Biostatistics (2023) **24**, 2, *pp.* 277–294
doi: 10.1093/biostatistics/kxab028
Advance Access publication on July 23, 2021

Bayesian adaptive model selection design for optimal biological dose finding in phase I/II clinical trials

RUITAO LIN*

Department of Biostatistics, The University of Texas MD Anderson Cancer Center, 1515 Holcombe Blvd, Houston, TX 77030, USA

rlin@mdanderson.org

GUOSHENG YIN

Department of Statistics and Actuarial Science, The University of Hong Kong, Pokfulam Road, Hong Kong, China

HAOLUN SHI

Department of Statistics and Actuarial Science, Simon Fraser University, 8888 University Dr, Burnaby, BC V5A 1S6, Canada

Summary

Identification of the optimal dose presents a major challenge in drug development with molecularly targeted agents, immunotherapy, as well as chimeric antigen receptor T-cell treatments. By casting dose finding as a Bayesian model selection problem, we propose an adaptive design by simultaneously incorporating the toxicity and efficacy outcomes to select the optimal biological dose (OBD) in phase I/II clinical trials. Without imposing any parametric assumption or shape constraint on the underlying dose–response curves, we specify curve-free models for both the toxicity and efficacy endpoints to determine the OBD. By integrating the observed data across all dose levels, the proposed design is coherent in dose assignment and thus greatly enhances efficiency and accuracy in pinning down the right dose. Not only does our design possess a completely new yet flexible dose-finding framework, but it also has satisfactory and robust performance as demonstrated by extensive simulation studies. In addition, we show that our design enjoys desirable coherence properties, while most of existing phase I/II designs do not. We further extend the design to accommodate late-onset outcomes which are common in immunotherapy. The proposed design is exemplified with a phase I/II clinical trial in chronic lymphocytic leukemia.

Keywords: Bayesian adaptive design; Bayesian model selection; Delayed response; Optimal biological dose; Phase I/II trial.

1. Introduction

Identification of the optimal biological dose (OBD) in early-phase cancer trials plays an increasingly important role in new drug development of targeted agents, immunotherapy, as well as chimeric

*To whom correspondence should be addressed.

antigen receptor (CAR) T-cell therapy. Compared with cytotoxic agents in oncology, these novel agents are generally tolerable, and thus the optimal doses may not be subject to severe side effects (Hoering *and others*, 2011). In contrast to the traditional phase I clinical trials, which often focus on determination of the maximum tolerated dose, the primary objective of dose finding for immunotherapy and targeted therapy is to locate the OBD which is defined as the lowest safe dose that has the highest efficacy (Yan *and others*, 2018). By considering both toxicity and efficacy outcomes simultaneously, phase I/II designs can seamless maximize the therapeutical effect of the new drug while controlling its adverse effect, which enhances the trial efficiency as well as the ethics of early drug development.

However, the OBD determination in phase I/II trials is associated with many challenges, among which the most prominent one is that the shape of the dose–efficacy curve is typically unknown and may not follow a monotonic pattern in contrast to the dose–toxicity relationship. In practice, the dose–efficacy curves for targeted agents and immunotherapy are often expected to have three main shapes: a monotone increasing or decreasing pattern; an umbrella shape which attains the maximum therapeutical effect in the middle of the dose range; or a plateau shape, that is, the dose–efficacy curve increases initially and then remains flat as the dose continues to increase. Without knowing *a priori* the specific form of the dose–efficacy relationship, it is difficult for the conventional toxicity-based dose-finding designs such as the $3+3$ design and the continual reassessment method (CRM) to identify the OBD accurately. Many phase I/II designs have been developed to capture the relationship between the efficacy probability and the dose level (Yin, 2012). For example, Thall and Cook (2004) adopted a quadratic form to quantify the non-monotone pattern of the dose–efficacy curve and introduced a toxicity–efficacy trade-off contour for dose finding. Zang *and others* (2014) proposed two designs: one is called the isotonic design, which continually estimates the dose–efficacy curve by a nonparametric isotonic regression approach; and the other has a semiparametric spirit as it assumes a logistic model for the dose–efficacy relationship locally around the current dose. Riviere *and others* (2018) utilized an adaptive randomization scheme to determine the plateau in phase I/II dose-finding trials. Takeda *and others* (2018) developed a Bayesian optimal interval design by considering both toxicity and efficacy outcomes. Other examples of phase I/II designs can be found in Braun (2002), Yin *and others* (2006), Wages and Tait (2015), Muenz *and others* (2019), Li *and others* (2020), Lin *and others* (2020), among others. Phase I/II designs have also be considered for more complex early-phase trials, for example, the drug-combination trials (Yuan and Yin, 2011), personalized dose finding (Guo and Yuan, 2017), incorporation of immune responses (Wang *and others*, 2019), accommodating time-to-event outcomes (Takeda *and others*, 2020), among others. For a comprehensive review on the existing methods of phase I/II trials, see Mandrekar *and others* (2010) and Yuan *and others* (2017).

Our research is motivated by a phase I/II dose-finding study in patients with relapsed or refractory chronic lymphocytic leukemia (CLL) (Van Den Neste *and others*, 2013). Preclinical studies have demonstrated *in vitro* that acadesine activates adenosine monophosphate-activated protein kinase and selectively induces apoptosis in B cells of CLL patients. The primary objective of this multicenter trial was to investigate the safety as well as efficacy of acadesine in CLL patients. Most of the existing phase I/II designs impose parametric assumptions to quantify the dose–toxicity and dose–efficacy curves, which however may not be reliable when the assumed models clearly deviate from the true dose–response profiles. To relax the rigidity in parametric model specifications, we propose an adaptive dose-finding design to search for the OBD under the framework of Bayesian model selection. A key feature that makes our design essentially different from existing ones is that our design estimates the optimal dose directly, instead of using an indirect way by estimating the dose–response curves. A set of curve-free models is specified for toxicity to examine which dose level is the cutoff for the tolerable dose set; and another set of curve-free models is imposed for efficacy to test which of the tolerable doses is the most efficacious one. Our prespecified models require minimum constraints yet can capture both monotone and nonmonotone dose–response relationships. As a result, the proposed design enjoys both robustness and adaptiveness in the sense that each dose assignment is determined optimally according to some criterion by integrating all the available

data in the trial. Furthermore, our design is shown to be coherent in the sequence of dose assignment, which enhances its ethical implications. To the best of our knowledge, the coherence property has never been studied for existing phase I/II trial designs. We further generalize the design to address the issue of late-onset outcomes, which is another prominent challenge with novel targeted agents and immunotherapy.

The remainder of this article is organized as follows. In Section 2, we formulate the Bayesian adaptive model selection design and present its theoretical properties. In Section 3, we examine the performance of the new design via simulation studies and make extensive comparisons with existing methods. As an illustration, we apply our design to a phase I/II dose-finding trial with acadesine therapy in Section 4, and Section 5 concludes with some remarks. Technical details and additional simulation studies are presented in the Supplementary material available at *Biostatistics* online. R codes to implement the proposed methods can be found from https://github.com/ruitaolin/BAMS.

2. Methodology

2.1. *Probability model*

Consider a phase I/II trial with J dose levels, let $p_{\mathrm{T}j}$ and $p_{\mathrm{E}j}$ denote the toxicity and efficacy probabilities at dose level j respectively, for $j = 1, \ldots, J$. Let $d_i \in \{1, \ldots, J\}$ denote the dose level assigned to the ith patient, and let $x_{\mathrm{T}i}$ and $x_{\mathrm{E}i}$ represent the corresponding binary dose-limiting toxicity (DLT) and efficacy outcomes, with a value of 1 indicating that the patient has experienced the corresponding outcome, and 0 otherwise. To facilitate the development of our method, we assume that the bivariate outcomes of all treated patients can be quickly observed without delay, while its extension to late-onset responses is given in Section 2.5. Let $D_n = \{(d_1, x_{\mathrm{E}1}, x_{\mathrm{T}1}), \ldots, (d_n, x_{\mathrm{E}n}, x_{\mathrm{T}n})\}$ represent the data accumulated up to the nth patient, and then $y_{\mathrm{E}j} = \sum_{i=1}^{n} x_{\mathrm{E}i} I\{d_i = j\}$, $y_{\mathrm{T}j} = \sum_{i=1}^{n} x_{\mathrm{T}i} I\{d_i = j\}$ and $m_j = \sum_{i=1}^{n} I\{d_i = j\}$ are the numbers of efficacy responses, DLTs and patients treated at dose level j, respectively, where $I\{\cdot\}$ is the indicator function.

In dose-finding trials, the toxicity probabilities are typically assumed to increase monotonically with dose levels, that is, $p_{\mathrm{T}1} < \cdots < p_{\mathrm{T}J}$, while the dose–efficacy relationship is not subject to such a monotonic constraint. Specifically, the dose–efficacy curve may be monotonically increasing, decreasing, umbrella-shaped, or plateau-shaped; that is, the efficacy probabilities may first increase and then decrease or remain flat in the plateau region as the dose increases. Therefore, any of the J dose levels may attain the largest efficacy probability, that is, $p_{\mathrm{E}1} \leq \cdots \leq p_{\mathrm{E}j}$ and $p_{\mathrm{E}j} \geq \cdots \geq p_{\mathrm{E}J}$, for any $j \in \{1, \ldots, J\}$. Let ϕ_{T} denote the prespecified maximum tolerable toxicity probability, and then the admissible set $\mathcal{A}_{\mathrm{T}}$ contains all the dose levels with the toxicity probabilities less than ϕ_{T}, $\mathcal{A}_{\mathrm{T}} = \{j : p_{\mathrm{T}j} \leq \phi_{\mathrm{T}}, j = 1, \ldots, J\}$. To handle possible ties, we define the OBD as the lowest dose level in $\mathcal{A}_{\mathrm{T}}$ that has the largest efficacy probability; that is

$$j^{\mathrm{OBD}} = \min\{\mathcal{S}\}, \quad \text{where } \mathcal{S} = \Big\{j : j = \underset{j' \in \mathcal{A}_{\mathrm{T}}}{\operatorname{argmax}}\{p_{\mathrm{E}j'}\}\Big\}.$$

Alternatively, the OBD can also be defined as the dose level that is the most efficacious in $\mathcal{A}_{\mathrm{T}}$ (Mozgunov and Jaki, 2019). The two definitions coincide when the efficacy curve is monotone or has an umbrella shape, while multiple optimal doses may satisfy the second definition for a dose–efficacy curve with a plateau.

To identify the admissible set $\mathcal{A}_{\mathrm{T}}$, we consider a Bayesian model selection procedure with $J+1$ toxicity models,

$$\mathcal{M}_{\mathrm{T}k} : \begin{cases} \mathcal{A}_{\mathrm{T}} = \varnothing, & k = 0, \\ \mathcal{A}_{\mathrm{T}} = \{1, \ldots, k\}, & k = 1, \ldots, J, \end{cases}$$

where $\mathcal{M}_{T0}$ indicates that none of the J toxicity probabilities satisfies $p_{Tj} \le \phi_T$, that is, all the J doses are excessively toxic, and $\mathcal{M}_{Tk}$ $(k = 1, \ldots, J)$ represents that the kth dose is the highest level in the admissible set $\mathcal{A}_T$. Under each toxicity model $\mathcal{M}_{Tk}$, $k = 0, \ldots, J$, we specify a sequential uniform prior distribution $\pi_T(p_{T1}, \ldots, p_{TJ}|\mathcal{M}_{Tk})$ to maintain the monotonically increasing pattern of the dose–toxicity curve,

$$\begin{cases} p_{Tj}|\mathcal{M}_{Tk} \sim \text{Unif}(0, \phi_T), & j = k, \\ p_{Tj}|\mathcal{M}_{Tk}, p_{T,j+1} \sim \text{Unif}(0,\ \min(p_{T,j+1}, \phi_T)), & j < k, \\ p_{Tj}|\mathcal{M}_{Tk}, p_{T,j-1} \sim \text{Unif}(\max(\phi_T, p_{T,j-1}),\ 1), & j > k, \end{cases} \tag{2.1}$$

where $p_{T0} = 0$ and $p_{T,J+1} = 1$ are the lower and upper boundaries of the prior distribution, respectively.

As the dose–efficacy curve may possess a nonmonotone pattern, we consider J models for the efficacy probabilities,

$$\mathcal{M}_{Ek} : p_{Ek} = \max_{1 \le j \le J} \{p_{Ej}\}, \quad k = 1, \ldots, J,$$

which indicates that under model $\mathcal{M}_{Ek}$ dose level k attains the largest efficacy probability among the J dose levels. The joint prior distribution $\pi_E(p_{E1}, \ldots, p_{EJ}|\mathcal{M}_{Ek})$ under model $\mathcal{M}_{Ek}$, $k = 1, \ldots, J$, is given by

$$\begin{cases} p_{Ej}|\mathcal{M}_{Ek} \sim \text{Unif}(\delta_E, 1), & j = k, \\ p_{Ej}|\mathcal{M}_{Ek}, p_{E,j+1} \sim \text{Unif}(0, p_{E,j+1}), & j < k, \\ p_{Ej}|\mathcal{M}_{Ek}, p_{E,j-1} \sim \text{Unif}(0, p_{E,j-1}), & j > k, \end{cases} \tag{2.2}$$

where δ_E is the exploratory efficacy cutoff for searching the maximum efficacious dose level, and it controls the degree of aggressiveness of the design. Usually, a large value of δ_E facilitates more exploration of untried tolerable doses, while a small value of δ_E might cause the trial to be trapped in some suboptimal doses. In practice, δ_E can be chosen as the response rate under the alternative hypothesis that would be used in a phase II trial. The default setting $\delta_E \in [0.25, 0.40]$ generally results in satisfactory performance of the proposed design given that the efficacy rate of the OBD is below 0.5, which is the case for most cancer treatments (Paoletti and Postel-Vinay, 2018). However, if it is expected that several doses may have large efficacy probabilities, for example, $p_{Ej} > 0.5$, we recommend to increase the value of δ_E to 0.5 or even a higher value.

Alternatively, we consider a data-adaptive choice for δ_E by treating δ_E as a function of the sample size as well as the observed efficacy rate at the current dose. Specifically, suppose that j is the current dose level, and we set δ_E as

$$\delta_E = \begin{cases} 0.5, & \text{if } m_j \ge N^* \text{ and } y_{Ej}/m_j \ge 0.25; \\ 0.25, & \text{otherwise.} \end{cases}$$

Such a data-adaptive choice for δ_E indicates that if the number of patients at the current dose is sufficiently large (say, greater than $N^* = 12$), and the corresponding observed efficacy rate is also acceptable, then we assign a larger value to δ_E to facilitate more exploration of other potential doses (i.e., untried doses or promising doses yet tested adequately). This scheme inherits the exploitation–exploration spirit of reinforcement learning, and thus it can prevent the trial from being stuck at some suboptimal doses without sacrificing too much on safety. We recommend using such a data-adaptive scheme as our numerical studies

show that it is more robust and can flexibly accommodate various toxicity/efficacy profiles. The value of δ_{E} can be further increased above 0.5 according to the trial's characteristics or the need of clinical investigators. For example, for trials that anticipate $\geq$ 70% response rates, setting the maximum value of δ_E larger than 0.5 may result in a better design performance. Nevertheless, even under the default setting, as demonstrated by extensive simulation studies, the proposed design still delivers a competitive performance in comparison with other designs.

Based on the observed data D_n, the marginal likelihood under model $\mathcal{M}_{lk}$ is

$$P(D_n|\mathcal{M}_{lk}) \propto \int \cdots \int \pi_l(p_{l1}, \ldots, p_{lJ}|\mathcal{M}_{lk}) \prod_{j=1}^{J} \left\{ p_{lj}^{y_{lj}} (1-p_{lj})^{m_j - y_{lj}} \right\} \mathrm{d}p_{l1} \cdots \mathrm{d}p_{lJ}, \quad l \in \{\mathrm{T}, \mathrm{E}\}$$

where $\pi_l(p_{l1}, \ldots, p_{lJ}|\mathcal{M}_{lk})$ represents the joint prior distribution of $\{p_{l1}, \ldots, p_{lJ}\}$ under model $\mathcal{M}_{lk}$. Thus, the posterior probabilities of model $\mathcal{M}_{\mathrm{T}k}$ and $\mathcal{M}_{\mathrm{E}k}$ are respectively given by

$$P(\mathcal{M}_{\mathrm{T}k}|D_n) = \frac{P(D_n|\mathcal{M}_{\mathrm{T}k})P(\mathcal{M}_{\mathrm{T}k})}{\sum_{j=0}^{J} P(D_n|\mathcal{M}_{\mathrm{T}j})P(\mathcal{M}_{\mathrm{T}j})}, \quad P(\mathcal{M}_{\mathrm{E}k}|D_n) = \frac{P(D_n|\mathcal{M}_{\mathrm{E}k})P(\mathcal{M}_{\mathrm{E}k})}{\sum_{j=1}^{J} P(D_n|\mathcal{M}_{\mathrm{E}j})P(\mathcal{M}_{\mathrm{E}j})},$$

where $P(\mathcal{M}_{lk})$ is the prior probability of $\mathcal{M}_{lk}, l \in \{\mathrm{T}, \mathrm{E}\}$. The prior information about the drug's safety and efficacy can be elicited from historical data or clinicians' expertise knowledge, which can be incorporated into the prior model probability $P(\mathcal{M}_{\mathrm{T}k})$. If there is little information *a priori*, we specify a discrete uniform distribution for the prior model probability, for example, $P(\mathcal{M}_{\mathrm{T}k}) = 1/(J+1)$ and $P(\mathcal{M}_{\mathrm{E}k}) = 1/J$.

We model the toxicity and efficacy outcomes independently, while it is possible to take their correlation into account, which however requires more complicated models and intensive computation. In the literature, it has been demonstrated that the approach of joint modeling does not improve the design performance (Yin *and others*, 2006), which is also confirmed via a sensitivity analysis in the Supplementary material available at *Biostatistics* online.

2.2. *Bayesian adaptive model selection design*

Most of the existing phase I/II designs cast dose finding into a regression problem of estimating the dose–toxicity and dose–efficacy curves. By contrast, we propose to determine the next dose level using the posterior model probabilities $P(\mathcal{M}_{lk}|D_n)$. The posterior probability of $p_{\mathrm{T}j} \leq \phi_{\mathrm{T}}$ can be derived using Bayesian model averaging (BMA),

$$P(p_{\mathrm{T}j} \leq \phi_{\mathrm{T}}|D_n) = \sum_{k=0}^{J} P(\mathcal{M}_{\mathrm{T}k}|D_n)P(p_{\mathrm{T}j} \leq \phi_{\mathrm{T}}|\mathcal{M}_{\mathrm{T}k}, D_n) = \sum_{k=j}^{J} P(\mathcal{M}_{\mathrm{T}k}|D_n),$$

where the second equality is inherited from the definition of $\mathcal{M}_{\mathrm{T}k}$. Given j, if $k < j$, $\mathcal{M}_{\mathrm{T}k}$ specifies that $p_{\mathrm{T}j} > \phi_{\mathrm{T}}$ and, as a result, $P(p_{\mathrm{T}j} \leq \phi_{\mathrm{T}}|\mathcal{M}_{\mathrm{T}k}, D_n) = 0$; if $k \geq j$, $\mathcal{M}_{\mathrm{T}k}$ specifies that $p_{\mathrm{T}j} \leq \phi_{\mathrm{T}}$ and, as a result, $P(p_{\mathrm{T}j} \leq \phi_{\mathrm{T}}|\mathcal{M}_{\mathrm{T}k}, D_n) = 1$.

The admissible set $\mathcal{A}_{\mathrm{T}}$ is estimated by $\widehat{\mathcal{A}}_{\mathrm{T}}(D_n) = \{j : P(p_{\mathrm{T}j} \leq \phi_{\mathrm{T}}|D_n) > \delta_{\mathrm{T}}\}$, and the highest dose level in $\widehat{\mathcal{A}}_{\mathrm{T}}(D_n)$ is denoted by $j_{\mathrm{T}n}^{\max} = \max\left\{\widehat{\mathcal{A}}_{\mathrm{T}}(D_n)\right\}$, where $\delta_{\mathrm{T}} > 0$ is a prespecified posterior probability cutoff for toxicity that determines the conservativeness of a trial. The smaller the value of δ_{T}, the more aggressive the dose allocation. When the OBD is located in the lower dose region, a larger value of δ_{T} is more desirable as it tends to rule out the higher dose levels. However, if the OBD is located in the higher dose region, an overly large value of δ_{T} may lead to exclusion of the OBD from the admissible set. In practice, the location of the OBD is unknown *a priori* and thus the optimal value of δ_{T} cannot be

determined. In general, we recommend $\delta_T \in [0.10, 0.20]$, which leads to effective exclusion of overly toxic dose levels.

For the efficacy outcome, the most efficacious dose level based on the current data D_n is the one that attains the largest posterior probability $\mathcal{M}_{Ej}$, that is, $j_{En}^{\max} = \arg\max_{1\le j\le J} P(\mathcal{M}_{Ej}|D_n)$. When multiple tolerable dose levels have the same largest posterior efficacy model probability, we would choose the lowest one from them as it has the smallest toxicity rate. By considering the toxicity and efficacy outcomes simultaneously, the optimal dose level based on the cumulative data is

$$j_n^* = \min(j_{Tn}^{\max}, j_{En}^{\max}). \tag{2.3}$$

If the efficacious dose level $j_{En}^{\max}$ is tolerable as noted by $j_{En}^{\max} \le j_{Tn}^{\max}$, then $j_{En}^{\max}$ is optimal; otherwise the optimal level is $j_{Tn}^{\max}$ due to the underlying umbrella-shaped dose–efficacy curve.

At the beginning of a trial, the posterior estimates are typically unstable due data sparsity, which often causes the dose movement to be trapped in some lower dose levels. To stabilize the estimates as well as accelerating the trial conduct, we implement a start-up phase as follows. The first cohort of patients is treated at the lowest dose level, and if no toxicity or efficacy outcome is observed, we escalate one dose level for the next cohort of patients. The start-up phase continues escalation until the first toxicity or efficacy outcome is observed or the highest dose level is reached, and then the trial enters seamlessly into the main phase. With a cohort size of three patients, when the dose assignment reaches dose level j at the end of the start-up phase, we would have at least one toxicity/efficacy outcome at dose level j, and a total of $3 \times j$ patients have been evenly allocated across the first j dose levels.

In the main phase, suppose that n patients have been treated and the latest cohort is assigned dose level j, then the optimal dose level j_n^* is determined according to (2.3). For the next cohort of patients, (i) if $j > j_n^*$, de-escalate to the lower dose level $j-1$; (ii) if $j < j_n^*$, escalate to the higher dose level $j+1$; and (iii), otherwise, retain at the same dose level j. The trial will be stopped after exhaustion of a total of N patients or be terminated early due to safety concerns.

As the proposed design can adaptively integrate all the observed data under the framework of Bayesian model selection for decision making, we name it as the Bayesian adaptive model selection (BAMS) design. In our proposal, the BAMS design adaptively allocates patients to a safe, highly efficacious dose that is determined using all the observed data. As the proposed model selection framework is flexible, alternative dose-finding rules can be incorporated. One choice is to adapt BAMS into a Bayesian decision-making framework by incorporating a loss function, as shown in the Supplementary material available at *Biostatistics* online.

2.3. *Coherence property*

The BAMS design enjoys several prominent features: First, the trial design is model-free, as no parametric assumption is imposed on either the dose–toxicity or dose–efficacy curves; Second, BAMS makes decisions based on Bayesian posterior model probabilities rather than posterior estimates of the toxicity and efficacy probabilities, which not only accounts for the observed number of toxicity and efficacy events but also reflects the sample size; Third, the dose movement of BAMS is coherent as indicated by the next theorem.

THEOREM 2.1 Suppose that the dose level for the nth patient is d_n.

(1) If $d_n = j_{T,n-1}^{\max}$ and $x_{Tn} = 1$, the BAMS design is coherent in dose escalation, that is, for the $(n+1)$th patient, $d_{n+1} \le d_n$, with probability one.

(2) If $d_n = j^{\max}_{\mathrm{E},n-1}$ and $(x_{\mathrm{T}n}, x_{\mathrm{E}n}) = (1, 1)$, the BAMS design is coherent in dose escalation, that is, for the $(n+1)$th patient, $d_{n+1} \le d_n$, with probability one.

(3) If $d_n = j^{\max}_{\mathrm{E},n-1}$ and $(x_{\mathrm{T}n}, x_{\mathrm{E}n}) = (0, 1)$, the BAMS design is coherent in dose retainment, that is, for the $(n+1)$th patient, $d_{n+1} = d_n$, with probability one.

In a further elaboration, (1) indicates that when the current dose is the highest one in the admissible set $\widehat{\mathcal{A}}_T(D_{n-1})$, if a DLT is observed for the most recent nth patient, then the highest dose in the updated set $\widehat{\mathcal{A}}_T(D_n) \subset \widehat{\mathcal{A}}_T(D_{n-1})$ satisfies $j^{\max}_{\mathrm{T}n} \le j^{\max}_{\mathrm{T},n-1}$, and thus the dose level for patient $(n+1)$ would not exceed that for the nth patient. And, (2) indicates that when $d_n = j^{\max}_{\mathrm{E},n-1}$, we have $P(\mathcal{M}_{\mathrm{E},d_n}|D_{n-1}) > P(\mathcal{M}_{\mathrm{E},d_n+1}|D_{n-1})$. If $x_{\mathrm{T}n} = 1$, then $\widehat{\mathcal{A}}_T(D_n) \subset \widehat{\mathcal{A}}_T(D_{n-1})$, and further if $x_{\mathrm{E}n} = 1$, then we have $P(\mathcal{M}_{\mathrm{E},d_n}|D_n) > P(\mathcal{M}_{\mathrm{E},d_n+1}|D_n)$. As a result, the next higher dose level will not be considered for treating the next patient by BAMS. In this case, whether to de-escalate or retain the dose then depends on whether or not the updated highest dose in $\widehat{\mathcal{A}}_T(D_n)$, $j^{\max}_{\mathrm{T}n}$, is lower than d_n. Similarly, (3) indicates that when the dose assigned to the nth patient has the largest posterior efficacy model probability based on D_{n-1}, if $x_{\mathrm{T}n} = 0$, then $\widehat{\mathcal{A}}_T(D_n) \supset \widehat{\mathcal{A}}_T(D_{n-1})$, and if further $x_{\mathrm{E}n} = 1$, the posterior model probability $P(\mathcal{M}_{\mathrm{E},d_n}|D_n)$ is strengthened even more. As a result, the current dose level is retained for the next patient.

The proof of Theorem 2.1 is provided in the Supplementary material available at *Biostatistics* online. The coherence principle was first proposed for toxicity-based dose-finding designs (Cheung, 2011). To the best of our knowledge, coherence has not been established for any existing phase I/II trial designs. This finite-sample property provides a theoretical insight into BAMS: Along the trial conduct, each dose assignment is reasonable and intuitive, which enhances practicality of the proposed method.

It is worth noting that Theorem 2.1 only establishes coherence for certain cases and there are other ambiguous cases that cannot be covered by this theorem. For example, the complement to (1) is the case with $d_n = j^{\max}_{\mathrm{T},n-1}$ and $x_{\mathrm{T}n} = 0$. Although we can show that $j^{\max}_{\mathrm{T},n-1} \le j^{\max}_{\mathrm{T}n}$, the determination of the next dose escalation/de-escalation decision requires to compare the posterior efficacy model probabilities among all dose levels. Moreover, in combination, (2) and (3) consider the cases with $d_n = j^{\max}_{\mathrm{E},n-1}$ where $(x_{\mathrm{T}n}, x_{\mathrm{E}n}) = (1, 1)$ or $(x_{\mathrm{T}n}, x_{\mathrm{E}n}) = (0, 1)$. For other cases with $d_n = j^{\max}_{\mathrm{E},n-1}$ where $(x_{\mathrm{T}n}, x_{\mathrm{E}n}) = (0, 0)$ or $(1, 0)$, coherence cannot be established in either dose escalation, de-escalation or retainment, because more information is needed. In particular, when $d_n = j^{\max}_{\mathrm{E},n-1}$ and $(x_{\mathrm{T}n}, x_{\mathrm{E}n}) = (0, 0)$, we can show that $j^{\max}_{\mathrm{T},n-1} \le j^{\max}_{\mathrm{T}n}$, but whether d_n still possesses the largest posterior efficacy model probability is unclear. Thus, decisions of dose escalation, retainment or de-escalation are all possible in this case. Similarly, if $d_n = j^{\max}_{\mathrm{E},n-1}$ and $(x_{\mathrm{T}n}, x_{\mathrm{E}n}) = (1, 0)$, we can show that $j^{\max}_{\mathrm{T},n-1} \ge j^{\max}_{\mathrm{T}n}$ and the posterior efficacy model probability of dose level d_n decreases (in turn, the posterior efficacy model probabilities of other dose levels increase). As a result, whether to escalate, retain or de-escalate the dose depends on the relative magnitude of the posterior efficacy model probabilities of the adjacent doses.

During the trial or at the time for dose estimation, we additionally implement a dose-elimination rule for safety when some doses are excessively toxic or unacceptably sub-therapeutic.

(1) If dose level j satisfies $P(p_{\mathrm{T}j} \le \phi_{\mathrm{T}} \mid D_n) < c_{\mathrm{T}}$, then this dose and all the higher dose levels should be eliminated from the trial due to excessive toxicity.

(2) If dose level j satisfies $P(p_{\mathrm{E}j} \ge \phi_{\mathrm{E}} \mid D_n) < c_{\mathrm{E}}$, then this dose is futile and should be eliminated from the trial.

Once a dose level is eliminated, it would not be considered in the subsequent dose assignments or in the final dose estimation. If all dose levels are eliminated during the trial, then the trial should be terminated early and no OBD is recommended. Here, ϕ_{E}, prespecified by clinicians, is the smallest efficacy probability that is considered to be clinically meaningful; c_{T} and c_{E} are two probability cutoffs, which can be calibrated through simulations so that the probability of incorrect elimination can be controlled at a reasonably low

level. For example, in our default settings with $\phi_T = 0.30$ and $\phi_E = 0.25$, and the cutoff combination $(c_T, c_E) = (0.05, 0.10)$ can generally ensure satisfactory performances.

2.4. *Optimal dose estimation*

After a total of N patients have been treated in the trial, we take a two-step procedure to estimate the OBD. First, we identify the highest safe dose level $j_{TN}^{\max}$ from the estimated admissible set $\widehat{\mathcal{A}}_T(D_N)$. Next, we consider the following J models for efficacy:

$$\begin{cases} p_{Ej}|\tilde{\mathcal{M}}_{Ek} \sim \text{Unif}(0,1), & j = k, \\ p_{Ej}|\tilde{\mathcal{M}}_{Ek}, p_{E,j+1} \sim (1-w)\text{Unif}(0, p_{E,j+1}) + w p_{E,j+1}, & j < k, \\ p_{Ej}|\tilde{\mathcal{M}}_{Ek}, p_{E,j-1} \sim (1-w)\text{Unif}(0, p_{E,j-1}) + w p_{E,j-1}, & j > k, \end{cases} \tag{2.4}$$

for $k = 1, \ldots, J$, where w is a prespecified prior probability that the two adjacent doses have the same efficacy probability. By this way, the plateau shape of a dose–efficacy curve can be accommodated. Another difference between $\tilde{\mathcal{M}}_{Ek}$ and $\mathcal{M}_{Ek}$ is that $\tilde{\mathcal{M}}_{Ek}$ no longer has the exploration parameter δ_E, due to the fact that we need not to try more doses at the final OBD estimation stage. In other words, the more aggressive model specification $\mathcal{M}_{Ek}$ facilitates dose exploration in the dose-finding stage so that more information about the dose–response curves can be collected. At the end of the trial when the OBD is estimated, $\tilde{\mathcal{M}}_{Ek}$ tends to favor lower doses. Given the high variation caused by the sparse data as well as the possibility of a dose–efficacy plateau relationship, we select a lower, efficacious dose level as follows:

$$j_{EN}^{\max} = \min\left\{k : P(\tilde{\mathcal{M}}_{Ek} \mid D_N) \geq \max_{j=1,\ldots,J}\{P(\tilde{\mathcal{M}}_{Ej} \mid D_N)\} - \epsilon, \text{ and } k \notin \mathcal{E}\right\},$$

where $\mathcal{E}$ is the set of dose levels that have been eliminated, and ϵ is a small positive number indicating an indifference or noninferiority margin, for example, $\epsilon = 0.05$. In other words, $j_{EN}^{\max}$ is the lowest admissible dose level that has the posterior model probability no less by ϵ than the maximum one. Due to the incorporation of the indifference parameter ϵ, the definition of $j_{EN}^{\max}$ in the dose-estimation stage differs from that of $j_{En}^{\max}$, $n < N$, in the dose-finding stage. This is because the primary objective of the dose-finding stage is to explore as many promising doses as possible in a reasonably greedy sense, while that of the dose-estimation stage is to identify the safest dose with the highest efficacy. The incorporation of ϵ in the dose-estimation stage can effectively identify a lower dose level that still possesses a high efficacy rate. This is particularly important in small trials where the underlying dose–efficacy curve has a plateau shape, because the posterior model probabilities for the high-efficacy doses tend to be indistinguishable. Finally, the OBD can be determined as

$$j_N^{\text{OBD}} = \min\{j_{TN}^{\max}, j_{EN}^{\max}\}.$$

In addition, the proposed framework also facilitates to estimate the toxicity and efficacy probabilities (p_{Tj}, p_{Ej}). For example, the BMA estimate of p_{Tj} at dose level j is given by

$$\bar{p}_{Tj} = \sum_{k=0}^{J} \hat{p}_{Tj}^{(k)} P(\mathcal{M}_{Tk}|D_N), \tag{2.5}$$

where $\hat{p}_{\mathrm{T}j}^{(k)}$ is the posterior mean of $p_{\mathrm{T}j}$ under model $\mathcal{M}_{\mathrm{T}k}$,

$$\hat{p}_{\mathrm{T}j}^{(k)} = \frac{\int \cdots \int p_{\mathrm{T}j} \pi_{\mathrm{T}}(p_{\mathrm{T}1}, \ldots, p_{\mathrm{T}J} | \mathcal{M}_{\mathrm{T}k}) \prod_{j'=1}^{J} \left\{ p_{\mathrm{T}j'}^{y_{\mathrm{T}j'}} (1 - p_{\mathrm{T}j'})^{m_{j'} - y_{\mathrm{T}j'}} \right\} \mathrm{d}p_{\mathrm{T}1} \cdots \mathrm{d}p_{\mathrm{T}J}}{\int \cdots \int \pi_{\mathrm{T}}(p_{\mathrm{T}1}, \ldots, p_{\mathrm{T}J} | \mathcal{M}_{\mathrm{T}k}) \prod_{j'=1}^{J} \left\{ p_{\mathrm{T}j'}^{y_{\mathrm{T}j'}} (1 - p_{\mathrm{T}j'})^{m_{j'} - y_{\mathrm{T}j'}} \right\} \mathrm{d}p_{\mathrm{T}1} \cdots \mathrm{d}p_{\mathrm{T}J}}.$$

The BMA estimate of $p_{\mathrm{E}j}$ can be derived in a similar way using the models specified in (2.4).

2.5. *Late-onset outcomes*

For many cancer treatments such as radiation therapy or novel targeted therapy, it may take a long period of time to observe the response post treatment. The delay in observing outcomes would cause serious logistical issues for implementing adaptive trial designs, which typically assume that the outcomes can be observed shortly after treatment so that the complete data of all previously treated patients can be used when making decisions for newly enrolled patients. A possible solution to late-onset outcomes is to suspend the accrual until the outcomes of all treated patients are observed. However, this would result in an extremely long duration and repeatedly interrupted trial, which is neither feasible nor ethical by holding incoming patients without treatment. To accommodate late-onset responses and shorten the trial duration, we adopt the idea of time-to-event CRM (Cheung, 2011) into the BAMS design and apply the approximation proposed by Lin and Yuan (2020) to obtain a standard binomial likelihood. More details on the time-to-event BAMS (TITE-BAMS) design are provided in the Supplementary material available at *Biostatistics* online.

3. Numerical studies

3.1. *Simulation configurations*

We conduct extensive simulation studies to examine the finite-sample performance of the proposed BAMS design and also compare it with other existing designs. We take the toxicity upper bound $\phi_{\mathrm{T}} = 0.30$ and the efficacy lower bound $\phi_{\mathrm{E}} = 0.25$. To account for a possible plateau in the dose–efficacy curve, we define the OBD as the lowest tolerable dose that achieves the highest efficacy probability; and define the suboptimal biological doses (s-OBDs) as all the tolerable doses with the highest efficacy, which certainly include the OBD. We consider five dose levels and set the maximum sample size $N = 30$ with a cohort size of three patients. As shown in Figure 1, eight dose–toxicity and dose–efficacy scenarios are specified with various true response probabilities, curve shapes, and locations of the OBD. The outcomes for each patient are simulated under the independence assumption between toxicity and efficacy; while the performance of the proposed method under different correlations is examined in the Supplementary material available at *Biostatistics* online. To avoid cherry-picked scenarios, the Supplementary material available at *Biostatistics* online also contains the same 12 scenarios in Mozgunov and Jaki (2019) with a maximum sample size of 60 (i.e., scenarios A1–A12 in Table S1 of the Supplementary material available at *Biostatistics* online), which were also examined in Wages and Tait (2015) and Riviere *and others* (2018). In particular, the 12 additional scenarios include flatter dose–response curves compared with the eight scenarios considered here.

We examine two versions of the proposed designs: $\mathrm{BAMS_f}$ uses a fixed exploratory efficacy cutoff $\delta_{\mathrm{E}} = 0.35$, and $\mathrm{BAMS_a}$ implements a data-adaptive cutoff as discussed in Section 2.1. For both BAMS designs, we set the posterior probability cutoff for toxicity $\delta_{\mathrm{T}} = 0.15$, the prior probability for a plateaued efficacy curve $w = 0.3$, the indifference margin $\epsilon = 0.05$, and the probability cutoffs for dose elimination $(c_{\mathrm{T}}, c_{\mathrm{E}}) = (0.05, 0.10)$ throughout all simulations. For comparison, we examine the operating characteristics of five

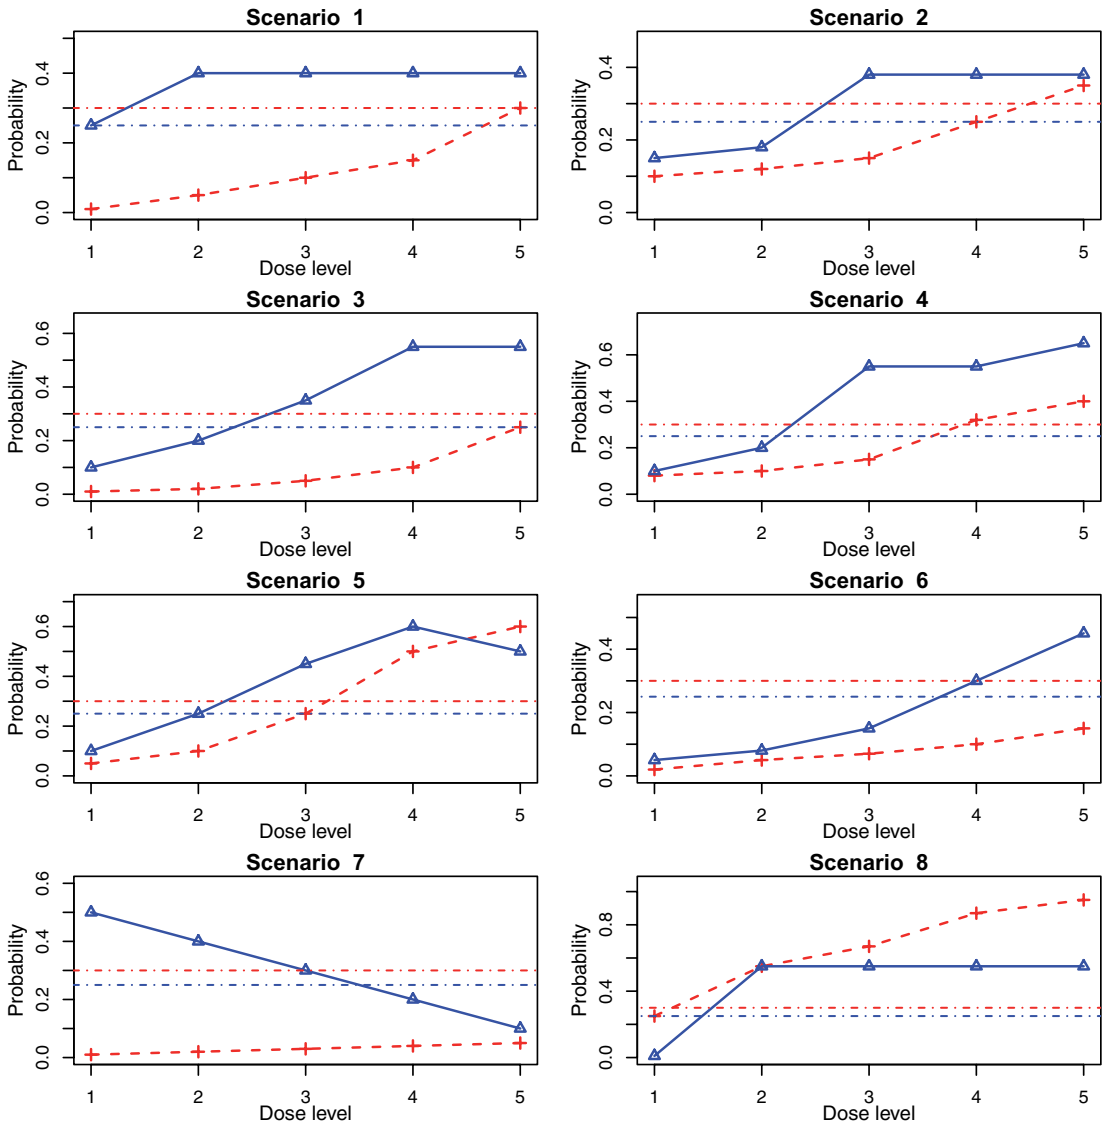

Fig. 1. Eight scenarios with various dose–toxicity (red dashed line) and dose–efficacy (blue solid line) curves in the simulation study. The red (or blue) dash–dotted line corresponds to the prespecified maximum tolerable toxicity probability $\phi_T = 0.30$ (or the minimum efficacy probability $\phi_E = 0.25$), which determines the admissible dose set.

existing dose-finding methods: the WT design (Wages and Tait, 2015), the MTA-RA design (Riviere *and others*, 2018), the L-logit design (Zang *and others*, 2014), the Iso design (Zang *and others*, 2014), and the BOIN-ET design (Takeda *and others*, 2018). The first three are parametric model-based designs and the last two can be viewed as model-free methods. Under all the considered designs, the trial starts treating the first cohort at the lowest dose level. Although different designs may have different utilities or objectives, we have performed extensive calibration to ensure that each design on average achieves the optimal performance across the considered scenarios. Detailed parameter specifications of the five competing designs are given

in the Supplementary material available at *Biostatistics* online. For a benchmark comparison, we also include the nonparametric optimal design (Cheung, 2014). The optimal design assumes, in an oracle way, that the outcomes for each patient are available at all dose levels, and thus it provides a practical upper bound for the accuracy of dose-finding designs. A comparison of the OBD selection percentage of the proposed design with the values under the optimal one shows how far the design performance is from the best possible selection percentage. Under each design, we replicate 5000 simulated trials for each scenario.

Supplementary material available at *Biostatistics* online also contains extensive sensitivity analyses to examine the performance of the proposed BAMS design with respect to different values of the posterior probability cutoff for toxicity δ_T, the exploratory efficacy cutoff δ_E, the correlation between toxicity and efficacy, as well as the prior probability w in the final efficacy models.

3.2. *Simulation results*

The simulation results based on the eight scenarios in Figure 1 are reported in Tables 1 and 2, and those based on the additional 12 scenarios are provided in Figures S1–S5 of the Supplementary material available at *Biostatistics* online. To evaluate the operating characteristics of different designs, we consider five performance statistics: the selection percentage of the OBD and those of the s-OBDs; the numbers of patients treated at the OBD and s-OBDs; and the number of patients treated at overly toxic dose levels as noted by $p_{Tj} \geq \phi_T + 0.05$. Not only do these five statistics quantify the accuracy and efficiency of a design, but they also reflect the ethical aspects of the trial. In general, it is more desirable to treat patients at the most efficacious dose level while controlling the risk of DLTs.

As shown in Table 1, when the highest efficacy rate is moderate, the performances of BAMS_a and BAMS_f are similar. In scenarios 1–4, the efficacy probability initially increases with the dose but plateaus later, and the OBD is lower than the Maximum Tolerated Dose (MTD). In scenario 1, compared with MTA-RA, BAMS_a yields a similar OBD selection percentage, but a 10% higher probability of selecting s-OBDs. In terms of patient allocation, BAMS_a treats most of the patients at the OBD, which in turn leads to a smaller ratio of the number of DLTs to the number of efficacy responses. In scenario 2, the efficacy curve plateaus at dose level 3, but the toxicity probability of dose level 5 exceeds the target toxicity rate. As a result, more aggressive designs, such as L-logit, would allocate more patients to dose level 5, and the WT and BAMS_a designs outperform the others in terms of the selection percentage of s-OBDs. In scenario 3, dose level 4 is the OBD and dose level 5 can be considered as an s-OBD. Although BAMS_a has a superior performance in selecting dose level 4, it tends to treat fewer patients at dose level 5, because BAMS_a uses a similar toxicity monitoring scheme as the design by Lin and Yin (2017). Such an overdose control procedure in BAMS_a also leads to the smallest number of DLTs. In scenario 4, the BOIN-ET and BAMS_a perform equally well, with the selection percentages of s-OBDs close to that of the optimal design. As shown in Table 2, the set of s-OBDs contains only one dose in scenarios 5–6. Under scenario 5, the OBD and the MTD coincides, and BAMS_a yields the best performance. Under scenario 6, the efficacy curve monotonically increases, and BAMS_a ranks the second best in terms of OBD identification. Scenario 7 has a monotonically decreasing efficacy curve, which is the opposite case of scenario 6. The MTA-RA has the highest OBD selection percentage, and BAMS_a performs similarly to the other designs except L-logit. It is desirable to terminate the trial early under scenario 8, because none of the five doses is admissible due to excessive toxicity or futility. According to the last column of Table 2, BAMS_a can quickly identify the inadmissibility of the selected doses with a limited sample size, and thus it has the highest early termination rate. In contrast, the L-logit and Iso designs do not have any futility stopping rules, leading to a large probability of selecting futile doses.

In the eight scenarios considered in Tables 1 and 2, the two proposed designs have similar performances when the efficacy response rate of OBD is not overly high (e.g., $<60\%$). The performances of BAMS_f

Table 1. *Operating characteristics of the BAMS designs in comparison with existing dose-finding methods under scenarios 1–4 with the optimal biological dose (OBD) in boldface and the suboptimal biological doses (s-OBDs) underlined; "Overdoses" meaning the overly toxic dose levels with* $p_{Tj} \geq 0.35$ *and "None" indicating the percentage of inconclusive trials. The maximum sample size is 30*

Design	Selection percentage (no. of patients) at dose level					No. of overdoses	No. of DLTs/Effs	None
	1	2	3	4	5			
				Scenario 1				
(p_T, p_E)	(0.01, 0.25)	**(0.05, 0.40)**	(0.10, 0.40)	(0.15, 0.40)	(0.30, 0.40)			
WT	14.8 (6.9)	**32.9 (9.1)**	26.2 (7.3)	18.5 (4.6)	7.6 (2.1)	0.0	2.5/11.0	0.0
MTA-RA	24.6 (7.7)	**38.4 (8.9)**	20.1 (6.5)	11.6 (4.6)	5.2 (2.3)	0.0	2.6/10.8	0.1
L-logit	11.1 (4.6)	**28.5 (7.9)**	26.4 (7.5)	25.1 (5.7)	9.0 (4.3)	0.0	3.3/11.5	0.0
Iso	17.6 (7.3)	**34.9 (9.9)**	23.9 (6.6)	17.7 (4.0)	5.9 (2.2)	0.0	2.5/10.9	0.0
BOIN-ET	6.5 (4.4)	**24.2 (7.9)**	25.8 (7.0)	27.2 (6.4)	16.3 (4.3)	0.0	3.4/11.3	0.0
$BAMS_f$	14.2 (6.2)	**35.1 (11.3)**	27.9 (7.8)	17.8 (3.6)	5.1 (1.2)	0.0	2.3/11.1	0.0
$BAMS_a$	14.3 (5.9)	**36.1 (11.7)**	28.8 (8.2)	16.8 (3.3)	4.1 (0.9)	0.0	2.3/11.1	0.0
Optimal	3.5	**42.9**	27.5	18.3	7.7	–	–	0.0
				Scenario 2				
(p_T, p_E)	(0.10, 0.15)	(0.12, 0.18)	**(0.15, 0.38)**	(0.25, 0.38)	(0.35, 0.38)			
WT	12.3 (7.9)	10.6 (5.9)	**36.9 (8.1)**	28.8 (5.6)	8.3 (2.2)	2.2	4.8/8.2	3.1
MTA-RA	13.0 (6.4)	13.4 (6.7)	**42.6 (9.0)**	21.9 (5.6)	7.6 (2.2)	2.2	5.0/8.6	1.5
L-logit	13.4 (6.0)	8.3 (5.7)	**37.7 (7.4)**	29.9 (6.4)	10.4 (4.4)	4.4	5.6/8.8	0.2
Iso	25.6 (9.0)	13.5 (6.1)	**31.4 (7.1)**	22.7 (4.9)	6.6 (2.8)	2.8	4.9/8.1	0.0
BOIN-ET	9.4 (4.9)	9.6 (5.3)	**38.6 (8.7)**	29.6 (7.0)	12.2 (4.1)	4.1	5.6/9.2	0.6
$BAMS_f$	8.6 (5.2)	6.2 (5.9)	**41.8 (10.0)**	33.3 (6.8)	8.7 (2.0)	2.0	5.1/8.9	1.5
$BAMS_a$	9.8 (5.4)	6.2 (6.3)	**41.3 (9.7)**	33.2 (6.6)	8.1 (1.8)	1.8	4.9/8.8	1.5
Optimal	0.9	2.4	**61.2**	26.5	8.1	–	–	1.0
				Scenario 3				
(p_T, p_E)	(0.01, 0.10)	(0.02, 0.20)	(0.05, 0.35)	**(0.10, 0.55)**	(0.25, 0.55)			
WT	3.9 (4.7)	7.0 (4.9)	12.8 (5.0)	**35.2 (7.4)**	40.6 (8.0)	0.0	3.0/11.7	0.5
MTA-RA	2.4 (4.1)	9.4 (5.2)	18.3 (6.2)	**41.9 (8.1)**	27.4 (6.3)	0.0	2.9/11.4	0.6
L-logit	4.3 (4.0)	6.5 (5.4)	15.2 (5.2)	**44.8 (6.8)**	29.2 (8.6)	0.0	3.3/11.8	0.0
Iso	14.6 (6.5)	14.5 (6.2)	13.7 (5.2)	**35.7 (6.6)**	21.4 (5.5)	0.0	2.5/10.4	0.0
BOIN-ET	0.8 (3.3)	2.3 (3.9)	10.4 (6.3)	**55.3 (11.2)**	31.3 (5.3)	0.0	2.8/12.4	0.0
$BAMS_f$	2.7 (4.0)	5.5 (5.2)	19.4 (7.6)	**53.7 (10.0)**	18.7 (3.1)	0.0	2.3/11.3	0.1
$BAMS_a$	4.1 (4.1)	6.6 (5.8)	16.0 (7.5)	**54.6 (9.6)**	18.7 (2.9)	0.0	2.1/11.1	0.1
Optimal	0.0	0.1	5.4	**67.3**	27.2	–	–	0.0
				Scenario 4				
(p_T, p_E)	(0.08, 0.10)	(0.10, 0.20)	**(0.15, 0.55)**	(0.32, 0.55)	(0.40, 0.65)			
WT	2.0 (5.5)	4.2 (4.7)	**38.6 (8.8)**	37.8 (7.6)	14.8 (3.1)	3.1	5.7/12.4	2.6
MTA-RA	1.8 (4.4)	5.6 (5.6)	**46.7 (9.5)**	27.8 (6.9)	17.4 (3.6)	3.6	5.9/12.7	0.7
L-logit	6.8 (4.8)	5.3 (4.8)	**42.1 (8.5)**	31.4 (6.7)	14.2 (5.2)	5.2	6.3/13.1	0.1
Iso	14.1 (6.5)	7.4 (4.8)	**45.0 (9.1)**	24.6 (5.9)	8.7 (3.7)	3.7	5.7/12.3	0.0
BOIN-ET	2.1 (3.7)	4.9 (4.5)	**62.4 (14.3)**	24.4 (5.6)	5.9 (1.8)	1.8	5.4/13.4	0.3
$BAMS_f$	1.2 (3.9)	2.4 (4.7)	**63.1 (14.5)**	28.5 (5.9)	4.1 (0.9)	0.9	5.3/13.0	0.8
$BAMS_a$	2.2 (4.0)	2.6 (5.1)	**60.8 (14.1)**	28.9 (5.7)	4.7 (0.9)	0.9	5.2/12.8	1.0
Optimal	0.0	0.44	**72.1**	16.7	10.4	–	–	0.4

Notes: $BAMS_f$ (or $BAMS_a$) is the proposed design with a fixed (or data-adaptive) exploratory cutoff; "Optimal" denotes the nonparametric optimal design (Cheung, 2014); five existing designs include WT (Wages and Tait, 2015), MTA-RA (Riviere *and others*, 2018), L-logit (Zang *and others*, 2014), Iso (Zang *and others*, 2014), and BOIN-ET (Takeda *and others*, 2018).

Table 2. *Operating characteristics of the BAMS designs in comparison with existing dose-finding methods under scenarios 5–8 with the optimal biological dose (OBD) in boldface and the suboptimal biological doses (s-OBDs) underlined; "Overdoses" meaning the overly toxic dose levels with* $p_{Tj} \geq 0.35$ *and "None" indicating the percentage of inconclusive trials. The maximum sample size is 30*

Design	Selection percentage (no. of patients) at dose level					No. of overdoses	No. of DLTs/Effs	None
	1	2	3	4	5			
				Scenario 5				
(p_T, p_E)	(0.05, 0.10)	(0.10, 0.25)	**(0.25, 0.45)**	(0.50, 0.60)	(0.60, 0.50)			
WT	2.5 (5.0)	12.0 (6.4)	**60.7 (11.3)**	22.0 (6.2)	0.7 (0.7)	6.9	7.3/11.3	2.1
MTA-RA	3.3 (4.6)	16.5 (6.9)	**48.9 (9.8)**	27.7 (6.8)	3.1 (1.8)	8.6	7.9/11.5	0.5
L-logit	5.9 (4.5)	22.9 (7.0)	**58.6 (10.2)**	11.6 (5.7)	1.1 (2.6)	8.3	9.0/11.5	0.0
Iso	13.3 (6.2)	26.1 (7.6)	**50.3 (9.1)**	9.8 (4.9)	0.5 (2.1)	7.0	7.1/10.6	0.0
BOIN-ET	7.0 (4.3)	20.8 (7.3)	**60.4 (12.7)**	10.9 (5.0)	0.3 (0.6)	5.6	7.0/11.3	0.7
$BAMS_f$	2.5 (4.0)	13.2 (6.8)	**66.9 (13.8)**	15.5 (5.1)	0.2 (0.2)	5.3	7.1/11.4	1.7
$BAMS_a$	3.2 (4.0)	13.0 (7.2)	**63.6 (13.5)**	18.5 (5.1)	0.2 (0.2)	5.3	6.9/11.4	1.6
Optimal	0.3	17.1	**70.9**	1.9	0.1	–	–	9.8
				Scenario 6				
(p_T, p_E)	(0.02, 0.05)	(0.05, 0.08)	(0.07, 0.15)	(0.10, 0.30)	**(0.15, 0.45)**			
WT	3.0 (4.5)	4.4 (4.8)	10.7 (4.9)	19.3 (5.6)	**55.0 (9.8)**	0.0	2.6/7.7	7.6
MTA-RA	1.4 (4.4)	3.6 (5.0)	12.8 (5.9)	29.1 (6.9)	**44.6 (7.3)**	0.0	2.6/6.9	8.5
L-logit	5.7 (3.9)	7.1 (5.1)	8.3 (5.3)	24.2 (5.4)	**54.7 (10.3)**	0.0	2.8/7.7	0.0
Iso	11.8 (5.7)	14.4 (6.1)	13.3 (5.0)	18.3 (4.9)	**42.2 (8.2)**	0.0	2.5/6.7	0.0
BOIN-ET	2.7 (3.7)	3.3 (4.2)	6.7 (4.5)	22.7 (6.2)	**64.6 (11.5)**	0.0	2.9/8.2	0.0
$BAMS_f$	3.4 (3.6)	2.9 (4.1)	6.3 (4.7)	23.8 (7.0)	**63.5 (10.6)**	0.0	2.9/8.0	0.2
$BAMS_a$	4.6 (3.8)	3.2 (4.5)	6.2 (4.9)	23.9 (6.8)	**61.9 (9.9)**	0.0	2.8/7.8	0.3
Optimal	0.0	0.0	0.6	20.9	**77.9**	–	–	0.6
				Scenario 7				
(p_T, p_E)	**(0.01, 0.50)**	(0.02, 0.40)	(0.03, 0.30)	(0.04, 0.20)	(0.05, 0.10)			
WT	**61.6 (16.0)**	27.3 (8.6)	8.9 (3.9)	1.8 (1.3)	0.4 (0.3)	0.0	0.5/12.8	0.0
MTA-RA	**82.9 (15.2)**	12.4 (5.7)	3.1 (3.9)	3.0 (0.6)	0.2 (0.2)	0.0	0.6/11.8	0.3
L-logit	**36.8 (10.8)**	30.0 (8.0)	21.2 (5.7)	8.0 (3.3)	1.6 (1.6)	0.0	9.4/12.7	2.5
Iso	**55.2 (14.8)**	30.3 (9.3)	11.6 (4.1)	2.3 (1.4)	0.5 (0.4)	0.0	0.5/12.7	0.0
BOIN-ET	**51.2 (12.1)**	27.0 (5.5)	15.1 (3.4)	6.0 (3.3)	0.7 (5.7)	0.0	0.8/10.5	0.0
$BAMS_f$	**61.3 (16.7)**	27.7 (8.7)	8.6 (3.4)	2.1 (1.0)	0.5 (0.3)	0.0	0.5/13.0	0.0
$BAMS_a$	**58.4 (15.8)**	29.4 (9.3)	9.7 (3.8)	2.1 (0.8)	0.5 (0.2)	0.0	0.5/13.0	0.0
Optimal	**86.8**	12.3	0.9	0.0	0.0	–	–	0.0
				Scenario 8				
(p_T, p_E)	(0.25, 0.01)	(0.55, 0.55)	(0.67, 0.55)	(0.87, 0.55)	(0.95, 0.55)			
WT	22.7 (12.9)	11.3 (6.0)	0.1 (1.0)	0.0 (0.1)	0.0 (0.0)	7.1	7.1/3.9	65.9
MTA-RA	29.4 (10.6)	22.4 (10.9)	0.3 (3.2)	0.0 (0.6)	0.0 (0.0)	14.7	11.3/8.2	47.8
L-logit	87.8 (20.6)	4.0 (5.4)	0.1 (1.9)	0.0 (0.6)	0.0 (0.0)	7.9	9.9/4.6	8.1
Iso	87.0 (17.1)	5.1 (7.2)	0.3 (3.4)	0.0 (0.9)	0.0 (0.1)	11.6	11.4/6.5	7.6
BOIN-ET	3.0 (9.1)	14.2 (9.5)	0.1 (0.9)	0.0 (0.1)	0.0 (0.0)	10.5	8.2/5.9	82.8
$BAMS_f$	5.2 (8.6)	7.9 (8.8)	0.2 (1.1)	0.0 (0.0)	0.0 (0.0)	9.9	7.6/5.4	87.8
$BAMS_a$	5.0 (8.6)	7.5 (8.8)	0.2 (1.0)	0.0 (0.0)	0.0 (0.0)	9.8	7.7/5.4	87.3
Optimal	0.0	0.4	0.0	0.0	0.0	–	–	99.6

and BAMS_a differ in the cases where multiple doses have high efficacy response rates, e.g., in scenarios A1, A6, and A7 in the Supplementary material available at *Biostatistics* online. In these scenarios, the fixed-cutoff BAMS_f design with $\delta_\text{E} = 0.35$ tends to be trapped in sub-optimal doses more often, because several sub-optimal doses also have adequate efficacy probabilities. In contrast, under the BAMS_a design, the value of δ_E is adaptively switched to 0.5 in these scenarios, which leads to more accurate selection of the OBD. Although the issue of the BAMS_f design can be addressed by utilizing a higher cutoff, say $\delta_\text{E} = 0.5$, BAMS_f is still not flexible enough to accommodate all the cases, because the range of the efficacy probabilities is typically unknown in early-phase trials.

In Supplementary material available at *Biostatistics* online, we report the simulation results of the six designs under the 12 additional scenarios in Table S1 of the Supplementary material available at *Biostatistics* online. In summary, the proposed BAMS designs perform competitively well in terms of trial accuracy and efficiency. In most scenarios, BAMS ranks either the first or second, and never does it yield the worst performance among the six designs in comparison. From the safety perspective, the BAMS design on average assigns a smaller number of patients to overdoses, which is an additional distinctive and appealing feature. Moreover, the sensitivity analyses and simulation studies with late-onset outcomes in Supplementary material available at *Biostatistics* online further demonstrate that BAMS and TITE-BAMS are robust to different toxicity–efficacy correlations as well as parameter specifications.

4. TRIAL APPLICATION

For illustration, we apply the proposed BAMS to redesign the aforementioned phase I/II clinical trial for dose finding of acadesine in patients with CLL (Van Den Neste *and others*, 2013; ClinicalTrials.gov identifier: NCT00559624). In the trial, five doses were considered: $\{50, 83.5, 139.5, 210, 315\}$ mg/kg. Severity of adverse events or toxicity was scored according to the National Cancer Institute Common Terminology Criteria for Adverse Events Version 3.0, and adverse events with grade ≥ 3 in severity were considered as DLTs. The primary efficacy outcome was complete response (CR), and the goal was to determine the MTD and OBD of acadesine.

At such an early stage of drug development, the prior information on the dose–response curves was very limited. It would be difficult to prespecify parametric models to capture the true response profiles, while the proposed BAMS design may well serve this purpose due to its robust feature. The target toxicity probability of the MTD was set at $\phi_\text{T} = 30\%$, and the response rate of the OBD of acadesine should be greater than $\phi_\text{E} = 25\%$. We implement the BAMS_a design with design parameters $(N^*, \delta_\text{T}, \epsilon, c_\text{T}, c_\text{E}) = (12, 0.15, 0.05, 0.05, 0.10)$, which are the same as those used in Section 3. We assume a total of 30 patients with a cohort size of three to be enrolled into the study, and use hypothetical data to illustrate the dose transition of the BAMS_a design. From the first to the tenth cohort, the hypothetical data are given by $(j, m_j, y_{\text{T}j}, y_{\text{E}j}) = \{(1,3,0,0), (2,3,0,0), (3,3,0,1), (4,3,2,1), (4,6,2,2), (5,3,2,0), (4,9,2,3), (4,12,3,4), (3,6,0,2), (4,15,3,7)\}$. At the beginning of the trial, a prephase stage is initiated to accelerate dose escalation, i.e., we escalate one dose level of acadesine until the first DLT or CR is observed. On the basis of this escalation rule, dose levels 1, 2, and 3 are sequentially tested on the patients until one CR (but no DLT) is observed at dose level 3. Thus, the main dose-finding phase starts from the fourth cohort. Figure 2 shows the posterior model probabilities at each decision-making time. After treating the first three cohorts, the posterior model probabilities of $\mathcal{M}_{\text{T}k}$ are $(0.00, 0.00, 0.02, 0.24, 0.34, 0.40)$, and those of $\mathcal{M}_{\text{E}k}$ are $(0.03, 0.04, 0.22, 0.38, 0.35)$. Although dose level 5 is safe according to the toxicity data, it does not achieve the highest efficacy. Therefore, our dose-finding rule recommends dose level 4 for the next cohort. As another example, after cohort 8, a total of 12 patients have been treated at dose level 4, leading to an estimated efficacy probability of 0.33. According to the data-adaptive model specification of BAMS_a, the exploratory cutoff is increased to $\delta_\text{E} = 0.5$. As shown in Figure 2, the posterior model probabilities of $\mathcal{M}_{\text{E}k}$ are immediately updated to $(0.01, 0.04, 0.56, 0.27, 0.13)$, resulting in de-escalation

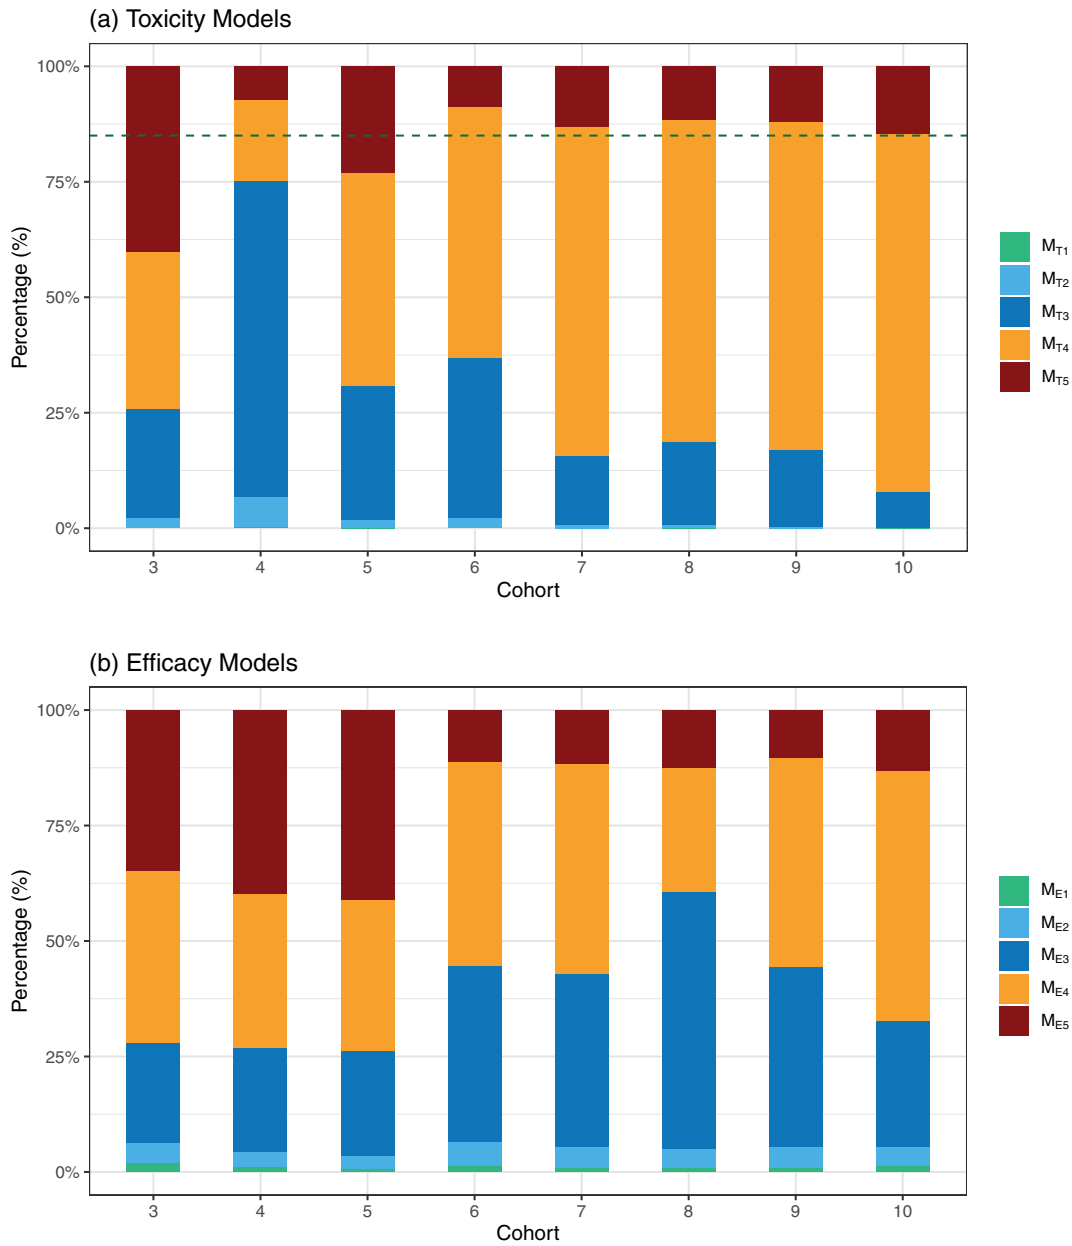

Fig. 2. Based on the hypothetical data accumulated up to cohorts 3–10 (after the prephase) in the redesigned acadesine trial, panel (a) shows the posterior probability for toxicity model $M_{Tk}, k = 1, \ldots, 5$, which specifies dose level k is the highest dose with a toxicity probability below $\phi_T = 0.30$. Because the posterior probability for M_{T0} is visually indifferent from 0, it is omitted from the plot. The dashed horizontal line corresponds to $1 - \delta_T = 0.85$. This shows that the highest admissible doses $j_{Tn}^{\max}$ are levels 5, 4, 5, 4, 4, 4, 4, 4, respectively. Panel (b) shows the posterior probability for efficacy models M_{Ek}, which specifies dose level k yields a higher efficacy probability than other doses. The most efficacious doses $j_{En}^{\max}$ are levels 4, 5, 5, 4, 4, 3, 4, 4, respectively. As the estimates of the highest admissible doses and the most efficacious doses are more meaningful when the trial has accumulated a moderate number of cohorts, only the values after the 3rd cohort (i.e., the start-up phase) are reported.

to dose level 3. Such a de-escalation decision is sensible, because dose level 3 has been tried by fewer patients yet possesses the same efficacy probability as dose level 4. Assigning more patients to promising doses but tested inadequately can increase the accuracy of OBD identification.

At the end of the trial, we have assigned 3, 3, 6, 15, and 3 patients respectively to dose levels 1, 2, 3, 4, and 5, and have observed five DLTs and ten CRs from these 30 patients. No dose is eliminated due to excessive toxicity or futility throughout the trial. The posterior model probabilities for toxicity are $(P(\mathcal{M}_{\mathrm{T0}}|D_N), \ldots, P(\mathcal{M}_{\mathrm{T5}}|D_N)) = (0.00, 0.00, 0.00, 0.08, 0.77, 0.15)$, and those for efficacy are $(P(\tilde{\mathcal{M}}_{\mathrm{E1}}|D_N), \ldots, P(\tilde{\mathcal{M}}_{\mathrm{E5}}|D_N)) = (0.01, 0.04, 0.28, 0.55, 0.13)$. According to our dose selection rule, dose level 4 is finally identified as the MTD as well as the OBD, which not only is tolerable but also has an optimal therapeutic effect. The posterior estimates of the toxicity and efficacy probabilities can be obtained via Bayesian model averaging, $(\bar{p}_{\mathrm{T1}}, \ldots, \bar{p}_{\mathrm{T5}}) = (0.01, 0.03, 0.07, 0.19, 0.58)$ and $(\bar{p}_{\mathrm{E1}}, \ldots, \bar{p}_{\mathrm{E5}}) = (0.08, 0.15, 0.33, 0.40, 0.22)$. The estimated toxicity and efficacy curves in Figure S8 of Supplementary material available at *Biostatistics* online further confirm the optimality of dose level 4.

It is worth noting that the toxicity rate at dose level 5 is estimated as $\bar{p}_{\mathrm{T5}} = 0.58$. Albeit being likely ove-toxic, it is computed based on a small sample size of three subjects only and thus associated with a high estimation uncertainty. Compared with the point estimate, the posterior model selection procedure can appropriately quantify such a high uncertainty: dose level 5 has a posterior model probability $P(\mathcal{M}_{\mathrm{T5}}|D_N) = 0.15$, the same as the decision boundary $\delta_{\mathrm{T}} = 0.15$. Although the choice of $\delta_{\mathrm{T}} = 0.15$ rules out dose level 5 in the admissible set $\widehat{\mathcal{A}}_{\mathrm{T}}(D_N)$, it should be cautious to interpret this dose as an overly toxic level. This is because a slightly larger value of δ_{T}, namely > 0.15, would retain dose level 5 back in $\widehat{\mathcal{A}}_{\mathrm{T}}(D_N)$. Based on the CRM with the skeleton probabilities (0.10, 0.20, 0.30, 0.40, 0.50), the final estimated toxicity rates are (0.02, 0.07, 0.13, 0.20, 0.30), and hence dose level 5 is deemed as acceptable. On the other hand, if dose level 5 has a sample size of 15 subjects and 10 DLTs, its posterior model probability $P(\mathcal{M}_{\mathrm{T5}}|D_N)$ would then be as low as 0.003, and surely such a dose level would be excluded from the admissible set and deemed as excessively toxic. This example further demonstrates the advantages of using posterior model probabilities in decision making, because not only does it consider the magnitude of the point estimate but it also accounts for the associated uncertainly in the estimation.

5. CONCLUDING REMARKS

We have proposed an adaptive phase I/II dose-finding design for determination of the OBD based on a Bayesian model selection framework. Our BAMS design does not require estimation of the underlying dose–response curves in contrast to the conventional methods. The BAMS design continually estimates the posterior model probabilities by incorporating all the available data. The prespecified curve-free (or nonparametric) models have meaningful interpretations of quantifying all possible profiles of the OBD. The adoption of the model selection scheme makes the BAMS design easy to implement. Due to the curve-free nature of the proposed design, it inherits the robustness property with respect to different dose–toxicity and dose–efficacy relationships. We have demonstrated the desirable performance of our design from both theoretical and numerical perspectives. Furthermore, the TITE-BAMS design can deal with late-onset responses by using the follow-up data to accelerate the trial without performance deterioration.

Upon a careful assessment of the operating characteristics of the design, it is worth emphasizing that the proposed design cannot achieve the optimal performance universally under all scenarios with a fixed set of design parameters. Specifically, when multiple (say ≥ 2) DLTs are observed among the first cohort of three patients at the optimal dose level j^{OBD} which has a toxicity probability $p_{\mathrm{T}j^{\mathrm{OBD}}} = \phi_{\mathrm{T}}$, the proposed design tends to assign relatively more patients to lower dose levels due to safety considerations. According to the proof of Theorem 2.1, as long as the OBD has not been eliminated by the dose-elimination rule, there is a positive probability to explore this "ambiguously toxic" dose in the subsequent dose assignment. However,

in some circumstances, such a turn-around probability might be relatively small under our default setting. As a result, a lower, suboptimal dose level tends to be selected more often. This issue can be resolved by elevating the value of the posterior probability cutoff δ_T, which needs to be further discussed with clinical investigators. Nonetheless, it is worth noting that such cases rarely happen in practice. For example, with $p_{T,OBD} = \phi_T = 0.3$, the probability that three (or two) DLTs are observed among the first three patients at such a dose level is only 0.027 (or 0.189). Moreover, the simulation results (e.g., scenarios A6 and A11 in Supplementary material available at *Biostatistics* online) indicate that in comparison with other designs, the proposed design still performs competitively in terms of the OBD selection percentage and patient allocation.

Supplementary Material

Supplementary material is available at Biostatistics online.

Acknowledgments

We thank the associate editor, the two referees, and the editor for their many constructive and insightful comments that have led to significant improvements in the article.

Conflict of Interest: None declared.

Funding

Lin's research was partially supported by award number 5P30CA016672 from the National Cancer Institute. Yin's research was partially supported by the Research Grants Council of Hong Kong (Grant number 17308420).

References

Braun, T. M. (2002). The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials* **23**, 240–256.

Cheung, Y. K. (2011). *Dose Finding by the Continual Reassessment Method*. Boca Raton, FL: Chapman & Hall/CRC.

Cheung, Y. K. (2014). Simple benchmark for complex dose finding studies. *Biometrics* **70**, 389–397.

Van Den Neste, E., Cazin, B., Janssens, A., Gonzlez-Barca, E., *and others* (2013). Acadesine for patients with relapsed/refractory chronic lymphocytic leukemia (CLL): a multicenter phase I/II study. *Cancer Chemotherapy and Pharmacology* **71**, 581–591.

Guo, B. and Yuan, Y. (2017) Bayesian phase I/II biomarker-based dose finding for precision medicine with molecularly targeted agents. *Journal of American Statistical Association* **112**, 508–520.

Hoering, A., LeBlanc, M., and Crowley, J. (2011). Seamless phase I–II trial design for assessing toxicity and efficacy for targeted agents. *Clinical Cancer Research* **17**, 640–646.

Li, P., Liu, R., Lin, J., Ji, Y. (2020). TEPI-2 and UBI: designs for optimal immuno-oncology and cell therapy dose finding with toxicity and efficacy. *Journal of Biopharmaceutical Statistics* **30**, 979–992.

Lin, R. and Yin, G. (2017). Nonparametric overdose control with late-onset toxicity in phase I clinical trials. *Biostatistics* **18**, 180–194.

Lin, R. and Yuan, Y. (2020). Time-to-event model-assisted designs for dose-finding trials with delayed toxicity. *Biostatistics* **21**, 807–824.

Lin, R., Zhou, Y., Yan, F., Li, D., and Yuan, Y. (2020). BOIN12: Bayesian optimal interval phase I/II trial design for utility-based dose finding in immunotherapy and targeted therapies. *JCO Precision Oncology* **4**, 1393–1402.

MANDREKAR, S. J., QIN, R., AND SARGENT, D. J. (2010). Model-based phase I designs incorporating toxicity and efficacy for single and dual agent drug combinations: methods and challenges. *Statistics in Medicine* **29**, 1077–1083.

MOZGUNOV, P. AND JAKI, T. (2019). An information theoretic phase I–II design for molecularly targeted agents that does not require an assumption of monotonicity. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **68**, 347–367.

MUENZ, D. G., TAYLOR, J. M., AND BRAUN, T. M. (2019). Phase I–II trial design for biologic agents using conditional auto-regressive models for toxicity and efficacy. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **68**, 331–345.

PAOLETTI, X. AND POSTEL-VINAY, S. (2018). Phase I–II trial designs: how early should efficacy guide the dose recommendation process? *Annals of Oncology* **29**, 540–541.

RIVIERE, M. K., YUAN, Y., JOURDAN, J. H., DUBOIS, F., AND ZOHAR, S. (2018). Phase I/II dose-finding design for molecularly targeted agent: plateau determination using adaptive randomization. *Statistical Methods in Medical Research* **27**, 466–479.

THALL, P. F. AND COOK, D. C. (2004). Dose-finding based on efficacy–toxicity trade-offs. *Biometrics* **60**, 684–693.

TAKEDA, K., TAGURI, M., AND MORITA, S. (2018). BOIN-ET: Bayesian optimal interval design for dose finding based on both efficacy and toxicity outcomes. *Pharmaceutical Statistics* **17**, 383–395.

TAKEDA, K., MORITA, S., AND TAGURI, M. (2020). TITE-BOIN-ET: time-to-event Bayesian optimal interval design to accelerate dose-finding based on both efficacy and toxicity outcomes. *Pharmaceutical Statistics* **19**, 335–349.

WAGES, N. A. AND TAIT, C. (2015). Seamless phase I/II adaptive design for oncology trials of molecularly targeted agents. *Journal of Biopharmaceutical Statistics* **25**, 903–920.

WANG, C., ROSNER, G. L., AND RODEN, R. B. (2019). A Bayesian design for phase I cancer therapeutic vaccine trials. *Statistics in Medicine* **38**, 1170–1189.

YAN, F., THALL, P. F., LU, K. H., GILBERT, M. R., AND YUAN, Y. (2018). Phase I–II clinical trial design: a state-of-the-art paradigm for dose finding. *Annals of Oncology* **29**, 694–699.

YIN, G. (2012). *Clinical Trial Design: Bayesian and Frequentist Adaptive Methods*. Hoboken, NJ, USA: Wiley.

YIN, G., LI, Y., AND JI, Y. (2006). Bayesian dose-finding in phase I/II trials using toxicity and efficacy odds ratio. *Biometrics* **62**, 777–784.

YUAN, Y. AND YIN, G. (2011). Bayesian Phase I/II adaptively randomized oncology trials with combined drugs. *Annals of Applied Statistics* **5**, 924–942.

YUAN, Y., NGUYEN, H. Q., AND THALL, P. F. (2017). *Bayesian Designs for Phase I-II Clinical Trials.* Boca Raton: CRC Press.

ZANG, Y., LEE J. J., AND YUAN, Y. (2014) Phase I dose-finding trial designs for identifying optimal biological dose for molecularly targeted agents. *Clinical Trials* **11**, 319–327.

[*Received October 15, 2020; revised May 26, 2021; accepted for publication June 6, 2021*]