Statistics in Biosciences (2024) 16:69–85
https://doi.org/10.1007/s12561-023-09379-5

A Unified Decision Framework for Phase I Dose-Finding Designs

Yunshan Duan[1] · **Shijie Yuan**[2] · **Yuan Ji**[3] · **Peter Mueller**[1]

Received: 22 July 2022 / Revised: 13 January 2023 / Accepted: 3 May 2023 /
Published online: 27 July 2023

Abstract

The purpose of a phase I dose-finding clinical trial is to investigate the toxicity profiles of various doses for a new drug and identify the maximum tolerate dose. Over the past three decades, various dose-finding designs have been proposed and discussed, including conventional model-based designs, new model-based designs using toxicity probability intervals, and rule-based designs. We present a simple decision framework that can generate several popular designs as special cases. We show that these designs share common elements under the framework, such as the same likelihood function, the use of the loss functions, and the nature of the optimal decisions as Bayes rules. They differ mostly in the choice of the prior distributions. We present theoretical results on the decision framework and its link to specific and popular designs like mTPI, BOIN, and CRM. These results provide useful insights into the similar theoretical foundations of these designs. We also show that the designs exhibit similar operating characteristics. Therefore, the choice of a design for a practical trial among the ones we reviewed may be up to the statistician's and clinician's own preference, such as preference of more model-based approach or more simple and transparent decisions.

Keywords Bayes rule · Phase I dose-finding designs · Mode-assisted designs · Decision-theoretic framework · Toxicity

✉ Yuan Ji
YJi@health.bsd.uchicago.edu

[1] Department of Statistics and Data Science, University of Texas at Austin, Austin, USA

[2] Laiya Consulting, Inc., Shanghai, China

[3] Department of Public Health Sciences, University of Chicago, Chicago, USA

1 Introduction

A phase I clinical trial is the first stage of in-human investigation of a new drug or therapy. Phase I dose-finding designs aim to identify the maximum tolerate dose (MTD) and to provide dose recommendation for later phase trials. In the vast majority of phase I trials, a set of ascending candidate doses is tested for toxicity and the dose toxicity probability is assumed to be monotonically increasing with the dose level. Typically, the MTD is defined as the highest dose with a dose limiting toxicity (DLT) probability closest to, or not higher than a target toxicity probability p_T. Usually p_T ranges from 0.17 and 0.3. In addition, some designs include the notion of an equivalence interval (EI) to allow for variations in the definition of the MTD. For example, one may choose to set $p_T = 0.3$ and $\text{EI} = (p_T - \epsilon_1, p_T + \epsilon_2) = (0.25, 0.35)$. This means that the target DLT probability of the MTD is 0.3, but doses with DLT probabilities between 0.25 and 0.35 can also be considered as the MTD. In other words, the EI allows investigators to consider doses with toxicity probabilities within the EI interval as appropriate MTD candidates.

A variety of statistical designs for phase I dose-finding trials has been discussed in the literature. A design consecutively assigns patients to recommended dose levels based on the observed DLT outcomes from previously enrolled patients. Existing designs can broadly be divided into two categories, rule-based designs and model-based designs. Among model-based designs, some use simple models and are sometimes called "model-assisted" designs. See Fig. 1 for an illustration. We provide a brief introduction of the designs in Fig. 1 next.

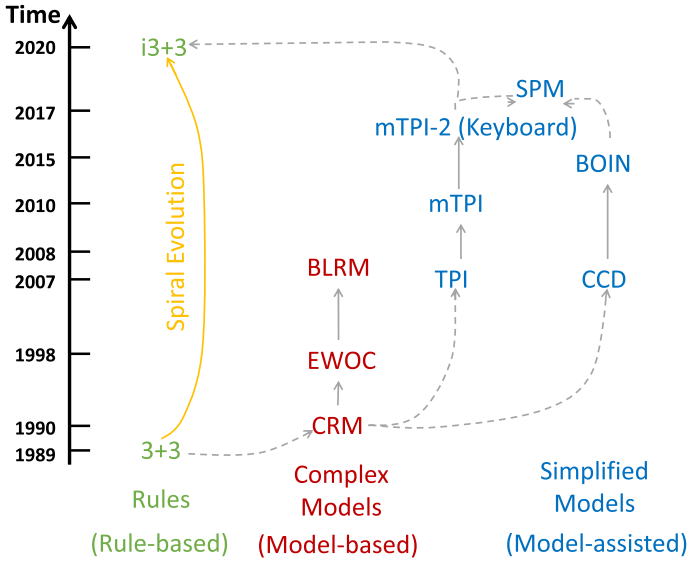

Fig. 1 Illustration of some Phase I designs. Dotted lines connect designs across different categories, and solid lines connect designs within the same category. Arrows imply chronological orders

The 3+3 design [1] is rule-based and consecutively assigns patients to the current dose or the adjacent higher or lower doses based on observed DLT outcomes. For example, if current dose at which patients are assigned is d, then 3+3 assigns the next patient cohort to doses $(d + 1)$, $(d - 1)$, or d itself. This is called the "up-and-down" rule. Based on the same up-and-down rule, a smarter rule-based design, i3+3, is proposed in Liu et al. [2] by accounting for higher sampling variability when the sample size at a dose is small. The i3+3 design maintains simplicity of rule-based designs, and exhibits operating characteristics comparable to more complex model-based/assisted designs.

The continual reassessment method (CRM) [3], as the first model-based design in the literature, is based on an inference model with a parsimoniously parameterized dose-response curve. During the trial, CRM continuously updates the estimated dose-response curve based on the observed DLT data throughout the trial. CRM has motivated subsequent important work on model-based designs, including the Bayesian logistic regression method (BLRM) in Neuenschwander et al. [4] and the escalation with over-dose control (EWOC) in Tighiouart and Rogatko [5], among many others, all leveraging parametric dose-response models for statistical inference.

Recently, a class of designs, collectively known as "interval-based designs" take advantage of the notion of an EI to simplify statistical modeling and decision making for phase I trials. Notable examples include the toxicity probability interval (TPI) design [6] and its two modifications, mTPI [7], mTPI-2 [8] (equivalently, Keyboard [9]), the cumulative cohort (CCD) design [10], and the Bayesian optimal interval (BOIN) design [11]. These designs use simple models such as the beta/binomial hierarchical model and assume independence across dose toxicity probabilities, without attempting to explicitly model a dose-response curve. While the independence model assumption is apparently not true because dose toxicity is assumed to be monotonically increasing, it does not affect the operating characteristics of the interval-based designs due to various reasons like the safety restrictions in practice (e.g., no skipping in dose escalation). As a result, interval-based designs show robust performance based on simple up-and-down rules, restricting dosing decisions to be no more than one dose-level change from the current dose used in the trial. In other words, a simple independent beta/binomial model coupled with a simple up-and-down rule leads to desirable simulation performance that justifies the application of these designs. Importantly, the interval-based designs often generate a decision table that greatly simplifies the trial conduct and allow investigators to easily execute the dosing decisions provided in the table.

In the past three decades, many designs (CRM, mTPI, mTPI-2, BOIN, etc.) have been successfully applied to real-world trials. It is natural to wonder which design or designs are suitable for a practical trial. Recent reviews [12, 13] provide some assessment of these designs, mainly from the perspective of simulation performance. Occasionally, conflicting conclusions might arise from different reviews based on the criteria used for design evaluation, or from different scenarios considered in the comparison. While simulation results can provide important information on the numerical performance of the designs, we argue that a theoretical investigation would complement the simulation results. In this article, we show that using the same optimal decision rule under the proposed decision

framework, one can generate several published designs as special cases. In other words, we show a theoretical connection across different designs. These designs include mTPI, mTPI-2, BOIN, and CCD. In addition, based on the proposed decision framework, we develop a new version of the CRM design, called Int-CRM, that is founded on the same model assumption with the original CRM design but a different decision rule. We show that Int-CRM achieves comparable simulation performance as the original CRM design and other interval-based designs. The general decision framework provides insight into the similarities and differences across various designs and may assist investigators to select the right design for their specific needs.

The remainder of the paper is structured as follows. In Sect. 2, we introduce the unified decision framework and its main components. Section 3 shows how known designs fit into this framework, including the mTPI, mTPI-2, BOIN, CCD, and Int-CRM designs. In Sect. 4, we conduct simulation studies to assess the operating characteristics of the designs using the i3+3 and CRM designs as benchmarks. Finally, we conclude and end the paper with a discussion in Sect. 5.

2 Decision Framework

2.1 Overview

We cast the problem of dose finding as an optimization in a decision problem. In particular we focus on the myopic decision problem of selecting the dose for the next patient (cohort). It is myopic because the problem does not address the global decision of stopping the trial and dose selection; instead, the problem only concerns about the local optimal decision of finding the next dose for future patients. The main components of the decision framework have been briefly illustrated in Guo et al. [8] for the mTPI and mTPI-2 designs. A decision problem is characterized by an action space for decisions a, a probability model for all unknown quantities, and a loss function [14]. Table 1 shows a summary of these components.

Table 1 Components of the proposed decision framework

Component	Notation	Notes
Probability model	$f(\boldsymbol{y} \mid \boldsymbol{\theta})\pi(\boldsymbol{\theta} \mid m)\pi(m)$	A hierarchical model with parameters $\boldsymbol{\theta}$ and m
Action	a	Up-and-down dosing decisions, including D (de-escalate), S (stay), and E (escalate)
Loss function	$\ell(a, \boldsymbol{\theta})$	The loss for taking action a where θ is the true parameter
Optimal rule	$\mathcal{R} = \arg\min_a \int \ell(a, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid \boldsymbol{y}) d\boldsymbol{\theta}$	Bayes' rule that chooses the action with the minimal posterior expected loss

2.2 General Framework

In dose-finding trials with binary DLT endpoints, the parameter of interest is a set of toxicity probabilities, $\boldsymbol{\theta} = (p_1, ..., p_T)$ at dose levels x_d, $d = 1, \dots, T$, where T is the number of dose levels, and p_d is the toxicity probability at dose d. Let y_d denote the number of patients who experience DLTs out of n_d patients treated at dose d, and let $\boldsymbol{y} = (y_1, ..., y_T)$. For all methods in the upcoming discussion the sampling model $f(\boldsymbol{y} \mid \boldsymbol{\theta})$ in Table 1 is a binomial distribution with parameter p_d, i.e.,

$$y_d \mid p_d \sim \text{ Bin } (n_d, p_d), \quad d = 1, ..., T$$

implying a likelihood function,

$$f(\boldsymbol{y} \mid \boldsymbol{\theta}) \propto \prod_{d=1}^{T} p_d^{y_d} (1 - p_d)^{n_d - y_d}$$

For model-based designs with a dose-response curve, toxicity probabilities are modeled as a function of the dose levels x_d. For example, a version of the CRM assumes $p_d = q_d^{exp(\alpha)}$, and a single parameter $\theta = \alpha$ (the values q_d are fixed and are known as the "skeleton"). The BLRM design uses $p_d = logit^{-1}(\alpha + \beta x_d)$ with parameters $\boldsymbol{\theta} = (\alpha, \beta)$.

The proposed decision framework uses a concept of probability intervals. The parameter of interest is p_d, and the parameter space of p_d is $I = [0, 1]$. Consider a set of intervals within I, denoted as $\Omega = \{I_k,\ k = 1, ..., K\}$, which form a partition of the parameter space I. That is, $\bigcup_{k=1}^{K} I_k = I$ and $I_k \cap I_{k'} = \emptyset,\ \ k \neq k'$. The true value of p_d belongs to one and only one of the intervals. For example, $\Omega = \{I_1 = [0, 0.5], I_2 = (0.5, 1)\}$ is a partition, and if $p_d = 0.3$, $p_d \in I_1$. We introduce a latent indicator m_d (or, for short, just m) with $m_d = k$ if $p_d \in I_k$, and define a hierarchical model prior $\pi(m)$ and $\pi(p_d \mid m)$. For example, $\pi(m = k) = \frac{1}{K}$, $\quad k = 1, ..., K$, and $\pi(\boldsymbol{\theta} \mid m = k) \propto \prod_{d=1}^{T} \text{ Be } (\alpha, \beta)\delta(p_d \in I_k)$, a truncated beta distribution. Here, $\delta(\cdot)$ is an indicator function. That is, p_d are conditionally independent with pdf

$$p(p_d \mid m = k) = \frac{beta(p_d; \alpha, \beta)\delta(p_d \in I_k)}{\int_{I_k} beta(p_d; \alpha, \beta) dp_d}$$

where $beta(p_d; \alpha, \beta) = \frac{p(\alpha+\beta)}{p(\alpha)p(\beta)} p_d^{\alpha-1}(1 - p_d)^{\beta-1}, \alpha > 0, \beta > 0$ is a Be (α, β) p.d.f.

We consider a special partition $\Omega = \{I_1, I_2, I_3\}$ where $I_2 \triangleq I_S = EI = (p_T - \epsilon_1, p_T + \epsilon_2)$, $I_1 \triangleq I_E = [0, p_T - \epsilon_1]$, and $I_3 \triangleq I_D = [p_T + \epsilon_2, 1]$. Therefore, $K = 3$ and we use notations I_S, I_E, and I_D to associate the intervals with corresponding up-and-down dose-finding decisions S, E, and D, respectively. We summarize the proposed decision framework below.

Likelihood:

$$f(\boldsymbol{y} \mid \boldsymbol{\theta}) \propto \prod_{d=1}^{D} p_d^{y_d} (1 - p_d)^{n_d - y_d} \tag{1}$$

where p_d is the toxicity probability for dose d, $d = 1, \dots, D$.

Prior: We assume p_d are *a priori* independent and

$$\pi(p_d \mid m = k) \propto g(p_d)\delta(p_d \in I_k), \quad k = 1, ..., K,$$
$$\pi(m = k) = \frac{1}{K}, \quad k = 1, ..., K.$$

For example, $g(p_d) = beta(p_d;\alpha, \beta)$.

Partition: $\Omega = \{I_E(I_1),\ I_S(I_2),\ I_D(I_3)\}$, where $I_E = [0, p_T - \epsilon_1], I_S = (p_T - \epsilon_1, p_T + \epsilon_2)$, and $I_D = [p_T + \epsilon_2, 1]$, and $I_1 = I_E$, $I_2 = I_S$, and $I_3 = I_D$.

Actions:

The actions are the three up-and-down decisions for dose-finding, i.e.,

$$a \in A = \{E, S, D\},$$

where A denotes the action space. Here E, S, D denote the dosing decisions "Escalation", "Stay", and "De-escalation," respectively. In particular, if the last patient was assigned dose d, then E, S, or D means treating future patients at dose $(d + 1)$, d, or $(d - 1)$, respectively.

Loss: We proceed with a myopic perspective, focusing on the decision for the respective next patient (cohort), and therefore specify a loss function for the next dose assignment a only.

We use a 0–1 loss function,

$$\ell(a, p_d) = \begin{cases} 1, & p_d \notin I_a \\ 0, & p_d \in I_a \end{cases}, \quad a \in A = \{E, S, D\}. \tag{2}$$

In words, when the action corresponds to an interval which contains the true parameter, the loss takes the value 0; otherwise, the loss equals 1. The loss function $\ell(a, p_d)$ is stated in Table 2.

In other words, the loss function $\ell(a, p_d)$ defines a 0–1 estimation loss for m, i.e., the interval that contains p_d.

Two more comments about the loss function and the setup of the decision problem. First, in general a loss (or, equivalently, utility) function could also be an argument of the outcome y_d. This is relevant, for example, if instead of inference loss we focus on the patients' preferences. However, the intention of this discussion is only to highlight common structure in existing dose finding methods, for which we only need this restricted inference loss. Another important limitation is the myopic nature of the setup. We consider the dose allocation for each patient (or patient cohort) in

Table 2 The 0–1 loss function $\ell(a, \theta)$

$\ell(a, p_d)$		$p_d \in$		
		$[0, p_T - \epsilon_1]$	$(p_T - \epsilon_1, p_T + \epsilon_2)$	$[p_T + \epsilon_2, 1]$
$a =$	D	1	1	0
	S	1	0	1
	E	0	1	1

isolation, ignoring that dose allocation now might help later decisions. That is, we ignore the sequential nature of the problem. Again, for the upcoming exposition of common underlying structure for the considered dose finding methods we will only refer to this myopic decision problem.

Bayes' rule: The optimal decision rule for dose d is the Bayes' rule, defined as

$$\mathcal{R}_d \triangleq \arg\min_{a\in A} \int \ell(a, p_d) p(p_d \mid \boldsymbol{y}) dp_d, \tag{3}$$

which minimizes the posterior expected loss. Here, $p(p_d \mid \boldsymbol{y})$ is the posterior distribution of p_d.

In general, under a 0–1 estimation loss for a discrete parameter the Bayes rule is simply the posterior mode. The following result states this in the context of our problem. The Bayes' rule is equivalent to the result of finding the interval with the maximal posterior probability.

Proposition 1 *Denote* $\Omega \triangleq \{I_1, I_2, I_3\} = \{I_E, I_S, I_D\}$, *where* $I_1 = I_E$, $I_2 = I_S$, $I_3 = I_D$. *Suppose dose d is the current dose. Let* $\{m = k\}$ *be an equivalent event to* $\{p_d \in I_k\}$, $k \in \{1, 2, 3\}$. *Let* $A = \{E, S, D\}$. *Assume* $\pi(m = k) = \frac{1}{3}$, $k \in \{1, 2, 3\}$. *The Bayes' rule under the* 0–1 *loss in Eq.* (3) *is given by*

$$\mathcal{R}_d = \arg\max_{a\in A} Pr(p_d \in I_a \mid \boldsymbol{y}) = \arg\max_{k\in\{1,2,3\}} Pr(m = k \mid \boldsymbol{y}) \tag{4}$$

See Appendix D in Supplementary for a proof.

3 Design Examples

We show how various designs fit as special cases into this framework. That is, we provide examples of the decision framework that give rise to well-known designs including mTPI, BOIN, CCD, mTPI-2, and a new version of CRM, called the Int-CRM design.

3.1 Interval-Based Designs

We first introduce the connection between the decision framework and the interval-based designs, mTPI, mTPI-2, BOIN and CCD. These designs share some common components under the framework, but also include some elements specific to each design.

Common components: Likelihood, Prior $\pi(m)$, Loss function, and the nature of the defined dose allocation as Bayes' rule.

Individual components: Prior $\pi(p_d \mid m)$, the specific partition $\Omega = \{I_k,\ k = 1, ..., K\}$, and the definition of the action set A.

All four interval-based designs use the binomial sampling model. And the designs share the same discrete uniform prior $\pi(m)$, the 0–1 loss function and the use of Bayes' rule to select a decision. They divide the [0, 1] parameter space of

p_d into different intervals and use different priors. See Table 3 as a summary. We discuss details for each design next.

3.1.1 The mTPI Design

For the mTPI design, given the equivalence interval EI= $(p_T - \epsilon_1, p_T + \epsilon_2)$, the [0, 1] parameter space is naturally partitioned into three intervals $\Omega = \{I_1 = I_E = [0, p_T - \epsilon_1], I_2 = I_S = (p_T - \epsilon_1, p_T + \epsilon_2), I_3 = I_D = [p_T + \epsilon_2, 1]\}$ that correspond to the actions $A = \{E, S, D\}$, as shown in Table 3. The mTPI decision is equivalent to the Bayes' rule $\mathcal{R}_d$ under the decision framework. See Corollary 1 below for a formal mathematical description.

Corollary 1 *The mTPI decision in Ji et al.* [7] *is given by*

$$R_{mTPI} = \arg \max_{a \in \{E,S,D\}} \text{UPM}\ (I_a),$$

where UPM stands for "unit probability mass" and $\text{UPM}\ (I_a) = Pr^*(p_d \in I_a)/||I_a||$; *here* $||I_a||$ *is the length of* I_a, *and*

$$Pr^*(p_d \in I_a) = \int B(y_d + 1, n_d - y_d + 1) \cdot p_d^{y_d}(1 - p_d)^{n_d - y_d} \cdot \delta(p_d \in I_a) dp_d$$

is calculated based on $p_d \sim \text{Be}\ (y_d + 1, n_d - y_d + 1)$. *Let* $I_1 = I_E = [0, p_T - \epsilon_1]$, $I_2 = I_S = (p_T - \epsilon_1, p_T + \epsilon_2)$, *and* $I_3 = I_D = [p_T + \epsilon_2, 1]$. *Then* $R_{mTPI} = \mathcal{R}_d$, *the Bayes rule under*

$$\pi(p_d \mid m = k) = C_k \cdot beta(p_d; 1, 1)\delta(p_d \in I_k),$$

where $beta(\cdot; 1, 1)$ *denotes the density function of* $\text{Be}\ (1, 1)$ *distribution and* $C_k = \frac{1}{\int_{I_k} beta(p;1,1)dp}$ *is a normalizing constant.*

See Appendix D in Supplementary for a proof.

Table 3 Individual components of the proposed decision framework for some interval-based designs

	mTPI	**mTPI-2**	**BOIN/CCD**
Actions	$A = \{E, S, D\}$	$A = \{1, ..., K\}$	$A = \{E, S, D\}$
Intervals	$I_E = [0, p_T - \epsilon_1]$ $I_S = (p_T - \epsilon_1, p_T + \epsilon_2)$ $I_D = [p_T + \epsilon_2, 1]$	$I_E = I_{E,1} \cup \cdots \cup I_{E,K_1}$ $I_S = (p_T - \epsilon_1, p_T + \epsilon_2)$ $I_D = I_{D,1} \cup \cdots \cup I_{D,K_2}$	$I_E = [0, \phi_E]$ $I_S = (\phi_E, \phi_D)$ $I_D = [\phi_D, 1]$
Priors	$\pi(p_d \mid m = k) \propto$ $Be(1, 1)\delta(p_d \in I_k)$	$\pi(p_d \mid m = k) \propto$ $Be(1, 1)\delta(p_d \in I_k)$	$\pi(p_d \mid m = k) =$ $\delta(p_d = \phi_k)^*$

*See Theorem 1 for details

3.1.2 The mTPI-2 Design

Ockham's razor is a principle in statistical inference calling for an explanation of the facts to be no more complicated than necessary (Thorburn [15]; Jefferys and Berger [16]). In the context of model selection, the Ockham's razor prefers parsimonious models that describe the data equally well as more complex models. In the proposed decision framework, $\{m = k\}$, $k = 1, ..., K$, is equivalent to K models $\{M_k : m = k\}$, and choosing the value of m is equivalent to a model selection problem. Bayesian model selection chooses the model with the largest posterior probability (compare Proposition 1), i.e., $Pr(m = k \mid y)$, and models are automatically penalized for their complexity. In other words, Bayes' rule $\mathcal{R}_d = \arg\max_{k\in\{1,2,3\}} Pr(m = k \mid y)$ implements Ockham's razor if we define model complexity as $||I_k||$. Therefore, when the three models are $I_1 = I_E = [0, p_T - \epsilon_1]$, $I_2 = I_S = (p_T - \epsilon_1, p_T + \epsilon_2)$, $I_3 = I_D = [p_T + \epsilon_2]$, the "simplest" model is $I_2 = I_S$, since ϵ_1 and ϵ_2 are typically small probabilities (≤ 0.05).

Guo et al. [8] explain mTPI-2 as aiming to blunt Ockham's razor by redefining a (finer) partition $\Omega^* = \{I_E^*, I_S = EI, I_D^*\}$, where $I_E^* = \{I_{E,1}, ..., I_{E,K_1}\}$ and $I_D^* = \{I_{D,1}, ..., I_{D,K_2}\}$. Probability intervals $\{I_{E,k}\}_{k=2}^{K_1}$ and $\{I_{D,k}\}_{k=2}^{K_2}$ have the same length as $I_S = EI$. Let $K = K_1 + K_2 + 1$. The selected model m under the mTPI-2 design can then be shown to be Bayes' rule $\mathcal{R}_d$, under an action set $A_m = \{1, ..., K\}$. Corollary 2 next summarizes the results.

Corollary 2 *Under mTPI-2,* $\Omega^* = \{I_E^* = \{I_{E,1}, \cdots, I_{E,K_1}\}, I_S = (p_T - \epsilon_1, p_T + \epsilon_2), I_D^* = \{I_{D,1}, \cdots, I_{D,K_2}\}\}$, $A_m = \{1, ..., K\}$. *Assume the prior on* p_d *is conditionally independent and given by*

$$\pi(p_d \mid m = k) = C_k \cdot beta(p_d; 1, 1)\delta(p_d \in I_k).$$

Then Bayes' rule is

$$\mathcal{R}_d = \arg\max_{m\in A_m} Pr(m = k \mid y) = \arg\max_{m\in A_m} UPM(I_m).$$

See Appendix D in Supplementary for a proof. Corollary 2 establishes the Bayes' rule $\mathcal{R}_d$ as an action in $A_m = \{1, ..., K\}$. To see the connection to the dose-finding decision in mTPI-2, we refer to the next result.

Corollary 3 *Let*

$$\mathcal{R}_{\text{mTPI-2}} = \begin{cases} E, & \arg\max_{I_m} \text{UPM}(I_m) \subset (0, p_T - \epsilon_1), \\ S, & \arg\max_{I_m} \text{UPM}(I_m) = (p_T - \epsilon_1, p_T + \epsilon_2), \\ D, & \arg\max_{I_m} \text{UPM}(I_m) \subset (p_T + \epsilon_2, 1). \end{cases}$$

Then

$$\mathcal{R}_{\text{mTPI-2}} = \begin{cases} E, \ I_{\mathcal{R}_d} \subset (0, p_T - \epsilon_1), \\ S, \ I_{\mathcal{R}_d} = (p_T - \epsilon_1, p_T + \epsilon_2), \\ D, \ I_{\mathcal{R}_d} \subset (p_T + \epsilon_2, 1). \end{cases}$$

In other words, if $\mathcal{R}_d = m^*$, then I_{m^*} is the interval with the largest UPM, for $m^* \in A_m$. And if I_{M^*} is below, equal to, or above the EI= $(p_T - \epsilon_1, p_T + \epsilon_2)$, the decision is E, S, or D, respectively. This is the same as the up-and-down rule in the mTPI-2 design in Guo et al. [8]. Proof of Corollary 3 is immediate and omitted.

3.1.3 The BOIN Design

The BOIN design (Liu and Yuan [11]) uses a decision rule

$$R_{BOIN} = \begin{cases} E, \ \hat{p}_d \le \lambda_1, \\ S, \ \lambda_1 < \hat{p}_d < \lambda_2, \\ D, \ \hat{p}_d \ge \lambda_2, \end{cases}$$

where $\hat{p}_d = \frac{y_d}{n_d}$, $\lambda_1 = \xi(\phi_E;\phi_S)$, $\lambda_2 = \xi(\phi_D;\phi_S)$, and

$$\xi(\phi_i;\phi_j) = \frac{log\left(\frac{1-\phi_i}{1-\phi_j}\right)}{log\left(\frac{\phi_j(1-\phi_i)}{\phi_i(1-\phi_j)}\right)}. \tag{5}$$

In particular, $\phi_S = p_T$, $\phi_E (< p_T)$ and $\phi_D (> p_T)$ are pre-specified values. Here, ϕ_E and ϕ_D play a similar rule as $(p_T - \epsilon_1)$ and $(p_T + \epsilon_2)$ in the mTPI and mTPI-2 designs, which defines the boundaries of an initial equivalence interval elicited from clinicians. We show that the decision rule R_{BOIN} is also a Bayes' rule under the proposed decision framework next.

Theorem 1 *Assume* $y_d \mid p_d \sim \text{Bin}(n_d, p_d)$, $\pi(p_d \mid m = k) = \delta(p_d = \phi_k)$, *for* $k = 1, 2, 3$, *and* $\phi_1 = \phi_E$, $\phi_2 = \phi_S = p_T$, *and* $\phi_3 = \phi_D$. *Under the* 0–1 *loss* $\ell(a, p_d)$ *in Eq.* (2), *the Bayes' rule is equivalent to* R_{BOIN}, *i.e.*,

$$\mathcal{R}_d = R_{BOIN} = \begin{cases} E, \ \hat{p}_d \le \lambda_1, \\ S, \ \lambda_1 < \hat{p}_d < \lambda_2, \\ D, \ \hat{p}_d \ge \lambda_2, \end{cases}$$

where $\hat{p}_d = \frac{y_d}{n_d}$, $\lambda_1 = \xi(\phi_E;\phi_S)$, $\lambda_2 = \xi(\phi_D;\phi_S)$, *and* ξ *is defined as in Eq.* (5)

See Appendix D in Supplementary for a proof. By Theorem 1 the BOIN design takes the form of the Bayes' rule under the same decision framework using the 0–1 loss. BOIN uses a point-mass prior for p_d on three values, ϕ_E, ϕ_S, ϕ_D, while mTPI/mTPI-2 using truncated beta priors instead. Next, we show that the BOIN design is almost the same as the CCD design. This is easiest seen under the perspective of the

proposed decision framework. The difference between the two designs are the locations of the point-mass priors.

3.1.4 The CCD design

The CCD design compares $\hat{p}_d$ with $(p_T - \epsilon_1)$ and $(p_T + \epsilon_2)$, and uses the following up-and-down rule,

$$\mathcal{R}_{\text{CCD}} = \begin{cases} E & \hat{p}_d \le p_T - \epsilon_1 \\ S & p_T - \epsilon_1 < \hat{p}_d < p_T + \epsilon_2 \\ D & \hat{p}_d \ge p_T + \epsilon_2 \end{cases}$$

Corollary 4 shows that the decision of the CCD design is the same Bayes' rule in the same framework as BOIN but with a different prior distribution.

Corollary 4 *The CCD decision* $\mathcal{R}_{\text{CCD}} = \mathcal{R}_d$ *with*

$$\pi(p_d \mid m = k) = \delta(p_d = \phi'_k), \quad k = E, S, D,$$

where $\phi'_E = \xi^{-1}(p_T - \epsilon_1)$, $\phi'_D = \xi^{-1}(p_T + \epsilon_2)$, $\phi'_S = \phi_S = p_T$, *and* $\xi(\phi) \equiv \xi(\phi, p_T)$ *in Eq.* (5).

See Appendix D in Supplementary for a proof.

Corollary 4 shows that BOIN and CCD are identical designs with the only difference being that BOIN uses a point-mass prior $\pi(p_d \mid m = k) = \delta(p_d = \phi_k)$, whereas CCD uses $\pi(p_d \mid m = k) = \delta(p_d = \phi'_k)$.

3.2 The Int-CRM Design

Using the same decision framework, we propose a variation of the CRM design, called Int-CRM. We assume the same parametric dose-response model as in the CRM design [3], with the probability of toxicity monotonically increasing with dose. Let d_i denote the dose for the ith patient, $d_i \in \{1, ..., T\}$, and Y_i the binary indicator of DLT. The dose-response curve is assumed to be the power model as in the CRM,

$$F(d, \theta) = q_d^{exp(\theta)}$$

where $(q_1, .., q_T)$ ("skeleton") are *a priori* pre-specified dose toxicity probabilities. Other sensible dose-response models, such as a logit model, may be considered as well. The toxicity rates are dependent across doses through the dose-response curve and the inference is based on the parameter θ. The likelihood function is given by

$$f(\mathbf{y} \mid \theta) \propto \prod_{i=1}^{n} F(d_i, \theta)^{y_i} \{1 - F(d_i, \theta)\}^{1 - y_i}$$

where n is the number of patients in the trial.

Following Cheung and Chappell [17], we define an interval $[A_1, A_{T+1}]$ for θ that is wide enough to allow for a wide range of dose-response curves. For example, set A_1 and A_{T+1} so that $q_1^{exp(A_1)} > 1 - 10^{-5}$ and $q_T^{exp(A_{T+1})} < 10^{-5}$, which correspond to response curves constantly equal to 1.0 and 0.0, respectively, and $\theta \in [A_1, A_{T+1}]$ allows choices in-between these extremes. Using A_1 and A_{T+1}, we define sub-intervals for θ as the set of values that imply d_k having toxicity probability closest to p_T,

$$I_k = \{\theta \in [A_1, A_{T+1}] : |F(k, \theta) - p_T| < |F(d, \theta) - p_T|, \forall d \neq k\}, \quad k = 1, \ldots, T$$

As shown in Cheung and Chappell [17], I_k is an interval, denoted as $I_1 = [\psi_1 = A_1, \psi_2), \quad I_k = [\psi_k, \psi_{k+1}),\ k = 2, ..., (T-1), \quad I_T = [\psi_T, \psi_{T+1} = A_{T+1}]$, where ψ_k is implicitly defined as the solution of

$$F(k-1, \psi_k) + F(k, \psi_k) = 2p_T, \quad k = 2, ..., T.$$

Given the "skeleton" $(q_1, ..., q_T)$, we can obtain the numerical result of the interval boundaries ϕ_k's by solving the equation above. See Appendix B in Supplementary for details. Each interval consists of a set of θ values where dose k's toxicity probability is the closest to p_T among all the doses. We use these intervals I_k's in our framework for Int-CRM. We propose hierarchical priors

$$\pi(m = k) = \frac{1}{T}, \quad k = 1, ..., T$$

and

$$\pi(\theta \mid m = k) = \frac{\phi(\theta)\delta(\theta \in I_k)}{\int_{I_k} \phi(\theta) d\theta}$$

where $\phi(\theta)$ is the density function of the normal distribution $N(0, \sigma^2)$.

The action space of the Int-CRM design is $A = \{1, ..., T\}$, corresponding to the dose for treating the next patient. Following the proposed decision framework, we use the 0–1 loss function and the Bayes' rule that minimizes the posterior expected loss for the Int-CRM decision.

Theorem 2 *Under the* 0–1 *loss*, *i.e.*,

$$\ell(a, \theta) = \begin{cases} 1, & \theta \notin I_a \\ 0, & \theta \in I_a \end{cases}, \quad a \in A = \{1, ..., T\}$$

the Int-CRM decision is the Bayes' rule

$$\begin{aligned} \mathcal{R}_{Int-CRM} &= \arg\max_{k \in A} Pr(m = k \mid y) = \arg\max_{k \in A} \int p(y \mid \theta)\pi(\theta \mid m = k) d\theta \\ &= \arg\max_{k \in A} \int \prod_{i=1}^{n} F(d_i, \theta)^{y_i} \{1 - F(d_i, \theta)^{1-y_i}\} \frac{\phi(\theta)\delta(\theta \in I_k)}{\int \phi(\theta)\delta(\theta \in I_k) d\theta} d\theta. \end{aligned}$$

The proof is immediate by the definition of Bayes' rule. Below is the proposed Int-CRM dose-finding algorithm.

The Int-CRM Algorithm:

Dose Finding Rules: After each cohort of patients completes the DLT follow-up period, the dose to be assigned is the $\mathcal{R}_{\text{Int-CRM}}$, the Bayes' rule, unless the following safety rules apply.

Safety Rules: Four additional rules are applied for safety.

Rule 1: Dose Exclusion: If the current dose is considered excessively toxic, i.e., $Prob\{p_d > p_T \mid data\} > \xi$ (see below about evaluating this probability), where the threshold ξ is close to 1, say 0.95, the current and all higher doses will be excluded in the remainder of the trial to avoid assigning any patients to those doses.

Rule 2: Early Stopping: If the current dose is the lowest dose (first dose) and is considered excessively toxic, i.e., $p\{p_1 > p_T \mid data\} > \xi$, where the threshold ξ is close to 1, say 0.95, stop the trial early and declare no MTD.

To evaluate $p\{p_d > p_T \mid data\}$ in Rules 1 and 2 we use a $\text{Be}(\alpha_0 + y_d, \beta_0 + n_d - y_d)$ distribution with $\alpha_0 = \beta_0 = 1$.

Rule 3: No-Skipping Escalation: If the dose-finding rule recommends escalation, such escalation shall not increase the dose by more than one level. Dose-escalation cannot increase by more than one level. That is, suppose the current dose is d. If the next recommended dose $\mathcal{R}_{\text{Int-CRM}}$ is such that $(\mathcal{R}_{\text{Int-CRM}} - d) > 1$, escalate to dose $(d+1)$ instead.

Rule 4: Coherence: No escalation is permitted if the empirical rate of DLT for the most recent cohort is higher than p_T, according to the coherence principle [18].

Trial Termination: The trial proceeds unless any of the following stopping criteria is met:

- If the pre-specified maximum total sample size n is reached.
- Rule 2 above.

MTD Selection: Once all the enrolled patients complete the DLT observation and the trial is not stopped early, the last dose level $\mathcal{R}_{\text{Int-CRM}}$ is selected as the MTD.

4 Simulation Studies

4.1 Simulation Settings

We set up simulation studies to evaluate the operating characteristics of the different designs that we have shown to be special cases of the proposed general framework,

including the mTPI, mTPI-2, BOIN, CCD and the Int-CRM designs. We aim to show the similarity of designs' performance since they are based on the same decision framework. We also compare to the i3+3 design and the original CRM design as benchmarks.

4.1.1 Fixed Scenarios

We use a total of 15 scenarios, with a set of $T = 4, 5$ or 6 doses. Assume the target toxicity probability $p_T = \phi_S = 0.3$ (ϕ_S is the notation in BOIN), and maximum sample size is 30. For all designs we apply the same safety rules as in the mTPI, mTPI-2, and Int-CRM designs. See Appendix A in Supplementary for details. For interval-based designs we use EI $= (p_T - \epsilon_1, p_T + \epsilon_2)$, $\epsilon_1 = \epsilon_2 = 0.05$. For the Int-CRM and CRM design, the skeleton q_d is generated using the approach proposed in Lee and Cheung [19], which selects the skeleton based on indifference intervals for the MTD. Also, we set the half width of the indifference intervals, $\delta = 0.05$. The coherence principle [18] is applied, avoiding immediate escalation after toxic outcomes.

For the BOIN design, we set $\lambda_1 = p_T - \epsilon_1$, $\lambda_2 = p_T + \epsilon_2$. This is equivalent to setting $\phi_E = \xi^{-1}(p_T - \epsilon_1)$, and $\phi_D = \xi^{-1}(p_T + \epsilon_2)$. By Theorem 1, these values for λ_1 and λ_2 make the BOIN decision identical to the CCD design, leading to same operating characteristics of the two designs.

4.1.2 Random Scenarios

We generate additional 1000 random scenarios to further evaluate the designs. Scenarios are generated based on the pseudo-uniform algorithm in Clertant et al. [20]. Figure 2 plots the first 20 scenarios. Other settings of the designs are the same as the fixed scenarios, such as p_T and EI, λ_1, λ_2 for the BOIN design, and δ for the Int-CRM and CRM designs.

4.2 Simulation Results

We evaluate the performance of the phase I designs through a few metrics, based on their ability to identify the MTD and the safety in dose selection and patient allocation. Table 4 summarizes the means and standard deviations of key performance metrics for the simulation with 1,000 scenarios. All designs show remarkable similarity with the largest mean difference across designs only about 0.02. This highlights the underlying connection of these designs and echoes our findings based on the unified decision framework that can generate most designs as special cases. Table 4 and Tables 5 and 6 in Appendix C present the simulation results of the 15 fixed scenarios.

In general, the five designs tested in the simulation studies exhibit remarkably similar performances. Specifically, they show comparable probabilities (across repeated simulation) of allocating patients to the true MTD, and similar risk of allocating patients to overly toxic doses. The BOIN/CCD and Int-CRM designs

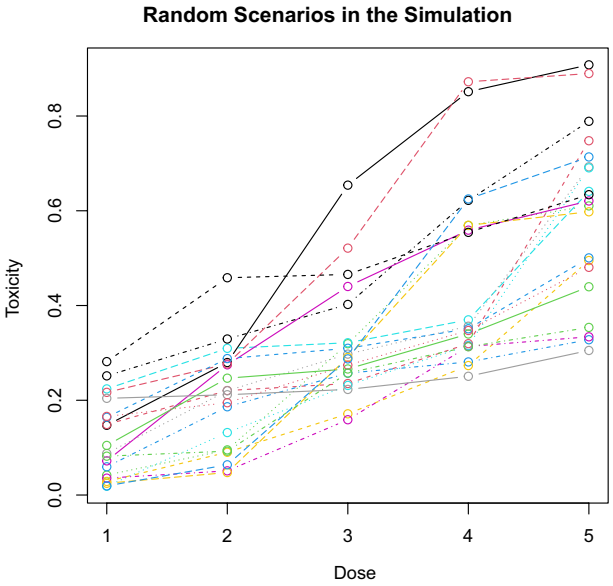

Fig. 2 Illustration of the first 20 random scenarios of true toxicity probabilities in the simulation

Table 4 Simulation results of 1000 random scenarios. Entries are the $\boldsymbol{mean}_{sd}$ (across the 1000 scenarios) proportion of simulated trials for each design and metric. For example, the first entry 0.60 means 60% of the trials correctly select the true MTD under the mTPI design. The last two metrics Tox. and None Sel. are the percentage of patients experienced DLT in all simulated trials and the proportion of simulated trials in which no dose is selected as the MTD, respectively

Metrics	mTPI	mTPI-2	BOIN/CCD	iCRM	CRM	i3+3
Correct Sel. of MTD	$0.60_{(0.15)}$	$0.62_{(0.14)}$	$0.62_{(0.14)}$	$0.62_{(0.14)}$	$0.62_{(0.14)}$	$0.62_{(0.14)}$
Sel. over MTD	$0.10_{(0.11)}$	$0.11_{(0.11)}$	$0.12_{(0.12)}$	$0.12_{(0.12)}$	$0.11_{(0.11)}$	$0.11_{(0.11)}$
Pat. at MTD	$0.50_{(0.22)}$	$0.50_{(0.21)}$	$0.50_{(0.20)}$	$0.51_{(0.20)}$	$0.51_{(0.21)}$	$0.50_{(0.21)}$
Pat. over MTD	$0.11_{(0.10)}$	$0.11_{(0.10)}$	$0.12_{(0.11)}$	$0.12_{(0.11)}$	$0.12_{(0.10)}$	$0.11_{(0.10)}$
Tox	$0.26_{(0.05)}$	$0.26_{(0.05)}$	$0.26_{(0.05)}$	$0.26_{(0.05)}$	$0.26_{(0.05)}$	$0.26_{(0.05)}$
None Sel	$0.04_{(0.07)}$	$0.04_{(0.07)}$	$0.04_{(0.07)}$	$0.04_{(0.07)}$	$0.04_{(0.07)}$	$0.04_{(0.07)}$

yield slightly higher PCS (probability of correct selection of MTD) in some cases, such as scenarios 1 and 4 in Table 5. However, they also report a higher risk in selecting doses beyond the true MTD. For example, in scenarios 2, 3, and 5 in Table 5, the probabilities of over-dosing selection under BOIN, CCD and Int-CRM are higher compared to the other designs. However, these differences are small compared to the reported standard deviations and therefore could be due to random noise and the arbitrarily generated scenarios. In summary, all designs exhibit remarkable operating characteristics in our simulations.

5 Discussion

We have developed a general decision framework for phase I dose-finding designs. We have shown that interval-based designs, like mTPI, mTPI-2, BOIN, CCD, and the model-based design Int-CRM fit into this unified framework.

All designs use the same 0–1 loss function, and all interval designs assume a binomial likelihood function. The prior construction for some designs involves the notion of candidate models. Candidate models are specified assuming different toxic profiles for the doses. Given the model, the mTPI, mTPI-2 and Int-CRM assume continuous prior distributions, using beta or normal distributions truncated to the limited parameter space implied by the given model. The BOIN and CCD designs use a different approach with a discrete prior on p_d, supported at three distinct values. Choosing those atoms is challenging and may be difficult to interpret. However, through numerous simulations conducted and published in the literature, the BOIN and CCD designs perform very well in Phase I trials with relatively small sample size.

Additionally, different loss functions can be considered in the proposed framework penalize undesirable actions and outcomes. For examples, the loss for mistakenly making an escalation decision may be larger than for a wrong de-escalation. However, such loses usually lead to more complex and less interpretable decision rules.

It is demonstrated that the designs in this paper perform similarly with comparable reliability and safety. The i3 + 3 rule-based design is not a part of this framework, but also generates similar operating characteristics, comparable with the other designs. The i3 + 3 design shares a practically important feature with interval based designs. One can pre-tabulate decision tables, which is a critical feature for the implementation in actual trials. Clinicians can choose a desirable design for phase I clinical trial based on their preference, including the model-based design CRM, the interval designs, mTPI, mTPI-2, BOIN and CCD, and the rule-based design i3+3.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s12561-023-09379-5.

Acknowledgements Peter Müller's research is partly supported by NSF Grant DMS-1952679; Yuan Ji's research is partly supported by NSF Grant DMS-1953340.

Declarations

Conflicts of interest Yuan Ji is Co-Founder of Bayesoft Inc., a statistical company providing consultation and software service for pharmaceutical and biotech industry. He is an IDMC member for Astellas, Boehringer Ingelheim, Lyell, and has research contracts with Pfizer and Sanofi.

References

1. Storer Barry E (1989) Design and analysis of phase I clinical trials. Biometrics 1989:925–937

2. Liu M, Wang SJ, Ji Y (2020) The i3+ 3 design for phase I clinical trials. J Biopharm Stat 30(2):294–304
3. John O, Margaret P, Lloyd F (1990) Continual reassessment method: a practical design for phase 1 clinical trials in cancer. Biometrics 1990:33–48
4. Neuenschwander B, Branson M, Gsponer T (2008) Critical aspects of the Bayesian approach to phase I cancer trials. Stat Med 27(13):2420–2439
5. Mourad T, André R (2010) Dose finding with escalation with overdose control (EWOC) in cancer clinical trials. Stat Sci 2010:217–226
6. Yuan J, Yisheng L, Nebiyou BB (2007) Dose-finding in phase I clinical trials based on toxicity probability intervals. Clin Trial 4(3):235–244
7. Yuan J, Ping L, Yisheng L, Nebiyou BB (2010) A modified toxicity probability interval method for dose-finding trials. Clin Trial 7(6):653–663
8. Guo W, Wang SJ, Yang S, Lynn H, Ji Y (2017) A Bayesian interval dose-finding design addressingOckham's razor: mTPI-2. Contemp Clin Trial 58:23–33
9. Fangrong Y, Mandrekar Sumithra J, Ying Y (2017) Keyboard: a novel Bayesian toxicity probability interval design for phase I clinical trials. Clin Cancer Res 23(15):3994–4003
10. Ivanova A, Flournoy N, Chung Y (2007) Cumulative cohort design for dose-finding. J Stat Plan Inference 137(7):2316–2327
11. Suyu L, Ying Y (2015) Bayesian optimal interval designs for phase I clinical trials. J Royal Stat Soc Series C Appl Stat 2015:507–523
12. Ying Y, Hess Kenneth R, Hilsenbeck Susan G, Gilbert Mark R (2016) Bayesian optimal interval design: a simple and well-performing design for phase i oncology trials. Clin Cancer Res 22(17):4291–4301
13. Jablonski HB, Wages Nolan A, Conaway Mark R (2017) Performance of toxicity probability interval based designs in contrast to the continual reassessment method. Stat Med 36(2):291–300
14. Berger James O (2013) Statistical decision theory and Bayesian analysis. Springer Science and Business Media, Cham
15. Thorburn William M (1918) The myth of Occam's razor. Mind 27(107):345–353
16. Jefferys William H, Berger James O (1992) Ockham's razor and Bayesian analysis. Am Sci 80(1):64–72
17. Ying Kuen C, Rick C (2002) A simple technique to evaluate model sensitivity in the continual reassessment method. Biometrics 58(3):671–674
18. Ying Kuen C (2011) Dose finding by the continual reassessment method. CRC Press, Boca Raton
19. Shing ML, Kuen CY (2011) Calibration of prior variance in the Bayesian continual reassessment method. Stat Medicine 30(17):2081–2089
20. Clertant M, O'Quigley J et al (2017) Semiparametric dose finding methods. J Royal Stat Soc Series B 79(5):1487–1508