Received: 21 April 2020 | Revised: 17 November 2020 | Accepted: 6 February 2021

DOI: 10.1002/sim.8922

RESEARCH ARTICLE

Statistics in Medicine WILEY

uTPI: A utility-based toxicity probability interval design for phase I/II dose-finding trials

Haolun Shi[1] | **Jiguo Cao**[1] | **Ying Yuan**[2] | **Ruitao Lin**[2]

[1]Department of Statistics and Actuarial Science, Simon Fraser University, Burnaby, British Columbia, Canada

[2]Department of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, Texas, USA

Correspondence
Ruitao Lin, Department of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, TX 77030, USA.
Email: rlin@mdanderson.org

Funding information
National Institutes of Health, Grant/Award Numbers: P30 CA016672, P50 CA098258, P50 CA221703

Unlike chemotherapy, the maximum tolerated dose (MTD) of molecularly targeted agents and immunotherapy may not pose significant clinical benefit over the lower doses. By simultaneously considering both toxicity and efficacy endpoints, phase I/II trials can identify a more clinically meaningful dose for subsequent phase II trials than traditional toxicity-based phase I trials in terms of risk-benefit tradeoff. To strengthen and simplify the current practice of phase I/II trials, we propose a utility-based toxicity probability interval (uTPI) design for finding the optimal biological dose, based on a numerical utility that provides a clinically meaningful, one-dimensional summary representation of the patient's bivariate toxicity and efficacy outcome. The uTPI design does not rely on any parametric specification of the dose-response relationship, and it directly models the dose desirability through a quasi binomial likelihood. Toxicity probability intervals are used to screen out overly toxic dose levels, and then the dose escalation/de-escalation decisions are made adaptively by comparing the posterior desirability distributions of the adjacent levels of the current dose. The uTPI design is flexible in accommodating various dose desirability formulations, while only requiring minimum design parameters. It has a clear decision structure such that a dose-assignment decision table can be calculated before the trial starts and can be used throughout the trial, which simplifies the practical implementation of the design. Extensive simulation studies demonstrate that the proposed uTPI design yields desirable as well as robust performance under various scenarios.

KEYWORDS

Bayesian quasi likelihood, dose desirability, dose finding, optimal biological dose, phase I/II trials, utility

1 | INTRODUCTION

The development of immunotherapy and molecularly targeted agents is one of the most innovative and remarkable advances in modern cancer treatment. It not only revolutionizes the oncological therapeutics and patient care in almost every kind of cancer but it also has led to a major paradigm shift from a drug development perspective. Unlike conventional chemotherapy, which involves cytotoxic agents that treat cancer by inhibiting and destroying the growth mechanism of tumor cells, immunotherapy uses and strengthens a patient's innate immune system to combat cancer.

Compared with the conventional cytotoxic therapies, this novel class of treatments has distinct efficacy and toxicity profiles, that is, the optimal therapeutic effect is not necessarily achieved at the highest tolerated dose level. These characteristics pose new challenges to the design of early-phase clinical trials.

Conventional phase I designs for chemotherapy commonly aim to identify the maximum tolerated dose (MTD), defined as the dose level that has the toxicity probability closest to the prespecified target level. This relies on the assumption that toxicity and efficacy increase monotonically with the dose such that a meaningful therapeutic treatment effect can be achieved at the MTD. However, the dose-response monotonicity assumption might not hold for immunotherapy and molecularly targeted agents, of which the dose-efficacy curves may take various shapes.[1,2] For example, the dose-efficacy curve of an immunotherapy may reach a plateau at a dose level lower than the MTD, that is, further improvement in efficacy is no longer possible after the drug exposure reaches a certain saturation threshold. Therefore, for immunotherapy and targeted agents, the MTD may no longer be a meaningful recommended phase 2 dose for further studies.[1] Novel early-phase dose-finding designs that accommodate both efficacy and toxicity endpoints are thus needed, with the primary goal of identifying the optimal biological dose (OBD), defined as a safe dose that renders the greatest utility and benefit to cancer patients.[3,4]

This article is motivated by a recent early-phase immunotherapy trial of a combination of two HER-2 peptide B-cell epitope vaccines.[5] This vaccine therapy is capable of functioning as "endogenous" trastuzumab and pertuzumab, which have been approved for clinical use but are not well tolerated in cancer patients. The primary objective of the trial was to assess the safety profile, as well as to identify the optimal immunologic/biological dose of the vaccine. In total, four dose levels of the peptide vaccine were considered, and eligible patients with advanced solid tumors were enrolled. In order to better evaluate the OBD, patients were required to receive a total of three consecutive inoculations of the vaccine therapy at 3-week intervals. Although the primary objective was to identify the OBD, the standard 3+3 design based only on toxicity data was adopted in the dose-escalation phase. The trial showed that the vaccine was well tolerated overall with no dose-limiting toxicities observed. In terms of efficacy, an immunogenic response was defined if the patient had achieved a response graded as at least "++" in both MVF-HER-2 (266-296) and MVF-HER-2 (597-626). Each of the four doses had six patients who had completed the full regimens and were evaluable for efficacy. Among the dose levels, dose 1 had zero response, dose 2 (with four responses observed) yielded the highest number of responses. The higher doses 3 and 4, respectively, had three and one responses and were not as good as dose 2, indicating a non-monotone dose-efficacy relationship. As a result, dose level 2 was finally selected as the OBD. According to Bekaii-Saab et al,[5] most of the patients (65.3%) were allocated to the higher dose levels (ie, doses 3 and 4), while the final recommended dose level (dose 2) only treated 18% of the enrolled patients. This example shows a typical toxicity-only design flaw in finding the OBD, of which the identification depends on both toxicity and efficacy jointly instead of toxicity solely. When all doses were well tolerable, the toxicity-only designs would always recommend dose escalation without considering the dose-efficacy relationship. Such a practice usually leads to a higher likelihood to miss the OBD, while also causing more patients to be treated at subtherapeutic doses. The non-monotone dose-efficacy relationship, where efficacy improvement can no longer be observed beyond a certain dose level, is often seen in the early-phase trials of targeted therapies, for example, Horn et al[6] and Patil et al.[7]

To more accurately identify the OBD, we propose an optimal biological dose-based toxicity probability interval (utility-based toxicity probability interval [uTPI]) design, based on a utility function that summarizes the desirability of each patient's toxicity and efficacy outcomes. The proposed design is built upon the keyboard design,[8] which is also known as the modified toxicity probability interval design.[9] The uTPI design defines the dose desirability using the joint toxicity and efficacy probabilities. It models the numerical utility outcomes observed at each dose based on a quasi-binomial likelihood,[10] which simplifies the posterior sampling and alleviates the computational burden. Following the idea of using intervals for dose finding in the toxicity probability interval (TPI) design,[11] the cumulative cohort design[12] and the keyboard design,[8] the uTPI design also constructs a toxicity probability interval and a desirability interval for decision making. In particular, the toxicity probability intervals are adopted to screen out overly toxic dose levels, and then the dose escalation/de-escalation decisions are made adaptively by comparing the posterior desirability distributions of the adjacent levels of the current dose. The uTPI design does not rely on any parametric dose-response model for decision making and its dose escalation/de-escalation scheme can be feasibly enumerated onto a pretabulated decision table. This simplifies the design implementation and trial conduct.

Various phase I/II trial designs that simultaneously model toxicity and efficacy endpoints to identify the OBD have been developed in recent years. In general, depending on whether or not a parametric curve is specified for the dose-response relationship, the existing methods for phase I/II trials can be categorized into two classes: curve-based designs and curve-free designs. The curve-based designs posit parametric models to characterize the relationship between

the dose and the joint toxicity and efficacy outcomes, and they can usually produce desirable operating characteristics after careful calibration.[3,13] Braun[14] extended the continual reassessment method to simultaneously account for both toxicity and efficacy. Thall and Cook[15] proposed a Bayesian model for the bivariate toxicity and efficacy endpoints and established priors based on efficacy-toxicity tradeoff contours (EffTox). Yin et al[16] developed a phase I/II design that uses the efficacy/toxicity odds ratio to model the tradeoff between the two endpoints. Zang et al[17] introduced an isotonic regression-based (Iso) design to identify the OBD. Mozgunov and Jaki[18] classified the bivariate toxicity and efficacy outcomes into three categories and used an information theoretic approach to determine the best dose without the assumption of monotonicity. Other approaches along the direction of bivariate modeling for phase I/II designs also have been proposed, such as Hoering et al,[19] Riviere et al,[20] and Liu and Johnson.[21]

As an alternative to curve-based designs, curve-free designs have been proposed to simplify the use of novel adaptive designs in real trials while making them more robust. Unlike the curve-based designs, this class of designs does not assume any parametric dose-response curve for dose finding. Li et al[22] proposed a toxicity and efficacy probability interval (TEPI) design that separately models the toxicity and efficacy as it adaptively assigns new patients to a dose level with a promising efficacy rate and a tolerable toxicity profile. Lin and Yin[23] developed a robust dose-finding design (hereafter referred to as STEIN) using the pool-adjacent-violators algorithm and model averaging. Takeda et al[24] presented a Bayesian optimal interval design (hereafter referred to as BOIN-ET) that accommodates both efficacy and toxicity endpoints. To incorporate the possible correlation between the toxicity and efficacy outcomes, Zhou et al[25] recently developed a two-stage utility-based Bayesian optimal interval (U-BOIN) design using a Multinomial-Dirichlet model to jointly model toxicity and efficacy.

Curve-free designs usually require less intensive computation than curve-based designs, while the performances are still competitive. However, the dose selection criteria of most of the existing curve-free designs are centered around efficacy, aiming to find the most efficacious dose among all the doses with tolerable toxicities; whereas very few designs make dose-assignment decisions based upon a utility function that explicitly quantifies the toxicity-efficacy preference. As a curve-free design, the proposed uTPI design seeks to fill this gap in research by taking into account both toxicity and efficacy outcomes in decision making through a defined dose desirability function from the start of the trial, instead of treating them independently. Similar to many other curve-free designs,[8,22,24,25] the uTPI design has a transparent decision structure such that a dose-assignment decision table can be calculated ahead of the trial start and can be readily used throughout the trial. To the best of our knowledge, the other curve-free design that also uses a utility function to account for the tradeoff between toxicity and efficacy is the U-BOIN design.[25] It is formulated as a two-stage design with the first stage being toxicity-based, which usually requires a relatively larger sample size to ensure desirable performance.

The rest of this article is organized as follows. In Section 2, we present the methodology of the uTPI design for phase I/II dose-finding trials. The design implementation and illustration can be found in Section 3. We conduct numerical studies to assess the operating characteristics of the proposed uTPI design in Section 4, and Section 6 concludes this article with some remarks. Additional simulation studies and discussions can be found in the Supplementary Material.

2 | UTILITY-BASED TOXICITY PROBABILITY INTERVAL DESIGN

2.1 | Bivariate toxicity-efficacy outcomes

For illustrative purposes, we only consider the case when both the outcomes are binary and can be observed quickly before the arrival of new patients, though the proposed method can be readily generalized to accommodate ordinal outcomes and delayed responses (see Section 5). For binary toxicity and efficacy endpoints, there are four different possible events for each patient: both toxicity and efficacy (T, E), toxicity but no efficacy (T, E^c), efficacy but no toxicity (T^c, E), and neither toxicity nor efficacy (T^c, E^c). The joint probability of each outcome can be specified based on the following 2×2 probability table:

	Toxicity (π_T)	**No toxicity ($1 - \pi_T$)**
Efficacy (π_E)	p_1	p_3
No efficacy ($1 - \pi_E$)	p_2	p_4

with the summation of the four probabilities equal to 1. Based on these four nonoverlapping events, the marginal probability for efficacy is, $\pi_E = p_1 + p_3$, and for toxicity, it is $\pi_T = p_1 + p_2$.

Based on the bivariate outcomes, we define a four-level multinomial random variable Y, with $Y = 1$, if (T, E); $Y = 2$, if (T, E^c); $Y = 3$, if (T^c, E); and $Y = 4$, if (T^c, E^c). Suppose that n patients have been treated at a given dose with x_k patients having experienced outcome $Y = k$, where $n = \sum_{k=1}^{4} x_k$. In addition, let $D = (x_1, x_2, x_3, x_4)$ denote the observed data, and let $n_T = x_1 + x_2$ and $n_E = x_1 + x_3$ denote the numbers of observed toxicity and efficacy outcomes, respectively. Under the Bayesian paradigm and assuming a Dirichlet prior for (p_1, p_2, p_3, p_4), the marginal posterior distributions of π_T and π_E follow the Beta-Binomial model,

$$\pi_T|D \sim \text{Beta}(\alpha_T + n_T, \beta_T + n - n_T),$$
$$\pi_E|D \sim \text{Beta}(\alpha_E + n_E, \beta_E + n - n_E),$$

where α_T, α_E, β_T, and β_E are hyperparameters. In this article, we propose to specify $\alpha_E = \alpha_T = \beta_E = \beta_T = 1$, that is, the priors for π_T and π_E are both non-informative Unif(0, 1) distributions.

2.2 | Modeling dose desirability under a quasi-binomial distribution

To balance the toxicity-efficacy tradeoffs in dose finding, we define a numerical, one-dimensional utility that characterizes the desirability of each of the four possible outcomes of Y. The utility function should be negatively related to toxicity and positively related to efficacy. In general, any utility function that can project the bivariate toxicity-efficacy outcomes to the (0, 1) domain is suitable for the proposed method. For illustrative purposes, we hereby define the numerical utility based on the following 2×2 risk-benefit preference table, and this further renders the quantification of dose desirability.[25-27] Generalization of the proposed method using other specifications of the dose desirability can be found in the Supplementary Material.

	Toxicity	No toxicity
Efficacy	w_1	w_3
No efficacy	w_2	w_4

In general, $w_3 \in [0, 1]$ should have the maximum value among the four specified values, since the outcome (T^c, E) is the most desirable, while $w_2 \in [0, 1]$ should be the smallest, as it indicates the worst-case scenario. For the sake of simplicity, we take $w_3 = 1$ and $w_2 = 0$, that is, we assume a value of 1 for the best outcome and a value of 0 for the worst outcome. The remaining two quantities, w_1 and w_4, should be non-negative and lie in $[0, 1]$. A specification of $w_1 > w_4$ suggests that enjoying the drug's efficacious treatment effect, while suffering from the negative effect of toxicity is better than having no effects at all, indicating that efficacy is taken to be more important than toxicity; vice versa for $w_1 < w_4$.

Based on the 2×2 utility preference table, the observed one-dimensional numerical utility value for each patient can be written as $U = \sum_{k=1}^{4} 1\{Y = k\} w_k$, where $1\{\cdot\}$ denotes the indicator function with the value equal to 1 if the event is observed, and 0 otherwise. Therefore, the expected utility of the patient's outcome, denoted as EU, is

$$EU = p_3 + p_1 w_1 + p_4 w_4. \tag{1}$$

In the dose-finding trials, the joint probability vector (p_1, p_2, p_3, p_4) depends on the dose level. In other words, the expected value EU can be used to characterize the desirability of the dose at which the patient is treated. Thus, we use EU to represent the dose desirability in this article. As a note, when $w_1 + w_4 = 1$, the dose desirability can be further simplified as

$$EU = \pi_E w_1 + (1 - \pi_T) w_4, \tag{2}$$

so that it quantifies the tradeoffs between the marginal toxicity and efficacy probabilities of a certain dose. Furthermore, when $w_1 = 1$ and $w_4 = 0$, maximizing EU is equivalent to maximizing the marginal efficacy probability π_E.

In the dose-finding rule of the proposed design, we use the posterior desirability distribution of EU, denoted as $f(EU \mid D)$, to decide the dose level for new patients. To further simplify the computation, we treat the numerical utility outcome U as a pseudo binomial observation and use quasi-likelihood for inference.[10] In particular, let u_i denote the observed numerical utility for patient i, $i = 1, \ldots, n$, and depending on the bivariate outcome of efficacy and toxicity, u_i can take a value from (w_1, w_2, w_3, w_4), which are all between 0 and 1. The quasi-likelihood for EU is constructed as

$$L(D|EU) = \prod_{i=1}^{n} EU^{u_i}(1 - EU)^{1-u_i},$$

where we treat EU as the unknown parameter. Under some regularity conditions, the above quasi binomial likelihood can produce consistent inference for dose desirability.[28,29] Using the beta-binomial model, we assign a beta prior for EU, that is, $EU \sim \text{Beta}(\alpha_U, \beta_U)$. As a result, the posterior distribution of EU is still a beta distribution,

$$EU|D \sim Beta\left(\alpha_U + \sum_{i=1}^{n} u_i, \beta_U + n - \sum_{i=1}^{n} u_i\right). \tag{3}$$

The posterior mean of EU is $\widehat{EU} = \frac{\alpha_U + \sum_{i=1}^{n} u_i}{\alpha_U + \beta_U + n}$. Under the 2×2 numerical utility preference table, we have $\sum_{i=1}^{n} u_i = \sum_{k=1}^{4} w_k x_k$. In the special case of $w_1 + w_4 = 1$, the posterior distribution of EU can be further simplified as $EU|D \sim \text{Beta}(\alpha_U + n_E w_1 + (n - n_T)w_4, \beta_U + (n - n_E)w_1 + n_T w_4)$, which depends on the number of patients n, the number of toxicity outcomes n_T, and the number of efficacy outcomes n_E. By default, we consider $\alpha_U = \beta_U = 1$, that is, a uniform distribution, to reflect a non-informative prior for EU. When there is prior information on EU, (α_U, β_U) can be calibrated by specifying the prior means and variances. It is worth noting that the marginally uniform priors for π_E and π_T are not equivalent to a uniform prior for EU. We choose to adopt a uniform prior for EU for the sake of conjugacy and computational simplicity. As long as the prior distributions are non-informative, we expect that such a discrepancy would not have a notable impact on the posterior estimation, as the information is primarily dominated by the data rather than the prior.

As a final remark, according to the construction (1), the conventional approach is to derive $f(EU \mid D)$ from the convolution of the joint posterior distribution of (p_1, p_2, p_3, p_4). However, such a convolutional approach involves calculation of (multiple) integrals (even in the simpler case of (2)), which is computationally intensive. In addition, it fails to render a regular and tractable posterior distribution of EU, which may possibly complicate the trial implementation. To simulate 10 000 trials, the convolutional integral-based approach takes more than 10 hours, whereas our proposed simplified computation of $f(EU|D)$ uses a standard Beta posterior and can be done in 7 seconds without deterioration in performance. Such a difference in computational speed is relevant when assessing and reporting the operating characteristics of the design through a large number of simulated trials. The details of this benchmark comparison can be found in Section S.3 of the Supplementary Material.

2.3 | Chessboard design

Consider a phase I/II dose-finding trial with J prespecified doses. We introduce the subscript j to indicate dose level j, $j = 1, \ldots, J$. In particular, we assume that $\pi_{1,T} < \ldots < \pi_{J,T}$, that is, the toxicity probability increases with the dose level. In contrast, we consider the case that the efficacy probability $\pi_{j,E}$ does not necessarily increase with j, but may increase initially and then reach a plateau from which no improvement or even decreasing efficacy may be seen with an increasing dose. The primary objective of the trial is to find the OBD that maximizes the dose desirability EU, with the constraint that it is not overly toxic. Specifically, let ϕ denote the target toxicity rate that is specified by the clinicians, let $j^\dagger$ denote the MTD. Since the OBD should not exceed the MTD, the OBD is defined as

$$j^* = \underset{j \le j^\dagger}{\arg\max}\{EU_j\}.$$

In other words, the OBD is the dose level that is no higher than MTD and yields the highest desirability.

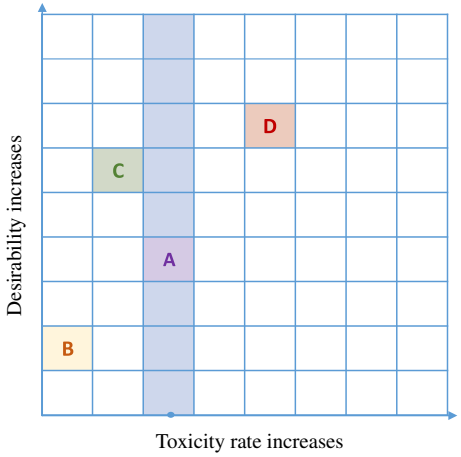

FIGURE 1 Dose-finding rule of the uTPI design based on a chessboard. The shaded interval corresponds to the target toxicity probability interval k^*. (1) If C is the current dose (or if A is the current dose level and $n_A < N^* = 9$), and suppose that B is the next lower dose and D is the next higher dose, we will allocate the next cohort of patients to the level among $\{C, B, D\}$ (or $\{A, B, D\}$) that has the highest desirability interval. In this case, we will select dose level D for the next cohort of patients. (2) If A is the current dose and $n_A \geq N^* = 9$, and suppose that B is the next lower dose and D is the next higher dose, we will allocate the next cohort of patients to the level among $\{A, B\}$ that has the highest desirability interval. In this case, we will select dose level A for the next cohort of patients [Colour figure can be viewed at wileyonlinelibrary.com]

To identify the OBD among the given J dose levels, we propose to jointly monitor the toxicity and desirability for each dose level. We partition the unit interval $(0, 1)$ into a series of intervals of equal width ϵ for the marginal toxicity probability $\pi_{j,T}$. Similarly, we obtain a series of intervals of equal width δ for EU_j. We denote the resulting intervals as $\mathcal{I}_{k,T}, k = 1, \ldots, K_T$, for toxicity, and $\mathcal{I}_{k,U}, k = 1, \ldots, K_U$, for desirability. For example, when the target toxicity rate $\phi = 0.3$, we specify 10 toxicity intervals of width $\epsilon = 0.1$, which results in $\mathcal{I}_{1,T} = [0, 0.1), \ldots, \mathcal{I}_{K_T,T} = [0.9, 1]$ with $K_T = 10$. We also construct 10 desirability intervals of width $\delta = 0.1$, leading to $\mathcal{I}_{1,U} = [0, 0.1), \ldots, \mathcal{I}_{K_U,U} = [0.9, 1]$ with $K_U = 10$. The parameter ϵ (or δ) denotes the indifference margin for toxicity (or desirability), which means that any two dose levels whose toxicity probabilities (or EU's) lying within the same toxicity (or desirability) interval are treated indifferently in terms of toxicity or expected desirability. The reason for specifying the indifference intervals is that the sample size in early-phase dose-finding trials is typically small, and thus it is difficult to differentiate the two dose levels that have similar toxicity probabilities (or EU). The construction of the toxicity/desirability intervals mimics the idea of specifying equally-spaced toxicity intervals in the keyboard design,[8] and we extend their proposed "keys" in a one-dimensional keyboard (of toxicity intervals) to "squares" on a two-dimensional chessboard (of both toxicity and desirability intervals), as shown in Figure 1. A similar idea of using probability interval for modeling has also been proposed in the toxicity probability interval design[11] and the cumulative cohort design.[12] In addition, the construction of equal-width probability intervals also avoids the issue of Ockham's razor, as discussed in Guo et al.[9]

Based on the observed data at dose level j, we can identify the strongest toxicity interval as that possesses the largest posterior probability using the marginal posterior distribution of $\pi_{j,T}$,

$$k_j^T = \underset{k=1,\ldots,K_T}{\arg\max}\{\Pr(\pi_{j,T} \in \mathcal{I}_{k,T} | D_j)\}.$$

Similarly, we identify the strongest desirability interval that has the largest posterior probability based on (3),

$$k_j^U = \underset{k=1,\ldots,K_U}{\arg\max}\{\Pr(EU_j \in \mathcal{I}_{k,U} | D_j)\}.$$

The collection of strongest toxicity and desirability intervals (k_j^T, k_j^U) can be treated as the vector of sufficient statistics of the proposed design in determining dose escalation/de-escalation. In other words, we traverse around the entire two-dimensional chessboard, attempting to select the square, (k_j^T, k_j^U), that best approximates the estimated combination of $\pi_{j,T}$ and EU_j.

Let k^* denote the location of the highest tolerable toxicity interval inside which $\pi_{j,T}$ is indifferent to the target toxicity rate ϕ. In the aforementioned example, $k^* = 4$ with $\mathcal{I}_{k^*,T} = [0.3, 0.4)$, meaning that any dose level $\pi_{j,T} \in \mathcal{I}_{k^*,T} = [0.3, 0.4)$ can be considered as acceptable, while any level with $\pi_{j,T} > 0.4$ is deemed as excessively toxic. As a result, in our dose-finding rule, if $k_j^T \leq k^*$, then the toxicity rate at dose level j is acceptable; if $k_j^T > k^*$, then dose level j is too toxic and dose escalation is not warranted.

Under the distributional assumptions for both toxicity probability $\pi_{j,T}$ and dose desirability EU_j, as described in the preceding sections, the proposed dose finding proceeds as follows (see also Figure 1).

Dose-finding rule:

1. Treat the first cohort of patients at the lowest dose level or the user-specified level.
2. Denote the current dose level as j.

- If $k_j^T > k^*$, de-escalate to dose level $j-1$ if $j>1$, or stay at the current dose level if $j=1$.
- If $k_j^T < k^*$, choose the dose level j' from the admissible set $\{j-1, j, j+1\}$ that has the largest $k_{j'}^U$ for the next cohort of patients.
- If $k_j^T = k^*$, we further consider two cases according to the number of patients treated at the current dose level:
  - If $n_j < N^*$, choose the dose level j' from the admissible set $\{j-1, j, j+1\}$ that has the largest desirability interval $k_{j'}^U$ for the next cohort of patients.
  - If $n_j \geq N^*$, choose the dose level j' from the admissible set $\{j-1, j\}$ that has the largest desirability interval $k_{j'}^U$ for the next cohort of patients.

  When the next dose level is out of the prespecified range, treat the next cohort of patients at the current dose. In the event of a tie, that is, when the largest $k_{j'}^U$ can be found in multiple dose levels, we select the dose j' that has maximum $\Pr(U_{j'} > \overline{\mathcal{I}}_{k_{j'}^U,U} | D_{j'})$, where $\overline{\mathcal{I}}_{k_{j'}^U,U}$ denotes the upper boundary of the $k_{j'}^U$th desirability interval.

3. Step 2 is repeated until the maximum sample size is reached.

The N^* is a tuning parameter that serves as a boundary between an early stage that encourages exploring higher dose levels, and a late stage that is more conservative in dose expansion. When $k_j^T = k^*$ and $n_j \geq N^*$, that is, our estimated toxicity interval is the same as that of the MTD, it indicates that (based on a relatively large sample size) it is highly likely that the current dose is indeed the MTD, and that higher dose levels would be overly toxic according to the assumption of a monotonically increasing dose-toxicity relationship. This means that further escalation would most likely results in unfavorable toxicity responses in patients. Therefore, for the sake of patient safety, we would prefer a conservative escalation rule when $k_j^T = k^*$ with $n_j \geq N^*$, that is, no further escalation needed. On the other hand, in the case where $k_j^T = k^*$ and $n_j < N^*$, we suggest a more aggressive escalation so that more untried dose levels can be explored for small n_j. As a default, we take $N^* = 9$.

At the end of the trial, we perform the pool-adjacent-violators algorithm to the observed toxicity rate $\hat{\pi}_{j,T} = n_{j,T}/n_j$ so that the isotonically regressed estimate $\tilde{\pi}_{j,T}$ is monotonically increasing with dose level j.[30] To select the OBD, we first identify the MTD level j^{MTD} as the dose that has the estimated toxicity rate $\tilde{\pi}_{j,T}$ closest to the target toxicity rate ϕ, that is, $j^{\text{MTD}} = \underset{j=1,\ldots,J}{\arg\min} |\tilde{\pi}_{j,T} - \phi|$. Since the OBD cannot exceed the MTD, the OBD, j^{OBD}, is then chosen as the level that has the maximum posterior mean of dose desirability EU_j among the levels less than or equal to j^{MTD}, that is,

$$j^{\text{OBD}} = \underset{j \leq j^{\text{MTD}}}{\arg\max} \{\widehat{EU}_j\}. \tag{4}$$

Furthermore, when the constraint $w_1 + w_4 = 1$ is imposed, the desirability in (2) is a function of the marginal toxicity rate $\pi_{j,T}$ and efficacy rate $\pi_{j,E}$. In this case, we can additionally obtain a model averaging estimate of $\pi_{j,E}$ by borrowing information across dose levels,[23] which we denote as $\tilde{\pi}_{j,E}$. As a result, the final OBD estimate based on (2) is given by

$$j^{\text{OBD}} = \underset{j \leq j^{\text{MTD}}}{\arg\max} \left\{ \frac{\alpha_U + n_j \tilde{\pi}_{j,E} w_1 + n_j (1 - \tilde{\pi}_{j,T}) w_4}{\alpha_U + \beta_U + n_j} \right\}. \tag{5}$$

The details of the estimation of the marginal toxicity and efficacy rates via the model averaging approach can be found in the Appendix.

In practice, for patients' safety and benefit, we can impose additional rules to eliminate overly-toxic doses or doses with unexpectedly low efficacy. In particular, let ϕ and ψ denote the target toxicity and efficacy probabilities, respectively.

Dose-elimination rule:

1. If $\Pr\left(\pi_{j,T} \geq \phi | D_j\right) > c_T$, dose level j and its higher doses are considered overly toxic and will be eliminated.
2. If $\Pr\left(\pi_{j,E} \leq \psi | D_j\right) > c_E$, dose level j is considered futile and will be eliminated.

Here, c_T and c_E are the probability cutoffs whose values can be calibrated based on simulations. When all doses are eliminated from the trial, the trial should be terminated early.

The objective of the proposed design is to find a dose level that possesses reasonably low toxicity yet promising efficacy so that the toxicity-efficacy tradeoff can better balance. To accomplish this goal, our dose-finding rule adaptively selects among the adjacent doses of the current level the one that possesses the highest desirability value as the next cohort's dose for the new patients. However, at the beginning of the trial when minimal data is available, estimation of the dose desirability EU_j is highly variable. As a result, the incorporation of the desirability in the dose assignment usually leads to overly conservative escalation, with many higher yet efficacious dose levels unexplored. To circumvent such an issue and accelerate dose escalation in the early stages, we propose an ad-hoc rule that suppresses the toxicity outcome in the desirability calculation. In particular, we consider the following modified n_j-dependent posterior score,

$$EU(n_j) = \begin{cases} \pi_{j,E} w_1 + w_4, & \text{if } n_j < N^*, \\ p_{j,3} + p_{j,1} w_1 + p_{j,4} w_4, & \text{if } n_j \geq N^*. \end{cases}$$

In this way, when $n_j < N^*$, the marginal efficacy rate plays the sole role in the calculation of dose desirability, while the effect of toxicity is temporarily suppressed. When the sample size is large, that is, $n_j \geq N^*$, the toxicity-efficacy tradeoff comes into play in the decision making.

In addition, when $n_j = 0$, that is, dose level j has been untried, we assign an initial value for k_j^U as $k_j^U = (2\psi w_1 + w_4)K_U$, which corresponds to a toxicity probability of 0 (as $n_j = 0 < N^*$) and an efficacy probability of 2ψ, with ψ being the physician-specified clinically uninteresting probability. Such a specification is chosen based on extensive simulation studies and can strike a well balance between exploration (ie, exploring more untried dose levels) and exploitation (ie, staying at a reasonably good dose level). In some situations when the lower dose levels have a very high efficacy probability, the initial setting of k_j^U may prevent the uTPI design exploring higher untried higher dose levels. In this case, a larger initial value, say $k_j^U = (3\psi w_1 + w_4)K_U$, can be specified to facilitate a more exploratory search in the dose space, reducing the probability that the trial is trapped at suboptimal doses. Moreover, in the initial stage of the trial, it is possible that dose level $j+1$ has been unexplored, but cannot be selected as the next assigned dose because its initial value of k_{j+1}^U is not as high as that of j or $j-1$. To avoid such a scenario and encourage further dose exploration to the untried levels, an additional dose exploration rule that pushes the dose to the next higher level when the trial might be trapped at the current dose can be adopted. Evaluation of the operating characteristics of this optional dose exploration rule can be found in Section S.2 of the Supplementary Material. In summary, the initial value of k_j^U for untried doses together with the additional dose exploration rule control the aggressiveness of the dose exploration. For most oncology trials where the response rates are relatively low, the default setting of the uTPI without the additional exploration rule can lead to satisfactory operating characteristics. In some cases such as some vaccine trials where the response rates are particularly high, it is more desirable to assign a higher initial value to the untried doses or impose the additional dose exploration rule.

3 | DESIGN IMPLEMENTATION

3.1 | Practical consideration

To implement the proposed uTPI design, one needs to specify several design parameters, including the target toxicity rate ϕ, the clinically uninteresting efficacy rate ψ, the widths of toxicity and desirability intervals ϵ and δ, the probability cutoffs c_T and c_E, and a sample size cutoff N^*. It is worth noting the specification of design parameters is relatively simplistic as the last five parameters ϵ, δ, c_T, c_E, and N^* can take prespecified default values.

In general, the target toxicity rate ϕ and the clinically uninteresting efficacy rate ψ can be specified by the clinicians. Based on extensive simulation studies, we recommend $\epsilon, \delta \in [0.05, 0.15]$, as it generally leads to desirable performance of the proposed method when the maximum sample size is less than 100. As a default, we set $\epsilon = \delta = 0.1$ for general use. The specification of c_T and c_E depends on the values of ϕ and ψ as well as the maximum sample size. In this article, we take the probability cutoffs c_T and c_E to be 0.95 and 0.90, respectively. The parameter N^* controls the aggressiveness of the design, the smaller the values of N^*, the more conservative the dose escalation. We set $N^* = 9$ as the default value, while our sensitivity analysis shows that $N^* = 6$ and $N^* = 12$ also yield similar operating characteristics for the proposed design.

3.2 | Decision table

The dose-finding rule of the proposed uTPI design depends on the locations of the strongest toxicity and desirability intervals, (k_j^T, k_j^U), which are derived based on the assumption of the Beta-Binomial model. Here, we consider the simplified dose desirability formulation (2), where the desirability EU_j depends only on the efficacy and toxicity probabilities, and correspondingly in the decision table, the number of efficacies $n_{j,E}$ and number of toxicities $n_{j,T}$. Given a maximum sample size, it is feasible to enumerate all possible combinations of $D_j = (n_j, n_{j,T}, n_{j,E})$, then calculate and sort their corresponding (k_j^T, k_j^U). For example, when the patients receive doses in a cohort of size 3, n_j can take the value of 0, 3, 6, and so on. For each possible value of n_j (eg, 3), we may enumerate the combination of $n_{j,T}$ and $n_{j,E}$, each ranging from 0 to n_j, as shown in Table 1. Under each D_j, we calculate the corresponding strongest toxicity and desirability intervals (k_j^T, k_j^U). As a more straightforward representation of the desirability interval, we can order all possible values of k_j^U from the smallest to the largest, and assign a desirability score (DS) equal to the rank of the desirability interval k_j^U divided by the total number of possible values of k_j^U. Thus, choosing a dose level with the highest DS is equivalent to choosing a combination with the largest value of k_j^U.

A decision table of the uTPI design can be pre-calculated, as shown in Table 1. At each decision point, the user only needs to update k_j^T and DS of the current dose by looking the values up in the decision table, and then makes the corresponding dose escalation/de-escalation decision. As an illustration, suppose the target toxicity rate of a trial is 0.30, and suppose that the 2×2 table is given by $(w_1, w_2, w_3, w_4) = (0.7, 0, 1, 0.3)$. Suppose the three dose levels of a drug have observed data $D_1 = (3, 0, 0)$, $D_2 = (9, 2, 5)$, and $D_3 = (3, 2, 1)$. Assuming that the current dose level is 2, according to Table 1, the corresponding strongest toxicity interval is 2, which is less than $k^* = 4$. As a result, the admissible set is $\{1, 2, 3\}$. By looking up Table 1 again, the corresponding DS's for the three doses are 12, 42, and 36, respectively. Since dose level 2 yields the largest DS, the current dose is retained at the same level.

Table 1 is generated based on the cohort size of 3 and $w_1 + w_4 = 1$. Due to the flexible nature of the uTPI design, similar decision tables with different cohort sizes and utility preference specifications can be generated accordingly. In addition, it is also feasible to compute the decision table based on (x_1, x_2, x_3, x_4) when $w_1 + w_4 \neq 1$. A decision table for a trial with a target toxicity rate $\phi = 0.30$, $\epsilon = \delta = 0.1$, and a preference structure specification $(w_1, w_2, w_3, w_4) = (0.4, 0, 1, 0.55)$ can be found in Table S.5 of the Supplementary Material.

3.3 | Trial illustration

We provide a trial illustration of the proposed uTPI design under the setting of the motivating trial.[5] Four dose levels of the conformational B-cell epitope vaccines are considered: 1.0, 1.5, 2.0, and 2.5 mg. In the actual trial, no toxicity was observed, and dose level 2 has the highest number of efficacy responses. The toxicity and efficacy outcomes are simulated in a way that the underlying efficacy and toxicity rates reflect the actual rates in the real trial. For simulating toxicity outcomes, as no DLT was observed at each dose level in the actual trial, we use a generic value of 0.1 as the toxicity probability for all four doses. For simulating efficacy results, we first use the log standardized dose level as the covariate, and fit a quadratic logistic regression model to the efficacy data in the real trial. The fitted probability is 0.01, 0.61, 0.57, 0.14 for dose level 1, 2, 3, 4, respectively; the dose level 2 has the highest efficacy, which is consistent with the results in the actual trial.

We assume that a total of 24 patients are enrolled and treated in cohorts of size three. Given the target toxicity rate of 0.3 and $(w_1, w_4) = (0.7, 0.3)$, the decision table of the uTPI design based on the default setting is provided in Table 1. As shown in Figure 2, the first cohort does not have any toxicity or efficacy, so the second cohort is escalated to dose level 2. Two efficacy responses and no toxicity are observed in cohort 2. According to Table 1, the highest DS is 56 at dose level 2 and thus the next dose is still dose level 2. With one DLT and no efficacy observed in cohort 3, the estimated toxicity interval is 2 for dose level 2, which is below the MTD interval of 4, and thus the admissible set of dose levels is $\{1, 2, 3\}$. Using Table 1 again, the DS's of these three admissible dose levels are (12, 30.5, 40). As dose level 3 yields the highest DS, the next cohort is escalated to dose level 3, where one efficacy and one toxicity occur in this cohort. The next dose set is updated to $\{2, 3, 4\}$ with dose level 4 having the highest DS, and the patients fifth cohort are allocated to level 4. In cohort 5, none of the patients have experienced efficacy or toxicity and the uTPI design identifies dose level 3 as the one with the highest DS, leading to the de-escalation to dose level 3. The remaining dose-assignment decisions of the trial based on the uTPI design are given in Figure 2. At the end of the trial, the desirability of each dose is estimated, based on the approach

TABLE 1 Toxicity interval and desirability table for the uTPI design with the target toxicity rate $\phi = 0.30$, $\epsilon = \delta = 0.1$, and the utility preference structure specification $w_1 = 0.70$, $w_2 = 0$, $w_3 = 1$, and $w_4 = 0.30$

Num. Patients	Num. Tox	Num. Eff	Toxicity Interval	Desirability Score	Num. Patients	Num. Tox	Num. Eff	Toxicity Interval	Desirability Score
0	0	0	0	40	9	0	4	1	41
3	0	0	1	12	9	0	5	1	50
3	0	1	1	36	9	0	6	1	62
3	0	2	1	56	9	0	7	1	67
3	0	3	1	81	9	0	8	1	81
3	1	0	4	12	9	0	9	1	81
3	1	1	4	36	9	1	0	2	E
3	1	2	4	56	9	1	1	2	9
3	1	3	4	81	9	1	2	2	17
3	2	0	7	12	9	1	3	2	26
3	2	1	7	36	9	1	4	2	38
3	2	2	7	56	9	1	5	2	48
3	2	3	7	81	9	1	6	2	53
3	3	≥0	10	E	9	1	7	2	65
6	0	0	1	6.5	9	1	8	2	74
6	0	1	1	20.5	9	1	9	2	81
6	0	2	1	30.5	9	2	0	3	E
6	0	3	1	45.5	9	2	1	3	3
6	0	4	1	59.5	9	2	2	3	15
6	0	5	1	71.5	9	2	3	3	24
6	0	6	1	81	9	2	4	3	33
6	1	0	2	6.5	9	2	5	3	42
6	1	1	2	20.5	9	2	6	3	51
6	1	2	2	30.5	9	2	7	3	63
6	1	3	2	45.5	9	2	8	3	68
6	1	4	2	59.5	9	2	9	3	81
6	1	5	2	71.5	9	3	0	4	E
6	1	6	2	81	9	3	1	4	2
6	2	0	4	6.5	9	3	2	4	10
6	2	1	4	20.5	9	3	3	4	18
6	2	2	4	30.5	9	3	4	4	27
6	2	3	4	45.5	9	3	5	4	39
6	2	4	4	59.5	9	3	6	4	49
6	2	5	4	71.5	9	3	7	4	54
6	2	6	4	81	9	3	8	4	66
6	3	0	6	6.5	9	3	9	4	75

(Continues)

TABLE 1 (Continued)

Num. Patients	Num. Tox	Num. Eff	Toxicity Interval	Desirability Score	Num. Patients	Num. Tox	Num. Eff	Toxicity Interval	Desirability Score
6	3	1	6	20.5	9	4	0	5	E
6	3	2	6	30.5	9	4	1	5	1
6	3	3	6	45.5	9	4	2	5	4
6	3	4	6	59.5	9	4	3	5	16
6	3	5	6	71.5	9	4	4	5	25
6	3	6	6	81	9	4	5	5	34
6	≥4	≥0	7	E	9	4	6	5	43
9	0	0	1	E	9	4	7	5	52
9	0	1	1	14	9	4	8	5	64
9	0	2	1	23	9	4	9	5	69
9	0	3	1	28	9	≥5	≥0	6	E

Note: "E" stands for elimination due to toxicity or futility. The toxicity interval is the strongest toxicity interval. Based on the target toxicity rate of 0.30, the location of the highest tolerable toxicity interval $k^* = 4$.

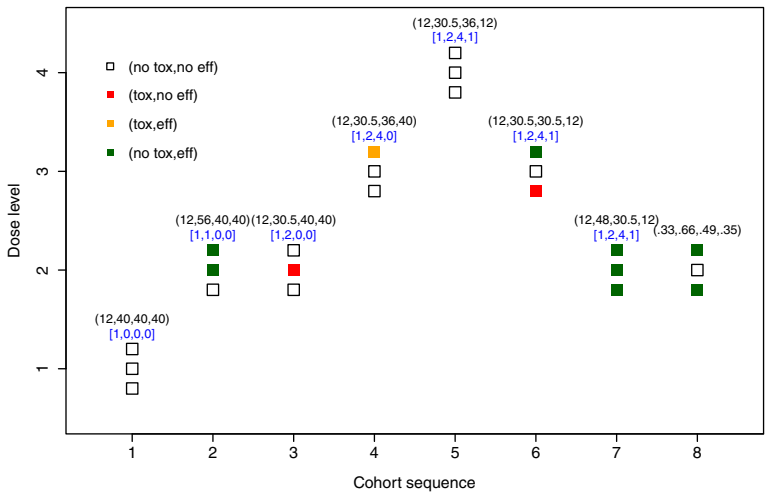

FIGURE 2 Trial illustration of the uTPI design under the setting of the HER-2 B-cell peptide vaccine trial.[5] The patients are enrolled in cohort size of three, and each square denotes a patient. For cohort 1-7, the numbers in parentheses correspond to the desirability scores of the four dose levels derived using Table 1, and the numbers in brackets correspond to the strongest toxicity intervals. For the last cohort, the vector in parentheses is the estimated desirability vector [Colour figure can be viewed at wileyonlinelibrary.com]

of Lin and Yin,[23] so that information can be borrowed across all the dose levels. Finally, dose level 2 is chosen as the OBD, as it possesses the highest desirability value with an acceptable toxicity rate below 0.3.

4 | NUMERICAL STUDIES

4.1 | Fixed scenarios

We compare the operating characteristics of the proposed uTPI design with those of other commonly used phase I/II dose-finding methods, including the TEPI design,[22] the Iso design,[17] and the EffTox design.[15] In particular, the EffTox and Iso designs are curve-based, while the uTPI and TEPI designs are curve-free. In the literature, there are also other designs proposed for phase I/II trials. We choose the aforementioned three designs here because all of them have available software to implement, and they also can be straightforwardly adapted to the utility-based setting. In the simulation setup, the target toxicity probability and the target efficacy probability are $\phi = 0.30$ and $\psi = 0.25$, respectively. A total of 36 subjects are sequentially enrolled into the trial in a cohort size of three, and allocated to one of the five dose levels. The design parameters and also the final OBD estimation procedures for each design are calibrated so that all considered approaches

have the same objectives and thus are comparable. In particular, we assume that the dose desirability formulation follows (1) and consider two sets of utility preference structure, (w_1, w_4)= (0.7, 0.3) and (0.4, 0.55), representing the cases where the efficacy is valued higher and lower than non-toxicity, respectively. In terms of the design parameters, the default values are chosen for the uTPI and TEPI designs. The simulation results of the Iso design are calculated by adapting the R code available from https://odin.mdacc.tmc.edu/~yyuan/Software/TargetAgent/targetAgentDF.r. We set the threshold for the posterior probability of toxicity to be 0.80, above which the dose would be deemed as inadmissible. The simulation results of the EffTox design are obtained based on the software available from https://biostatistics.mdanderson.org/SoftwareDownload. The dose desirability function of the EffTox design has been calibrated to be close to that of our design for the purpose of a fair comparison. In particular, the EffTox design requires specifying three points on the desirability contour (ie, having the same desirability values). These three points are calibrated such that they yield identical desirabilities in terms of (1). The implementation details of the EffTox design can be found in Section S.5 of the Supplementary Material. Similar dose-elimination rules for toxicity and efficacy are imposed for all four considered designs, with the probability cutoffs at $c_T = 0.95$ and $c_E = 0.90$. Furthermore, the OBD selection procedure (4) is imposed for the uTPI, TEPI, and Iso designs.

A total of 10 scenarios, each with different dose-response curves for toxicity/efficacy, are specified and simulated. In particular, scenario 10 has no OBD, as all the doses are either overly toxic or unexpectedly inefficacious. In this simulation, the efficacy and toxicity outcomes of each patient are simulated independently, while the performance of uTPI with correlated outcomes is examined in Section 4.3. The percentage of selection and the number of patients allocated to each dose are shown in Table 2 (for $(w_1, w_4) = (0.7, 0.3)$) and Table 3 (for $(w_1, w_4) = (0.4, 0.55)$). In addition, the average numbers of DLTs and efficacy outcomes, as well as the percentage of trials that are stopped early are reported. All the metrics are averaged over 10 000 trial replications.

In the first simulation where $(w_1, w_4) = (0.7, 0.3)$, as shown in Table 2, although in a few scenarios (eg, 1 and 9), the uTPI design is out-performed by Iso and EffTox, overall, it has reasonably robust and satisfactory performances. A notable comparison between TEPI and uTPI reveals that uTPI has higher OBD selection percentages and patient allocations than TEPI in most of the scenarios. It is worth noting that patient allocations in TEPI might be suboptimal in Scenarios 3 to 6 and 8, where the OBD is different from the MTD (ie, the dose level that has the target toxicity rate closest to the target ϕ). In these scenarios, we observe that the number of patients treated at the OBD in the uTPI design is more than double that under the TEPI design, and we also find that the number of patients treated at the higher doses is relatively larger under the TEPI design. A possible explanation for such a phenomenon is that the dose-finding algorithm of TEPI makes a dose-escalation decision as long as the current dose level has a small or moderate toxicity rate, regardless of the efficacy outcomes. Such an escalation decision is based on the implicit assumption that higher dose levels possess higher efficacy, which might not necessarily be true.

The EffTox design outperforms the uTPI design in terms of the correct selection percentage in Scenario 9. However, in other scenarios it tends to exhibit a larger variability in performances, and generally performs not as well as the uTPI design. For example, low correct dose selection percentages in the EffTox design are observed in Scenarios 4 to 6. This indicates that the curve-based EffTox design might be sensitive to the various dose-toxicity and dose-efficacy relationships. In particular, when all dose levels are tolerable, and the tumor response probability increases at low doses and then plateaus or decreases at higher doses, such as in Scenarios 4 and 5, the EffTox design treats most of the patients at the highest dose level. The Iso design appears to be more robust than the EffTox design; its performance is similar or superior to the uTPI design in Scenarios 1 and 5, and inferior to the uTPI design in other scenarios.

In the last scenario, where no dose is admissible, it is of interest to assess the early stopping percentage, as it is desirable to terminate the trial as early as possible. We observe that the uTPI design achieves a similar early stopping percentage to the other designs, indicating that it is competitively safe.

One major question in dose-finding problems is how to balance the tradeoff between exploration and exploitation. For untried dose level, we assign an initial value for k_j^U as $k_j^U = (0.5w_1 + w_4)K_U$ in the simulation study. In fact, this initial specification corresponds to a dose level with a toxicity probability of 0 and an efficacy probability of 0.5. Thus, the expected utility of untried dose levels is centered around $(0.5w_1 + w_4)$, which is actually very high. If k_{j-1}^U is greater than the initial high desirability value of an untried dose k_j^U, suggesting that the current dose level $j-1$ has higher desirability, the untried dose level j will not be explored. As an example, in Scenario 4, where the optimal dose level has low toxicity rates and very high efficacy rates that result in high desirability, as long as the dose assignment hits this dose level, the uTPI design can quickly identify the great benefit of assigning patients to these levels, rather than continuing to escalate to the higher levels. Therefore, it will favor exploitation and thus tend to make the dose retainment decisions instead of

TABLE 2 Comparison of the selection percentages (with the numbers of patients treated at each dose in parenthesis) of the uTPI, TEPI, Iso, and EffTox designs under 10 fixed scenarios, where the utility preference structure follows $(w_1, w_2, w_3, w_4) = (0.7, 0, 1, 0.3)$; the optimal biological dose with the highest desirability is denoted in bold

Design	Dose level 1	2	3	4	5	Number of DLTs/efficacy	% early stopping
	Scenario 1						
(π_T, π_E, EU)	**(0.20,0.40,0.52)**	(0.40,0.50,0.53)	(0.45,0.60,0.59)	(0.50,0.70,0.64)	(0.55,0.80,0.70)		
uTPI	**69.9 (19.6)**	21.5 (11.0)	4.1 (3.7)	0.5 (0.7)	0.0 (0.1)	10.4/16.1	3.9
TEPI	**66.8 (15.6)**	21.8 (13.8)	5.9 (4.4)	1.2 (0.9)	0.1 (0.1)	11.1/16.6	4.2
Iso	**75.2 (18.5)**	17.5 (9.6)	3.3 (4.6)	0.6 (1.9)	0.0 (0.7)	10.9/16.8	3.4
EffTox	**50.0 (18.2)**	24.0 (7.5)	13.0 (4.8)	5.0 (2.1)	4.0 (2.2)	11.1/17.1	4.0
	Scenario 2						
(π_T, π_E, EU)	(0.15,0.40,0.54)	**(0.30,0.60,0.63)**	(0.45,0.60,0.59)	(0.55,0.60,0.56)	(0.65,0.60,0.53)		
uTPI	39.7 (14.1)	**54.4 (17.3)**	4.3 (3.6)	0.3 (0.5)	0.0 (0.1)	9.3/18.6	1.3
TEPI	41.0 (9.4)	**49.8 (17.0)**	7.2 (7.7)	0.6 (1.4)	0.0 (0.1)	10.8/19.5	1.4
Iso	52.2 (13.6)	**42.8 (14.7)**	3.6 (5.4)	0.1 (1.6)	0.0 (0.4)	10.0/18.7	1.3
EffTox	43.0 (16.5)	**36.0 (9.6)**	16.0 (6.4)	3.0 (1.9)	1.0 (1.1)	10.0/18.0	2.0
	Scenario 3						
(π_T, π_E, EU)	(0.15,0.25,0.43)	**(0.20,0.55,0.63)**	(0.25,0.40,0.51)	(0.35,0.30,0.41)	(0.45,0.20,0.31)		
uTPI	11.4 (6.7)	**67.7 (18.2)**	15.7 (6.0)	3.1 (3.3)	0.2 (1.4)	7.9/15.4	1.9
TEPI	14.3 (6.4)	**61.6 (9.4)**	18.5 (10.0)	3.1 (7.1)	0.1 (2.5)	9.0/13.4	2.3
Iso	17.9 (7.4)	**65.5 (18.3)**	13.3 (7.3)	1.4 (2.2)	0.0 (0.5)	7.5/15.6	1.9
EffTox	42.0 (14.5)	**24.0 (7.4)**	18.0 (6.8)	5.0 (3.2)	7.0 (3.3)	8.0/12.0	4.0
	Scenario 4						
(π_T, π_E, EU)	(0.10,0.05,0.31)	(0.12,0.30,0.47)	**(0.15,0.60,0.68)**	(0.20,0.60,0.66)	(0.25,0.60,0.65)		
uTPI	0.3 (3.4)	4.5 (5.7)	**56.2 (17.0)**	27.6 (6.9)	10.1 (2.6)	5.6/17.8	1.3
TEPI	0.2 (4.7)	7.2 (5.0)	**37.9 (6.6)**	36.4 (7.2)	14.8 (11.6)	6.4/16.9	3.5
Iso	1.0 (4.1)	11.9 (6.7)	**49.4 (12.2)**	26.8 (8.2)	9.7 (4.5)	5.8/17.2	1.3
EffTox	0.0 (3.5)	1.0 (3.3)	**6.0 (4.5)**	10.0 (5.4)	82.0 (19.1)	7.3/18.6	1.0
	Scenario 5						
(π_T, π_E, EU)	(0.01,0.05,0.33)	(0.02,0.10,0.36)	**(0.03,0.35,0.54)**	(0.04,0.20,0.43)	(0.05,0.15,0.39)		
uTPI	2.2 (3.2)	5.7 (3.8)	**62.8 (11.9)**	20.7 (8.7)	8.5 (8.4)	1.2/7.7	0.0
TEPI	1.8 (3.2)	6.5 (3.3)	**56.4 (3.8)**	29.0 (7.7)	6.4 (17.9)	1.4/6.0	0.0
Iso	0.5 (4.9)	5.6 (5.5)	**62.5 (15.3)**	23.4 (6.8)	7.9 (3.5)	1.1/8.0	0.2
EffTox	5.0 (5.1)	1.0 (3.1)	**4.0 (3.7)**	6.0 (4.2)	47.0 (15.3)	1.2/5.0	38.0
	Scenario 6						
(π_T, π_E, EU)	(0.08,0.10,0.35)	(0.10,0.20,0.41)	**(0.15,0.70,0.75)**	(0.32,0.70,0.69)	(0.40,0.75,0.71)		
uTPI	1.1 (3.6)	2.2 (4.3)	**82.5 (23.4)**	12.5 (4.0)	1.1 (0.6)	5.8/20.9	0.6
TEPI	0.7 (4.6)	2.9 (4.6)	**61.8 (9.8)**	27.4 (11.2)	5.5 (5.5)	8.0/20.1	1.6
Iso	3.7 (6.2)	8.7 (4.6)	**62.2 (12.7)**	20.7 (8.2)	3.3 (4.1)	7.1/19.2	1.3
EffTox	1.0 (4.4)	1.0 (3.3)	**13.0 (5.7)**	33.0 (9.1)	51.0 (13.2)	9.7/21.4	1.0

(Continues)

TABLE 2 (Continued)

Design	Dose level 1	2	3	4	5	Number of DLTs/efficacy	% early stopping
	Scenario 7						
(π_T, π_E, EU)	(0.05,0.05,0.32)	(0.09,0.10,0.34)	(0.10,0.50,0.62)	**(0.25,0.65,0.68)**	(0.45,0.80,0.73)		
uTPI	0.7 (3.3)	1.1 (3.5)	42.5 (14.5)	**50.2 (12.2)**	4.9 (2.4)	6.0/17.6	0.6
TEPI	0.2 (4.3)	1.0 (4.2)	36.1 (6.1)	**49.9 (13.0)**	11.3 (8.1)	8.1/18.6	1.5
Iso	1.1 (4.9)	3.6 (4.3)	44.0 (9.7)	**44.3 (10.8)**	5.8 (6.1)	7.0/17.5	1.2
EffTox	0.0 (3.6)	0.0 (3.1)	4.0 (4.1)	**45.0 (9.7)**	50.0 (15.3)	10.2/21.1	0.0
	Scenario 8						
(π_T, π_E, EU)	(0.01,0.10,0.37)	(0.05,0.15,0.39)	(0.10,0.25,0.45)	**(0.12,0.50,0.61)**	(0.27,0.50,0.57)		
uTPI	0.8 (3.4)	2.6 (3.9)	7.6 (5.2)	**64.1 (15.3)**	24.7 (8.2)	4.8/14.0	0.2
TEPI	1.2 (3.5)	3.5 (4.1)	11.4 (4.6)	**60.5 (6.7)**	23.2 (17.2)	6.1/14.0	0.2
Iso	4.9 (7.3)	17.6 (7.9)	15.8 (5.8)	**47.5 (8.9)**	13.6 (6.1)	3.8/10.9	0.7
EffTox	3.0 (5.5)	1.0 (3.4)	6.0 (4.4)	**10.0 (4.6)**	79.0 (17.8)	6.0/13.4	1.0
	Scenario 9						
(π_T, π_E, EU)	(0.05,0.03,0.31)	(0.10,0.05,0.31)	(0.15,0.15,0.36)	(0.19,0.30,0.45)	**(0.28,0.55,0.60)**		
uTPI	2.7 (3.9)	3.7 (4.2)	13.0 (5.5)	30.5 (8.9)	**47.8 (13.3)**	6.8/11.1	2.3
TEPI	2.9 (4.8)	3.1 (5.6)	12.3 (6.3)	32.8 (6.9)	**44.7 (11.9)**	6.4/10.0	4.2
Iso	2.7 (4.4)	5.2 (5.2)	19.7 (7.2)	32.5 (8.1)	**35.4 (10.8)**	6.3/9.8	4.5
EffTox	0.0 (3.6)	0.0 (3.1)	1.0 (3.6)	6.0 (4.1)	**87.0 (20.2)**	7.5/13.1	6.0
	Scenario 10						
(π_T, π_E, EU)	(0.15,0.01,0.26)	(0.30,0.02,0.22)	(0.45,0.03,0.19)	(0.55,0.04,0.16)	(0.60,0.05,0.16)		
uTPI	8.7 (9.3)	5.4 (8.8)	4.0 (5.9)	1.0 (2.6)	0.1 (0.9)	8.7/0.6	80.8
TEPI	8.3 (9.2)	1.9 (8.5)	1.5 (4.6)	0.5 (1.3)	0.0 (0.1)	6.7/0.5	87.8
Iso	18.8 (6.5)	4.2 (7.8)	0.9 (6.5)	0.1 (3.9)	0.0 (2.0)	9.6/0.7	75.9
EffTox	0.0 (3.3)	0.0 (3.0)	0.0 (3.0)	2.0 (2.1)	1.0 (1.5)	4.8/0.3	96.0

Note: π_T denotes the marginal toxicity rate, π_E denotes and marginal efficacy rate, and EU represents the dose desirability.

exploring more dose levels. As a result, compared with other designs, uTPI design allocates fewer patients on average to the more toxic dose levels in this scenario.

In the second simulation, the utility preference structure of the desirability function shifts to $(w_1, w_4) = (0.4, 0.55)$, as shown in Table 3. For the purpose of a fair comparison of the correct selection percentage, we adjust and calibrate the other three designs such that they also have the same desirability functions as our design in the final selection of the optimal dose. As a result, a change in the preference structure (w_1, w_4) would result in different correct selection percentages of the three other designs. On the other hand, the patient allocation distributions of Iso and TEPI remain the same across different preference structure specifications (by comparing Tables 2 and 3). This is due to the fact that the patient allocation rules of these two designs do not depend on the dose desirability, thus a shift in utility preference structure would lead to identical patient allocation. In contrast, both the EffTox and uTPI designs are more flexible and can incorporate different preference structures for decision making, and thus different patient allocation values are observed. As in the first simulation, we observe a similar trend with regards to the TEPI and EffTox design, that is, the uTPI design outperforms the TEPI design in most scenarios and large variation is seen in the performance of the EffTox design.

The framework of the uTPI design can be generalized to incorporate other forms of dose desirability calculation. It is feasible to adapt our design by replacing the formula of EU based on a 2×2 utility preference table with a customized desirability function. Additional results on the adapted version of uTPI with the dose desirability

TABLE 3 Comparison of the selection percentages (with the numbers of patients treated at each dose in parenthesis) of the uTPI, TEPI, Iso, and EffTox designs under 10 fixed scenarios, where the utility preference structure follows $(w_1, w_2, w_3, w_4) = (0.4, 0, 1, 0.55)$; the optimal biological dose with the highest desirability is denoted in bold

Design	Dose level 1	2	3	4	5	Number of DLTs/efficacy	% early stopping
	Scenario 1						
(π_T, π_E, EU)	**(0.20,0.40,0.60)**	(0.40,0.50,0.53)	(0.45,0.60,0.54)	(0.50,0.70,0.56)	(0.55,0.80,0.57)		
uTPI	**79.1 (20.9)**	13.1 (9.6)	2.9 (3.5)	0.5 (0.8)	0.1 (0.1)	10.1/16.0	4.3
TEPI	**73.8 (15.6)**	17.0 (13.8)	4.2 (4.4)	0.7 (0.9)	0.1 (0.1)	11.1/16.6	4.2
Iso	**80.9 (18.5)**	14.5 (9.6)	1.1 (4.6)	0.1 (1.9)	0.0 (0.7)	10.9/16.8	3.4
EffTox	**87.0 (30.4)**	5.0 (2.7)	3.0 (1.1)	1.0 (0.5)	1.0 (0.4)	8.1/14.8	4.0
	Scenario 2						
(π_T, π_E, EU)	**(0.15,0.40,0.63)**	(0.30,0.60,0.63)	(0.45,0.60,0.54)	(0.55,0.60,0.49)	(0.65,0.60,0.43)		
uTPI	**55.9 (16.6)**	38.7 (14.3)	3.7 (3.9)	0.2 (0.7)	0.0 (0.1)	9.0/18.1	1.5
TEPI	**47.7 (9.4)**	45.4 (17.0)	5.2 (7.7)	0.2 (1.4)	0.0 (0.1)	10.8/19.5	1.4
Iso	**81.0 (13.6)**	13.8 (14.7)	3.6 (5.4)	0.2 (1.6)	0.0 (0.4)	10.0/18.7	1.3
EffTox	**84.0 (29.1)**	12.0 (4.8)	2.0 (1.2)	0.0 (0.3)	0.0 (0.1)	6.6/15.5	1.0
	Scenario 3						
(π_T, π_E, EU)	(0.15,0.25,0.57)	**(0.20,0.55,0.66)**	(0.25,0.40,0.57)	(0.35,0.30,0.48)	(0.45,0.20,0.38)		
uTPI	18.5 (7.8)	**62.9 (16.1)**	14.5 (6.5)	2.0 (3.7)	0.2 (1.5)	8.0/14.8	1.9
TEPI	19.1 (6.4)	**60.0 (9.4)**	15.7 (10.0)	2.8 (7.1)	0.1 (2.5)	9.0/13.4	2.3
Iso	63.6 (7.4)	**13.1 (18.3)**	17.4 (7.3)	3.8 (2.2)	0.2 (0.5)	7.5/15.6	1.9
EffTox	60.0 (21.4)	**29.0 (9.1)**	5.0 (3.1)	2.0 (1.0)	2.0 (0.9)	6.6/12.1	2.0
	Scenario 4						
(π_T, π_E, EU)	(0.10,0.05,0.52)	(0.12,0.30,0.60)	**(0.15,0.60,0.71)**	(0.20,0.60,0.68)	(0.25,0.60,0.65)		
uTPI	0.9 (3.6)	9.3 (6.3)	**53.0 (16.0)**	25.5 (7.0)	10.1 (2.9)	5.6/17.6	1.3
TEPI	0.7 (4.7)	11.0 (5.0)	**41.4 (6.6)**	27.0 (7.2)	16.4 (11.6)	6.4/16.9	3.5
Iso	42.2 (4.1)	31.7 (6.7)	**11.5 (12.2)**	9.5 (8.2)	3.8 (4.5)	5.8/17.2	1.3
EffTox	1.0 (4.6)	17.0 (8.3)	**34.0 (10.6)**	15.0 (5.0)	33.0 (7.4)	5.9/16.5	1.0
	Scenario 5						
(π_T, π_E, EU)	(0.01,0.05,0.56)	(0.02,0.10,0.58)	**(0.03,0.35,0.67)**	(0.04,0.20,0.61)	(0.05,0.15,0.58)		
uTPI	4.8 (3.2)	7.7 (3.7)	**58.3 (12.2)**	19.9 (8.9)	9.3 (8.1)	1.2/7.8	0.0
TEPI	6.9 (3.2)	10.9 (3.3)	**47.6 (3.8)**	22.3 (7.7)	12.2 (17.9)	1.4/6.0	0.0
Iso	51.4 (4.9)	16.2 (5.5)	**11.9 (15.3)**	16.9 (6.8)	3.4 (3.5)	1.1/8.0	0.2
EffTox	3.0 (4.6)	2.0 (4.1)	**18.0 (7.1)**	15.0 (6.5)	35.0 (10.2)	1.1/6.0	27.0
	Scenario 6						
(π_T, π_E, EU)	(0.08,0.10,0.55)	(0.10,0.20,0.58)	**(0.15,0.70,0.75)**	(0.32,0.70,0.65)	(0.40,0.75,0.63)		
uTPI	1.9 (3.7)	4.0 (4.7)	**80.9 (22.5)**	10.9 (4.2)	1.6 (0.7)	5.8/20.6	0.7
TEPI	1.8 (4.6)	4.2 (4.6)	**66.2 (9.8)**	21.6 (11.2)	4.6 (5.5)	8.0/20.1	1.6
Iso	48.9 (6.2)	37.7 (4.6)	**7.1 (12.7)**	4.3 (8.2)	0.7 (4.1)	7.1/19.2	1.3
EffTox	3.0 (6.2)	12.0 (6.8)	**48.0 (12.9)**	21.0 (6.1)	16.0 (3.8)	6.6/18.1	0.0

(Continues)

TABLE 3 (Continued)

Design	Dose level 1	2	3	4	5	Number of DLTs/efficacy	% early stopping
	Scenario 7						
(π_T, π_E, EU)	(0.05,0.05,0.54)	(0.09,0.10,0.54)	**(0.10,0.50,0.70)**	(0.25,0.65,0.67)	(0.45,0.80,0.62)		
uTPI	2.2 (3.4)	2.1 (3.7)	**52.8 (15.3)**	37.9 (11.0)	4.4 (2.6)	5.9/17.4	0.5
TEPI	1.2 (4.3)	2.2 (4.2)	**41.5 (6.1)**	44.7 (13.0)	8.9 (8.1)	8.1/18.6	1.5
Iso	49.7 (4.9)	23.3 (4.3)	**19.5 (9.7)**	5.8 (10.8)	0.6 (6.1)	7.0/17.5	1.2
EffTox	0.0 (4.1)	2.0 (4.7)	**38.0 (11.3)**	42.0 (10.6)	17.0 (5.2)	6.7/17.4	0.0
	Scenario 8						
(π_T, π_E, EU)	(0.01,0.10,0.58)	(0.05,0.15,0.58)	(0.10,0.25,0.60)	**(0.12,0.50,0.68)**	(0.27,0.50,0.60)		
uTPI	4.7 (3.5)	6.2 (4.1)	10.5 (5.6)	**61.1 (15.5)**	17.4 (7.3)	4.6/13.8	0.1
TEPI	6.5 (3.5)	8.2 (4.1)	13.7 (4.6)	**48.9 (6.7)**	22.5 (17.2)	6.1/14.0	0.2
Iso	45.2 (7.3)	19.9 (7.9)	22.8 (5.8)	**7.8 (8.9)**	3.6 (6.1)	3.8/10.9	0.7
EffTox	4.0 (6.0)	8.0 (6.3)	30.0 (9.4)	**30.0 (7.2)**	27.0 (6.9)	4.0/10.9	1.0
	Scenario 9						
(π_T, π_E, EU)	(0.05,0.03,0.53)	(0.10,0.05,0.52)	(0.15,0.15,0.53)	(0.19,0.30,0.57)	**(0.28,0.55,0.62)**		
uTPI	10.8 (4.1)	7.8 (4.6)	15.0 (6.2)	26.8 (9.3)	**37.0 (11.5)**	6.6/10.4	2.5
TEPI	9.9 (4.8)	6.5 (5.6)	14.6 (6.3)	25.9 (6.9)	**38.9 (11.9)**	6.4/10.0	4.2
Iso	54.4 (4.4)	9.5 (5.2)	15.5 (7.2)	13.3 (8.1)	**2.8 (10.8)**	6.3/9.8	4.5
EffTox	0.0 (3.7)	0.0 (3.7)	13.0 (7.1)	26.0 (7.8)	**53.0 (12.1)**	6.5/10.4	7.0
	Scenario 10						
(π_T, π_E, EU)	(0.15,0.01,0.47)	(0.30,0.02,0.39)	(0.45,0.03,0.31)	(0.55,0.04,0.26)	(0.60,0.05,0.24)		
uTPI	9.1 (9.3)	5.2 (8.8)	3.7 (5.8)	1.0 (2.6)	0.1 (0.9)	8.7/0.6	80.9
TEPI	8.9 (9.2)	1.6 (8.5)	1.3 (4.6)	0.5 (1.3)	0.0 (0.1)	6.7/0.5	87.8
Iso	21.9 (6.5)	1.6 (7.8)	0.5 (6.5)	0.1 (3.9)	0.0 (2.0)	9.6/0.7	75.9
EffTox	0.0 (3.2)	0.0 (3.2)	1.0 (3.2)	2.0 (2.2)	1.0 (1.3)	4.9/0.3	95.0

Note: π_T denotes the marginal toxicity rate, π_E denotes and marginal efficacy rate, and EU represents the dose desirability.

formulation being the same as the desirability function of the EffTox or TEPI design can be found in Section S.4 of the Supplementary Material.

4.2 | Random scenarios

To avoid selection bias in the assessment under "cherry-picked" fixed scenarios, we adopt a more objective approach that evaluates the operating characteristics based on a larger number of randomly generated scenarios. In particular, the probability vectors of toxicity and efficacy are generated via a variant of the pseudo-uniform approach, which is based on the ordered statistics of "quasi-uniform" random variables. The detailed algorithm for generating the random scenarios is provided in Appendix; a total of 10 000 scenarios are randomly generated with different toxicity and efficacy dose-response curves and various OBD locations. We have carefully tuned the parameters for generating the random scenarios to make sure that most of the generated scenarios have realistic toxicity and efficacy curves that may be encountered in real applications. For visualization, Figure 3 displays 20 scenarios from the 10 000 generated random scenarios.

In this simulation with random scenarios, the parameters and design configurations are the same as those in the fixed scenarios. We consider two sets of preference structure, $(w_1, w_4) = (0.7, 0.3)$ and $(0.4, 0.55)$, respectively. There exists

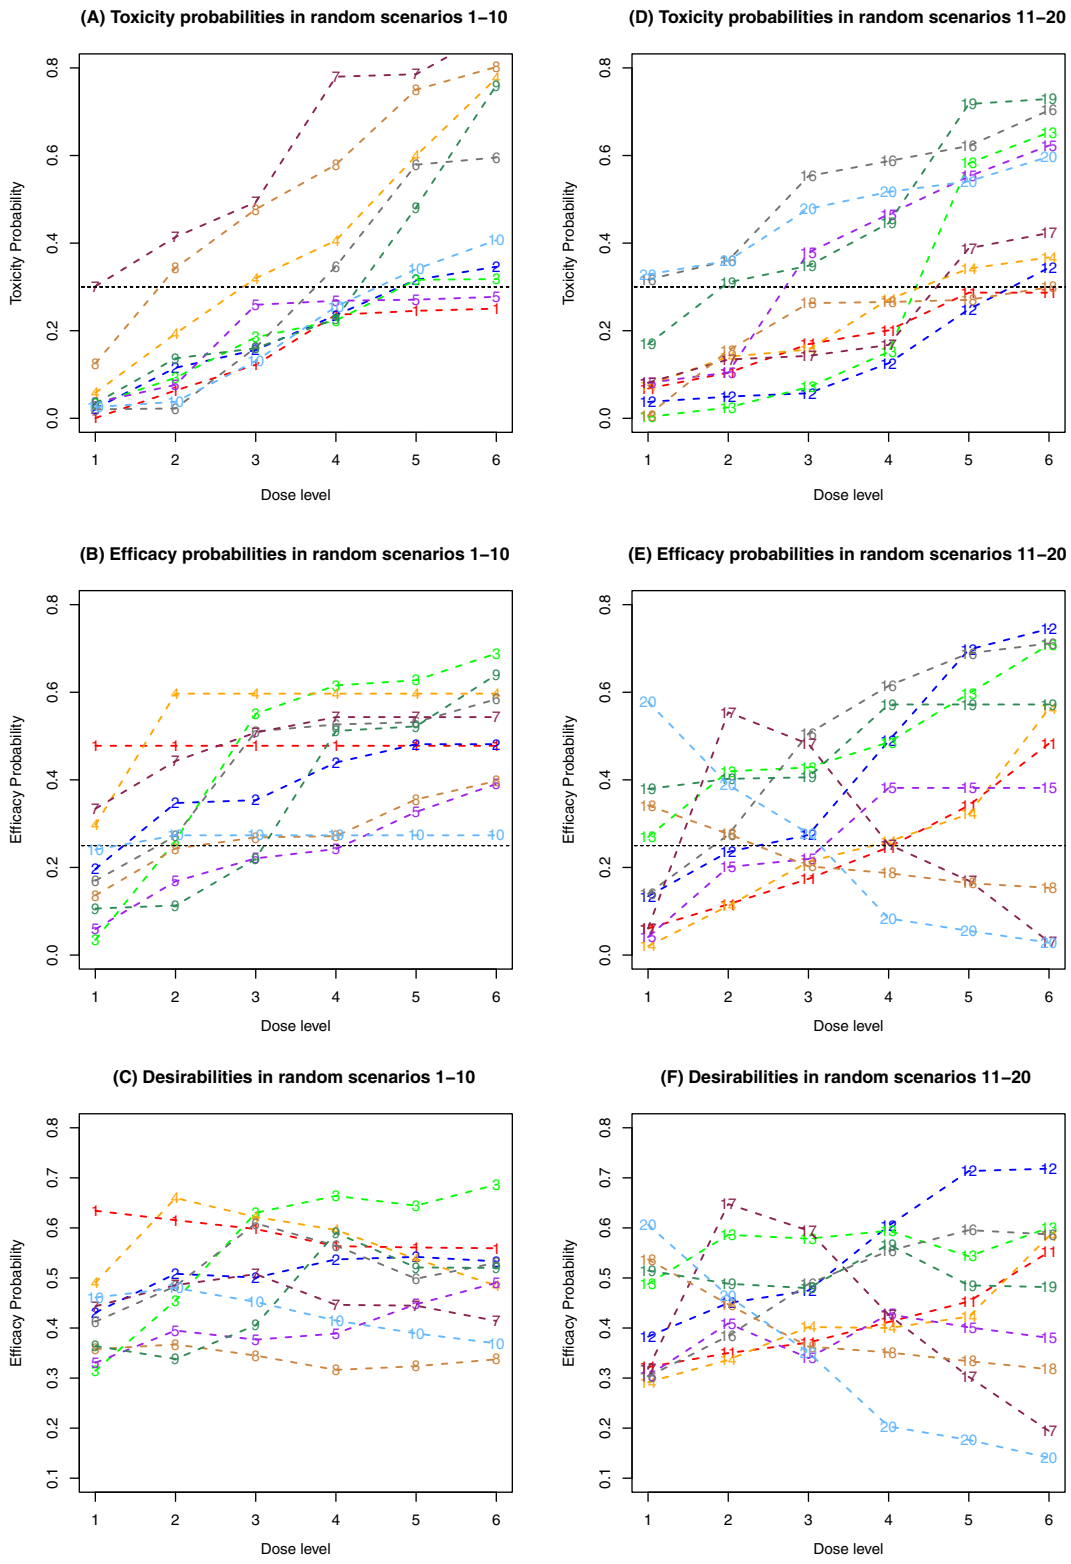

FIGURE 3 Twenty scenarios randomly selected from the 10 000 random scenarios used in the simulation study. The horizontal dashed lines correspond to the toxicity and efficacy targets $\phi = 0.30$ and $\psi = 0.25$, respectively [Colour figure can be viewed at wileyonlinelibrary.com]

no significant difference between the conclusions under the two preference structures, and thus only the results for $(w_1, w_4) = (0.7, 0.3)$ are shown. We examine the trend of various performance metrics with respect to increasing sample size N (from 30 to 120), averaged over all the 10 000 random scenarios. The performance metrics under consideration are: the selection percentage of the best dose (the tolerable dose with the maximum dose desirability), the selection percentage of favorable doses (the tolerable doses with desirability values around $\pm$ 5% of the highest desirability), percentage of patients allocated at the best dose, percentage of patients allocated at favorable doses, the poor allocation percentage (the proportion of trials that allocates less than $N/5$ patients to the best dose), and the number of patients allocated to overly toxic dose levels (the dose levels with toxicity probabilities larger than $\phi + \epsilon$). The relationships between the six performance metrics and increasing sample size are exhibited in Figure 4. It is evident that the uTPI design achieves desirable performance across all metrics. Among the four designs, uTPI has the highest selection and allocation percentages at the best/favorable doses, and the lowest percentages of poor allocation and overdose cases.

4.3 | Sensitivity analysis

The uTPI design has several default parameters, the sample size cutoff for dose expansion during the early stage N^*, as well as the widths of toxicity and desirability intervals ϵ and δ. We conducted a sensitivity analysis to see the performance of uTPI when the sample size cutoff N^* and the width δ take different values. In the first sensitivity analysis, we consider three values for N^*, that is, $N^* = 6$, 9 (default) , and 12, respectively, representing conservative, moderate, and aggressive exploration at the beginning of the trial. We assess the best dose's selection percentage and the number of patients allocated at the best dose across the 10 fixed scenarios considered in Section 4.1. The simulation results based on 10 000 replicated trials are provided in the upper panel of Figure 5. We observe that the performance of uTPI is quite robust to different values of N^*, especially in terms of the patient allocation distribution. With respect to the selection percentage of best dose, when the OBD locates in lower dose levels, the correct selection percentage is a little bit higher with $N^* = 6$; likewise, when the OBD locates in higher doses, $N^* = 9$ leads to higher correct selection. In the second sensitivity assessment, we examine the robustness of uTPI with respect to δ, the width of desirability intervals. We examine three cases: $\delta = 0.05$, 0.10 (default), and 0.15, which correspond to narrow, moderate, and wide desirability intervals, respectively. As shown in the middle panel of Figure 5, it is evident that the performance of the uTPI design is quite robust against the interval width.

Moreover, we further examine the performance of uTPI based on different correlations between the toxicity and efficacy outcomes. To do so, for each patient, we first simulate a bivariate normal random variable (z_1, z_2) based on zero mean vector and covariance matrix $\Sigma = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$, where ρ is the correlation parameter. Suppose the dose level for the patient is d, and denote $(\pi_{j,T}, \pi_{j,E})$ as the true toxicity and efficacy probability at dose level d. Based on this latent vector (z_1, z_2), if $z_1 < \Phi^{-1}(\pi_{j,T})$ $(z_2 < \Phi^{-1}(\pi_{j,E}))$, then toxicity (efficacy) is observed for this patient, where $\Phi^{-1}(\cdot)$ is the inverse cumulative distribution function of the standard normal random variable. In the simulation, we consider three different values for ρ, respectively indicating positive ($\rho = 0.5$), independent ($\rho = 0$), and negative ($\rho = -0.5$) correlations between toxicity and efficacy. We report the simulation results in the lower panel of Figure 5. It is still observed that the patient allocation of the uTPI design is robust to the correlation parameter ρ. However, since we assume that $w_1 > w_4$, the correct selection percentage under positive correlation between toxicity and efficacy is slightly larger in some scenarios. This is because, based on $w_1 > w_4$, the scenario with $\rho > 0$ facilitates more aggressive exploration of the dose space. As a result, the correct dose has a higher likelihood to be selected. Nonetheless, the variations between the correct selection percentages of uTPI with different toxicity and efficacy correlations are bounded by 3.5 percent, which still indicates that the proposed uTPI is robust across different situations.

5 | INCORPORATION OF DELAYED RESPONSES

For illustrative purposes, the previous sections employ the assumption that the efficacy/toxicity outcomes of the enrolled patients can be observed before the arrival of the new patient cohort. For certain endpoints in immunotherapy or targeted agent trials, this assumption may not hold true and the outcomes would be delayed. To deal with the case of such late-onset toxicity/efficacy, that is, a patient's outcome can be observed many weeks after receiving the treatment, we extend the uTPI design by treating the unobserved responses as missing data. As shown in the Supplementary

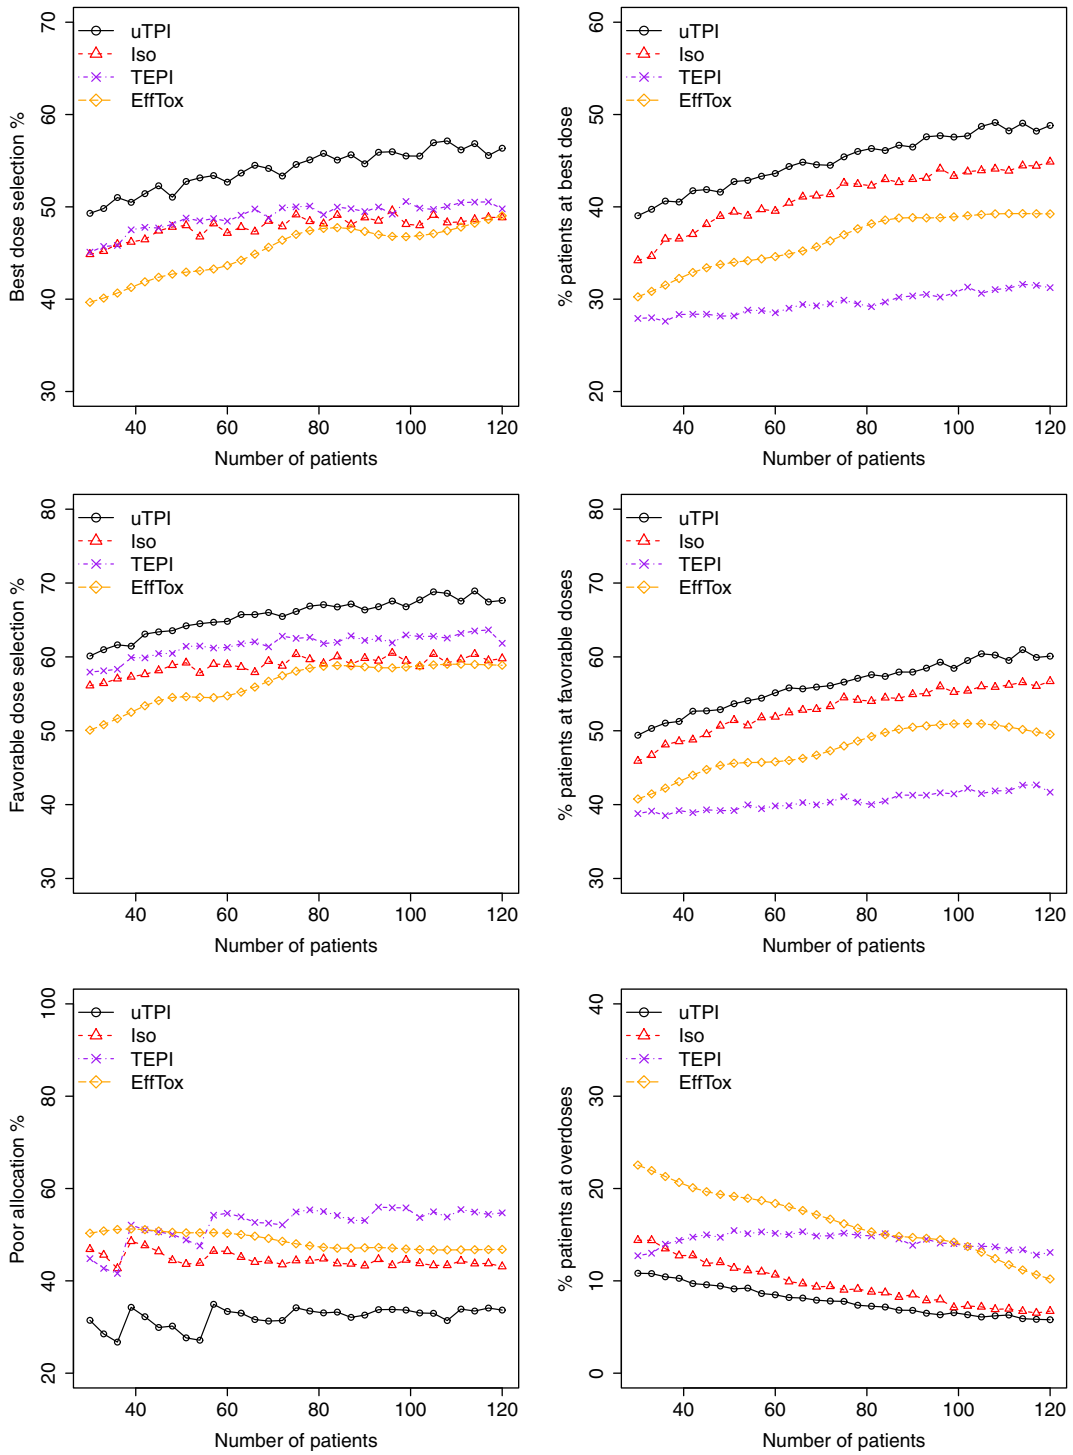

FIGURE 4 Comparison of the six performance metrics of uTPI, TEPI, Iso, and EffTox designs averaged over 10 000 randomly generated toxicity/efficacy scenarios with the utility preference structure specification $w_1 = 0.70$, $w_2 = 0$, $w_3 = 1$, and $w_4 = 0.30$ [Colour figure can be viewed at wileyonlinelibrary.com]

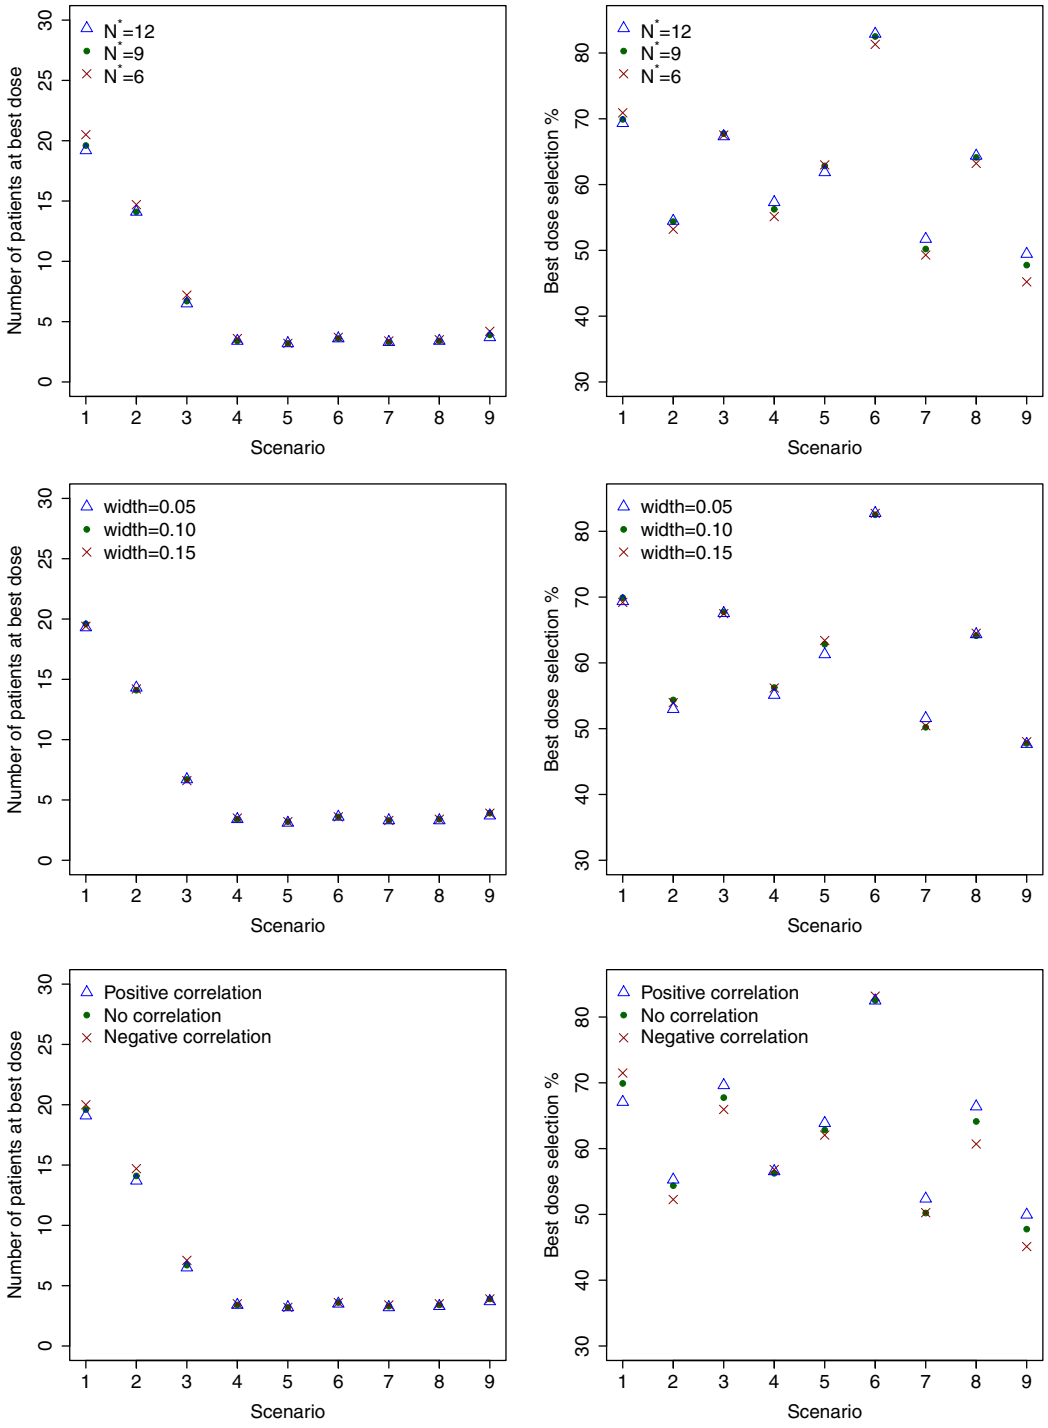

FIGURE 5 Sensitivity analysis of the selection percentage of the best dose and the number of patients allocated to the best dose [Colour figure can be viewed at wileyonlinelibrary.com]

Material, we adopt an approximated likelihood approach[31] and derive the posterior desirability distribution by incorporating the patient follow-up times. An additional simulation study reported in the Supplementary Material shows that the developed time-to-event uTPI (TITE-uTPI) design still can yield satisfactory operating characteristics across various scenarios. Compared to the standard uTPI design that requires repeated suspension of the patient accrual before assigning the next dose level for each new patient cohort, the TITE-uTPI design significantly shortens the trial duration due to the ability to make real-time decisions, and thus it is able to accelerate drug development.

6 | CONCLUSION

We have proposed an optimal biological dose-based toxicity probability interval design that is able to accurately identify the OBD, based on a one-dimensional, numerical utility value that summarizes the tradeoff between efficacy and toxicity outcomes. The proposed uTPI design accounts for both toxicity and efficacy outcomes, and uses quasi-binomial likelihood to simplify the modeling of desirability. The dose-assignment decisions of the uTPI design are adaptively made by maximizing the toxicity-efficacy tradeoffs. The design can be seen as a seamless improvement over the TEPI design, leading to a simpler, yet superior and more robust extension. Simulation studies have shown that the operating performance of the proposed design is both efficient and robust under various scenarios, and the application of our design to the immunotherapy trial has demonstrated the practicality of its use. We have also developed a TITE-uTPI design to address the issue of delayed responses. The `R` code to implement the uTPI design can be obtained from https://github.com/haoluns/uTPI. The web-based software for the uTPI design is currently under development, and it will be freely available soon from http://trialdesign.org.

The proposed uTPI has several important features in comparison with the existing phase I/II designs. For example, compared with the TEPI design, the proposed design uses an explicit utility function that accounts for the toxicity-efficacy tradeoffs by assigning different weighting to the efficacy/toxicity outcomes. Compared with the EffTox design, which employs a parametric dose-response model for inference, our design refrains from using a complicated function to model the dose desirability, and thus is more computationally appealing and easier to implement. Compared to the U-BOIN design, which uses a multinomial model for the toxicity and efficacy outcomes and the posterior mean of the dose desirability for making dose-assignment decisions, our design directly models the dose desirability and uses its full posterior distribution, instead of a single posterior mean estimate, for decision making.

In the proposed design, the toxicity probability is considered both in the modeling of the toxicity interval, and partially in the formulation of *EU*. Although a design purely based on the desirability interval is theoretically feasible, we incorporate the toxicity interval as an important dimension in the design because controlling the toxicity rate is a key motivation in early-phase trials, otherwise it is possible that too many patients might be allocated to an overly-toxic dose, which could lead to ethical concerns.

Upon a careful examination of the characteristics of the design, it is worth cautioning that there are a few possible dose-response cases where the design might not work optimally. For example, if the efficacy probabilities are all very high (eg, greater than 50%) at all the doses levels, it is possible that the design might be trapped at the initial doses and fail to escalate; adopting the optional dose escalation rule might help with avoiding such a scenario. Moreover, if the dose-desirability relationship follows an N-shaped curve and the OBD is located at the highest dose, our design might be trapped at a suboptimal dose level.

ACKNOWLEDGEMENTS

We thank the two referees, Associate Editor, and Editor for their many constructive and insightful comments that have led to significant improvements in the article. Lin's research was supported in part by grants P30 CA016672 and P50 CA221703 from the National Cancer Institute (NCI), and Yuan's research was supported in part by grants P50 CA098258 and P30 CA016672 from the NCI.

ORCID

Haolun Shi https://orcid.org/0000-0002-3620-9717
Jiguo Cao https://orcid.org/0000-0001-7417-6330
Ying Yuan https://orcid.org/0000-0003-3163-480X
Ruitao Lin https://orcid.org/0000-0003-2244-131X

REFERENCES

1. Cook N, Hansen AR, Siu LL, Abdul AR. Early phase clinical trials to identify optimal dosing and safety. *Mol Oncol*. 2015;9:997-1007.
2. Postel-Vinay S, Aspeslagh S, Lanoy E, Robert C, Soria JC, Marabelle A. Challenges of phase 1 clinical trials evaluating immune checkpoint-targeted antibodies. *Ann Oncol*. 2016;27:214-224.
3. Yuan Y, Nguyen HQ, Thall PF. *Bayesian Designs for Phase I-II Clin Trials*. Boca Raton, FL: Chapman & Hall/CRC Press; 2017.
4. Yan F, Thall PF, Lu KH, Gilbert MR, Yuan Y. Phase I–II clinical trial design: a state-of-the-art paradigm for dose finding. *Ann Oncol*. 2017;29:694-699.
5. Bekaii-Saab T, Wesolowski R, Ahn DH, Wu C, et al. Phase I immunotherapy trial with two chimeric HER-2 B-cell peptide vaccines emulsified in montanide ISA 720VG and Nor-MDP adjuvant in patients with advanced solid tumors. *Clin Cancer Res*. 2019;25:3495-3507.

6. Horn L, Infante JR, Reckamp KL, et al. Ensartinib (X-396) in ALK-positive non-small cell lung cancer: results from a first-in-human phase I/II, multicenter study. *Clin Cancer Res*. 2018;24:2771-2779.
7. Patil VM, Noronha V, Joshi A, et al. Phase I/II study of palliative triple metronomic chemotherapy in platinum-refractory/early-failure oral cancer. *J Clin Oncol*. 2019;37:3032-3041.
8. Yan F, Mandrekar SJ, Yuan Y. Keyboard: a novel Bayesian toxicity probability interval design for phase I clinical trials. *Clin Cancer Res*. 2017;23:3994-4003.
9. Guo W, Wang SJ, Yang S, Lynn H, Ji Y. A Bayesian interval dose-finding design addressing Ockham's razor: mTPI-2. *Contemp Clin Trials*. 2017;58:23-33.
10. Yuan Z, Chappell R, Bailey H. The continual reassessment method for multiple toxicity grades: a Bayesian quasi-likelihood approach. *Biometrics*. 2007;63:173-179.
11. Ji Y, Li Y, Nebiyou Bekele B. Dose-finding in phase I clinical trials based on toxicity probability intervals. *Clin Trials*. 2007;4:235-244.
12. Ivanova A, Flournoy N, Chung Y. Cumulative cohort design for dose-finding. *J Stat Plann Infer*. 2007;137:2316-2327.
13. Sverdlov O, Wong WK. Novel statistical designs for phase I/II and phase II clinical trials with dose-finding objectives. *Ther Innov Regul Sci*. 2014;48:601-612.
14. Braun TM. The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Contemp Clin Trials*. 2002;23:240-256.
15. Thall PF, Cook JD. Dose-finding based on efficacy–toxicity tradeoffs. *Biometrics*. 2004;60:684-693.
16. Yin G, Li Y, Ji Y. Bayesian dose-finding in phase I/II clinical trials using toxicity and efficacy odds ratios. *Biometrics*. 2006;62:777-787.
17. Zang Y, Lee J, Yuan Y. Adaptive designs for identifying optimal biological dose for molecularly targeted agents. *Clin Trials*. 2014;11:319-327.
18. Mozgunov P, Jaki T. An information theoretic phase I–II design for molecularly targeted agents that does not require an assumption of monotonicity. *J R Stat Soc Ser C Appl Stat*. 2019;68:347-367.
19. Hoering A, LeBlanc M, Crowley J. Seamless phase I-II trial design for assessing toxicity and efficacy for targeted agents. *Clin Cancer Res*. 2011;17:640-646.
20. Riviere MK, Yuan Y, Jourdan JH, Dubois F, Zohar S. Phase I/II dose-finding design for molecularly targeted agent: plateau determination using adaptive randomization. *Stat Methods Med Res*. 2016;27:466-479.
21. Liu S, Johnson VE. A robust Bayesian dose-finding design for phase I/II clinical trials. *Biostatistics*. 2016;17:249-263.
22. Li DH, Whitmore JB, Guo W, Ji Y. Toxicity and efficacy probability interval design for phase I adoptive cell therapy dose-finding clinical trials. *Clin Cancer Res*. 2017;23:13-20.
23. Lin R, Yin G. STEIN: a simple toxicity and efficacy interval design for seamless phase I/II clinical trials. *Stat Med*. 2017;36:4106-4120.
24. Takeda K, Taguri M, Morita S. BOIN-ET: Bayesian optimal interval design for dose finding based on both efficacy and toxicity outcomes. *Pharm Stat*. 2018;17:383-395.
25. Zhou Y, Lee JJ, Yuan Y. A utility-based Bayesian optimal interval (U-BOIN) phase I/II design to identify the optimal biological dose for targeted and immune therapies. *Stat Med*. 2019;38(28):5299–5316. https://doi.org/10.1002/sim.8361.
26. Houede N, Thall PF, Nguyen H, et al. Utility-based optimization of combination therapy using ordinal toxicity and efficacy in phase I/II trials. *Biometrics*. 2010;66:532-540.
27. Lee J, Thall PF, Ji Y, et al. Bayesian dose-finding in two treatment cycles based on the joint utility of efficacy and toxicity. *J Am Stat Assoc*. 2015;110:711-722.
28. Gourieroux C, Monfort A, Trognon A. Pseudo maximum likelihood methods: theory. *Econometrica*. 1984;52:681-700.
29. McCullagh P, Nelder J. *Generalized Linear Models*. 2nd ed. New York, NY: Chapman & Hall; 1989.
30. Bril G, Dykstra R, Pillers C, Robertson T, Algorithm AS. 206: isotonic regression in two independent variables. *J R Stat Soc Ser C Appl Stat*. 1984;33:352-357.
31. Lin R, Yuan Y. Time-to-event model-assisted designs for dose-finding trials with delayed toxicity. *Biostatistics*. 2019;21(4):807-824. https://doi.org/10.1093/biostatistics/kxz007.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Shi H, Cao J, Yuan Y, Lin R. uTPI: A utility-based toxicity probability interval design for phase I/II dose-finding trials. *Statistics in Medicine*. 2021;40:2626–2649. https://doi.org/10.1002/sim.8922

APPENDIX A

A.1 Estimation of toxicity and efficacy probabilities

At the end of the trial, the toxicity and efficacy probabilities can be estimated using the model-averaging approach proposed by Lin and Yin.[23] For the toxicity outcome, an isotonic regression is performed on the observed toxicity probabilities $\{\hat{\pi}_{j,T}\}_{j=1}^{J}$ through the pool-adjacent-violators algorithm, and let $\{\tilde{\pi}_{j,T}\}_{j=1}^{J}$ denote the isotonically transformed values. For the efficacy outcome, the dose-response relationship is typically unknown, and a model averaging approach is used to estimate the efficacy probabilities. Specifically, we enumerate all possible modes in the dose-efficacy curve, and perform J unimodal isotonic regressions on the observed efficacy probabilities $\{\hat{\pi}_{j,E}\}_{j=1}^{J}$. Given the mode being dose k, $k = 1, \ldots, J$, that is, the dose level k attains the highest efficacy at all dose levels, the unimodal-isotonically transformed values $\{\tilde{\pi}_{j,E}^{(k)}\}_{j=1}^{J}$ enjoy the following relationship:

$$\tilde{\pi}_{k,E}^{(1)} \leq \ldots \leq \tilde{\pi}_{k,E}^{(k)}, \quad \text{and} \quad \tilde{\pi}_{k,E}^{(k)} \geq \ldots \geq \tilde{\pi}_{J,E}^{(k)}.$$

By applying unimodal isotonic regression, we obtain J sets of possible isotonic estimates $\{\tilde{\pi}_{j,E}^{(k)}\}_{j=1}^{J}, k = 1, \ldots, J$. The pseudo likelihood based on the kth unimodal isotonic regression is given by

$$\tilde{L}_k = \prod_{j=1}^{J} \binom{n_j}{n_{j,E}} \tilde{\pi}_{j,E}^{(k)^{n_{j,E}}} \left(1 - \tilde{\pi}_{j,E}^{(k)}\right)^{n_j - n_{j,E}}.$$

The Akaike information criterion for pseudo likelihood is $\text{AIC}_k = -2 \log \tilde{L}_k + 2J$. The final model averaging estimate of the efficacy probability at dose level j is

$$\tilde{\pi}_{j,E} = \sum_{k=1}^{J} w_k \tilde{\pi}_{j,E}^{(k)},$$

where w_k is the weight function,

$$w_k = \frac{\exp(-\text{AIC}_k/2)}{\sum_{j=1}^{J} \exp(-\text{AIC}_j/2)} = \frac{\tilde{L}_k}{\sum_{j=1}^{J} \tilde{L}_j}.$$

Finally, the dose desirability can then be computed based on (2) and OBD can be selected based on (5).

A.2 Random scenario generator

To objectively assess the operating characteristics, we generate random toxicity and efficacy scenarios using a variant of the pseudo-uniform approach, the steps of which are detailed as follows.

1. The MTD is selected from the prespecified dose levels, $j^{\text{MTD}} \in \{1, \ldots, J\}$. Repeat the following steps until the toxicity probability of the dose level j^{MTD} is closest to ϕ and within ϵ_ϕ absolute distance from ϕ.
   a. Randomly generate a variable $M \sim \text{Beta}(\max(J - j^{\text{MTD}}, 0.5), 1)$, and then set the upper bound of the toxicity probability as $p_U = \max(\phi + 0.025, \phi + (1 - \phi)M)$.
   b. Construct the dose-toxicity curve by drawing an ordered sample of J independent uniform random variables from $\text{Unif}(0, p_U)$.
2. Let j^{MED} and j^{OBD} denote the MED (maximum efficacious dose), that is, the dose with largest efficacy probability, and OBD, that is, the dose with the highest desirability, respectively. The MED is selected from the prespecified dose levels, $j^{\text{MED}} \in \{1, \ldots, J\}$. Repeat the following steps until efficacy probabilities of the dose levels j^{MED} and j^{OBD} are greater than $\psi - \epsilon_\psi$.
   a. Let q_L and q_U denote lower and upper boundaries of the efficacy probability, respectively. Set the lower bound as q_L and draw the upper bound q_U from a uniform distribution $\text{Unif}(r_L, r_U)$.
   b. Randomly draw J independent uniform random variables from $\text{Unif}(q_L, q_U)$, and sort them such that the dose level j^{MED} has the largest efficacy probability.

c. To ensure that a certain proportion of the dose-efficacy curve has a plateau shape, with a probability of w, that is, the proportion of plateau-shape curves, we restrict $\pi_{j',E} = \pi_{j^{\text{MED}},E}$ for $j^{\text{MED}} \leq j' \leq J$.
d. The desirability vector is computed from (2) and the dose with the maximum desirability j^{OBD} is selected as the optimal biological dose.

We take $\epsilon_\phi = 0.10$ and $\epsilon_\psi = 0.05$, $q_L = 0.01$, $r_L = 0.4$, $r_U = 0.8$. Moreover, we assume that 25% of the dose-efficacy curves are plateau shaped, thus we set $w = 0.25$.