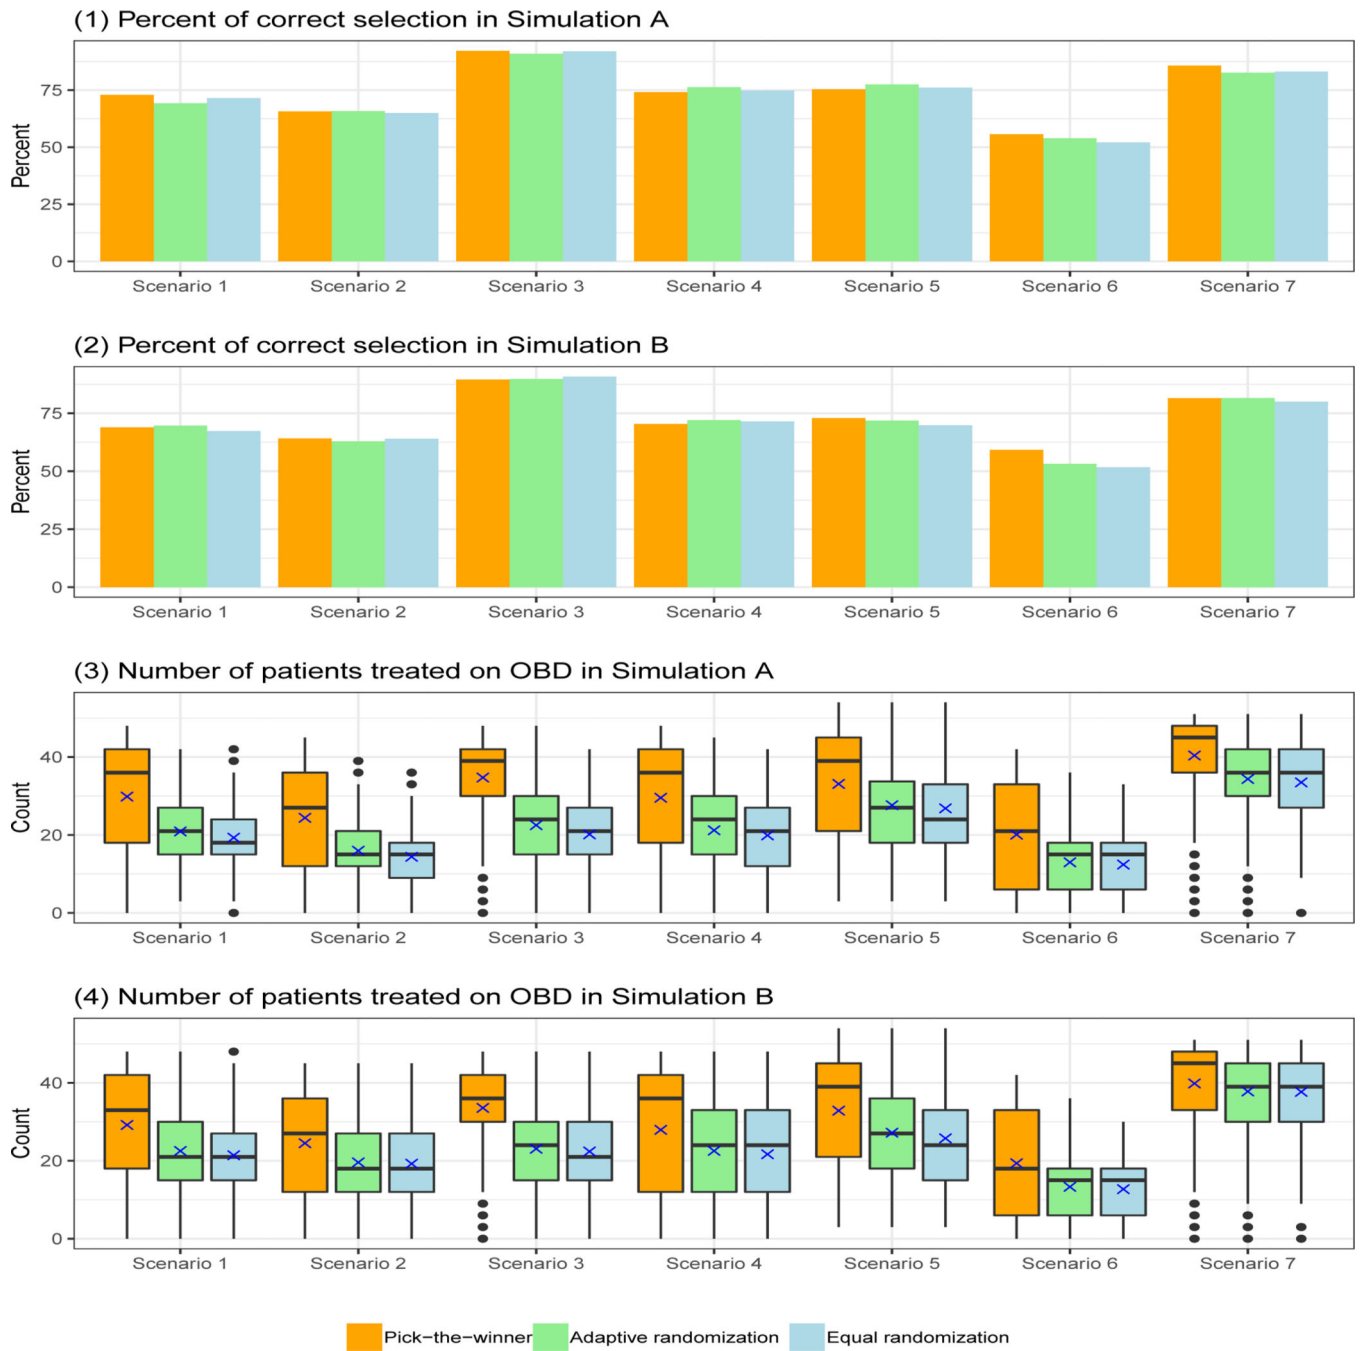


FIGURE 5. Results of sensitivity analysis for different patient allocation strategies: pick-the-winner (PW), adaptive randomization (AR), and equal randomization (ER). Scenario 8 is not included, as the OBD does not exist in that scenario.

TABLE 1

Examples of utility

(a) Example 1

	$Y_T = 1$	$Y_T = 0$
$Y_E = 0$	$\psi_1 = 0$	$\psi_2 = 30$
$Y_E = 1$	$\psi_3 = 50$	$\psi_4 = 100$

(b) Example 2

	$Y_T = 1$	$Y_T = 0$
$Y_E = 0$	$\psi_1 = 0$	$\psi_2 = 30$
$Y_E = 1$	$\psi_3 = 65$	$\psi_4 = 100$

(b) Example 3

	$Y_T = \text{severe}$	$Y_T = \text{moderate}$	$Y_T = \text{minor}$
$Y_E = \text{PD}$	$\psi_1 = 0$	$\psi_2 = 15$	$\psi_3 = 30$
$Y_E = \text{SD}$	$\psi_4 = 0$	$\psi_5 = 30$	$\psi_6 = 50$
$Y_E = \text{PR/CR}$	$\psi_7 = 15$	$\psi_8 = 45$	$\psi_9 = 100$

Note. PD: partial disease, SD: stable disease, PR: partial remission, CR: complete remission.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 2

Dose escalation and de-escalation boundaries of the BOIN design

Boundaries	Target DLT rate (ϕ_T)					
	0.15	0.20	0.25	0.30	0.35	0.40
λ_c (escalation)	0.118	0.157	0.197	0.236	0.276	0.316
λ_d (de-escalation)	0.179	0.238	0.298	0.358	0.419	0.480

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 3

Results of Simulation A, including the selection percentage (selection %), the average number of patients treated at each dose (No. of patients), and the percentage of early stopping. The optimal biological dose (OBD) is bolded. In scenario 8, the OBD does not exist, and thus the percentage of early stopping is bolded.

Design		Dose Level					% of early stopping
		1	2	3	4	5	
Scenario 1							
	DLT rate	0.02	0.15	0.30	0.45	0.60	
	Efficacy rate	0.20	0.65	0.65	0.65	0.65	
	Utility	43.0	69.0	63.0	56.0	50.0	
EffTox	Selection %	2.0	50.0	45.0	2.0	0.0	0.0
	No. of patients	4.3	22.7	23.8	2.8	0.4	
U-BOIN	Selection %	1.7	72.9	22.4	2.8	0.0	0.2
	No. of patients	6.2	29.9	13.8	3.5	0.5	
Scenario 2							
	DLT rate	0.03	0.08	0.15	0.28	0.40	
	Efficacy rate	0.10	0.22	0.60	0.60	0.60	
	Utility	36.0	43.0	66.0	60.0	55.0	
EffTox	Selection %	0.0	4.0	60.0	29.0	7.0	0.0
	No. of patients	3.4	4.9	26.4	14.2	5.1	
U-BOIN	Selection %	1.1	3.2	65.7	24.9	4.3	0.8
	No. of patients	4.9	7.5	24.4	12.7	4.3	
Scenario 3							
	DLT rate	0.05	0.15	0.30	0.45	0.60	
	Efficacy rate	0.08	0.46	0.25	0.20	0.10	
	Utility	34.0	56.0	37.0	29.0	18.0	
EffTox	Selection %	12.0	70.0	10.0	1.0	0.0	7.0
	No. of patients	10.8	25.5	11.5	2.9	1.0	
U-BOIN	Selection %	1.2	92.2	4.0	0.4	0.0	2.1
	No. of patients	6.6	34.7	8.9	2.6	0.4	
Scenario 4							
	DLT rate	0.15	0.25	0.40	0.45	0.50	
	Efficacy rate	0.15	0.45	0.30	0.25	0.20	
	Utility	36.0	52.0	36.0	32.0	27.0	
EffTox	Selection %	36.0	47.0	5.0	2.0	2.0	9.0
	No. of patients	19.4	19.0	7.5	2.6	2.3	
U-BOIN	Selection %	11.9	74.1	3.7	0.4	0.0	9.9
	No. of patients	15.0	29.5	5.1	1.1	0.2	
Scenario 5							
	DLT rate	0.10	0.30	0.50	0.55	0.65	

Design		Dose Level					% of early stopping
		1	2	3	4	5	
	Efficacy rate	0.45	0.45	0.45	0.45	0.45	
	Utility	58.0	50.0	42.0	40.0	36.0	
EffTox	Selection %	69.0	27.0	2.0	0.0	0.0	2.0
	No. of patients	29.7	17.3	4.9	1.0	0.3	
U-BOIN	Selection %	75.4	22.8	1.5	0.2	0.0	0.2
	No. of patients	33.1	16.5	3.7	0.5	0.1	
Scenario 6							
	DLT rate	0.05	0.07	0.10	0.12	0.16	
	Efficacy rate	0.35	0.45	0.50	0.55	0.75	
	Utility	53.0	59.0	61.0	64.0	75.0	
EffTox	Selection %	10.0	12.0	26.0	24.0	29.0	0.0
	No. of patients	8.6	7.9	14.2	11.1	12.1	
U-BOIN	Selection %	5.9	11.7	13.1	13.6	55.7	0.0
	No. of patients	7.0	8.8	9.0	9.1	20.1	
Scenario 7							
	DLT rate	0.03	0.16	0.27	0.45	0.55	
	Efficacy rate	0.15	0.38	0.45	0.60	0.70	
	Utility	40.0	51.0	51.0	53.0	55.0	
EffTox	Selection %	1.0	33.0	54.0	9.0	1.0	2.0
	No. of patients	4.2	13.8	25.9	7.0	2.3	
U-BOIN	Selection %	2.0	45.0	41.0	10.0	1.0	1.0
	No. of patients	5.1	20.4	20.0	6.7	1.2	
Scenario 8							
	DLT rate	0.22	0.45	0.55	0.65	0.70	
	Efficacy rate	0.03	0.10	0.20	0.35	0.40	
	Utility	25.0	23.0	25.0	30.0	31.0	
EffTox	Selection %	1.0	6.0	6.0	0.0	0.0	87.0
	No. of patients	4.6	6.2	6.3	2.3	1.2	
U-BOIN	Selection %	0.8	5.5	1.7	0.0	0.0	92.0
	No. of patients	14.3	9.7	1.6	0.1	0.0	

TABLE 4

Results of Simulation B, including the selection percentage (selection %), the average number of patients treated at each dose (No. of patients), the percentage of early stopping, and the trial duration. The optimal biological dose (OBD) is bolded. In scenario 8, the OBD does not exist, and thus the percentage of early stopping is bolded.

Design		Dose Level					% of early stopping	Duration (month)
		1	2	3	4	5		
Scenario 1								
EffTox	Selection %	2.0	50.0	45.0	2.0	0.0	0.0	45.9
	No. of patients	4.3	22.7	23.8	2.8	0.4		
U-BOIN	Selection %	0.9	68.9	24.1	2.9	0.3	2.8	20.8
	No. of patients	6.4	29.2	13.7	3.6	0.5		
Scenario 2								
EffTox	Selection %	0.0	4.0	60.0	29.0	7.0	0.0	43.8
	No. of patients	3.4	4.9	26.4	14.2	5.1		
U-BOIN	Selection %	0.0	0.2	64.1	22.4	4.9	8.3	20.0
	No. of patients	4.3	6.1	24.5	12.1	4.1		
Scenario 3								
EffTox	Selection %	12.0	70.0	10.0	1.0	0.0	7.0	45.6
	No. of patients	10.8	25.5	11.5	2.9	1.0		
U-BOIN	Selection %	0.8	89.6	6.0	0.8	0.0	2.8	20.7
	No. of patients	6.7	33.6	9.5	2.9	0.4		
Scenario 4								
EffTox	Selection %	36.0	47.0	5.0	2.0	2.0	9.0	47.8
	No. of patients	19.4	19.0	7.5	2.6	2.3		
U-BOIN	Selection %	6.8	70.4	3.4	0.4	0.0	19.1	18.8
	No. of patients	13.6	27.9	4.7	1.0	0.2		
Scenario 5								
EffTox	Selection %	69.0	27.0	2.0	0.0	0.0	2.0	50.0
	No. of patients	29.7	17.3	4.9	1.0	0.3		
U-BOIN	Selection %	72.9	22.7	2.2	0.2	0.0	2.0	20.7
	No. of patients	32.9	16.4	3.7	0.5	0.0		
Scenario 6								
EffTox	Selection %	10.0	12.0	26.0	24.0	29.0	0.0	43.4
	No. of patients	8.6	7.9	14.2	11.1	12.1		
U-BOIN	Selection %	4.2	10.3	12.2	13.6	59.2	0.4	21.0
	No. of patients	6.5	8.6	9.5	9.9	19.4		
Scenario 7								
EffTox	Selection %	33.0	45.0	45.0	43.0	40.0		46.5
	No. of patients	1.0	33.0	54.0	9.0	1.0		
U-BOIN	Selection %	5.0	45.6	33.1	10.5	1.7	4.1	20.7
	No. of patients	8.5	22.7	15.6	5.2	0.8		

Design		Dose Level					% of early stopping	Duration (month)
		1	2	3	4	5		
Scenario 8								
EffTox	Selection %	1.0	6.0	6.0	0.0	0.0	87.00	16.6
	No. of patients	4.7	6.1	6.3	2.4	1.3		
U-BOIN	Selection %	0.0	0.4	0.4	0.0	0.0	99.10	8.8

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript