

A Comparison of Model-Free Phase I Dose Escalation Designs for Dual-Agent Combination Therapies

Journal Title

XX(X):2-42

© The Author(s) 2016

Reprints and permission:

sagepub.co.uk/journalsPermissions.nav

DOI: 10.1177/ToBeAssigned

www.sagepub.com/

SAGE

Helen Barnett¹, Matthew George^{1,2}, Donia Skanji³,
Gaelle Saint-Hilary^{4,5}, Thomas Jaki^{6,7} and Pavel
Mozgunov⁶

Abstract

It is increasingly common for therapies in oncology to be given in combination. In some cases, patients can benefit from the interaction between two drugs, although often at the risk of higher toxicity. A large number of designs to conduct phase I trials in this setting are available, where the objective is to select the maximum tolerated dose combination (MTC). Recently, a number of model-free (also called model-assisted) designs have provoked interest, providing several practical advantages over the more conventional approaches of rule-based or model-based designs. In this paper, we demonstrate a novel calibration procedure for model-free designs to determine their most desirable parameters. Under the calibration procedure, we compare the behaviour of model-free designs to model-based designs in a comprehensive simulation study, covering a number of clinically plausible scenarios. It is found that model-free designs are competitive with the model-based designs in terms of the proportion of correct selections of the MTC. However, there are a number of scenarios in which model-free designs offer a safer alternative. This is also illustrated in the application of the designs to a case study using data from a phase I oncology trial.

Keywords

Dose-Finding; Combination Therapies; Model-Free Designs; Phase I Trials.

¹Department of Mathematics and Statistics, Lancaster University

²Phastar

³Servier

⁴Saryga

⁵Politecnico di Torino

⁶MRC Biostatistics Unit, University of Cambridge

⁷University of Regensburg

Corresponding author:

Helen Barnett

Email: helen.barnett@lancaster.ac.uk

Introduction

The aim of phase I clinical trials investigating a single therapy is to find the highest dose that can be administered whilst ensuring that patients are at a low risk of serious side effects. To offer patients a higher chance of successful treatment, there is willingness to accept a dose that leads to more toxic responses, commonly labelled as dose-limiting toxicities (DLTs). The highest dose for which the treatment has a pre-specified probability of leading to a toxic outcome (target toxicity) is called the maximum tolerated dose (MTD). In an analysis of over 400,000 clinical trials conducted between 2000 and 2015³⁰, it was found that 57.6% of all phase I oncology trials successfully progressed to phase II. It was found that in 73% of trials excluding oncology, treatments were successful in moving to phase II, thus demonstrating the importance of successful dose-finding methods in oncology, where drugs are clearly harder to develop.

In this work, we consider phase I oncology trials in which a combination of two therapies are investigated. Here the objective is to identify a maximum tolerated dose combination (MTC), the dose combination with a probability of toxicity closest to the target toxicity. Phase I oncology trials in this dual-agent setting have recently provoked notable interest³¹. In particular, it was found that immunotherapy, a targeted agent that stimulates the immune system to fight cancerous cells³, can provide benefit to patients when administered in combination with chemotherapy or another targeted agent²⁵. One difficulty in the dual-agent setting is that the order of toxicity is unknown for some combinations – if the amount of one compound in the combination is increased while another is decreased, it is unknown whether the overall toxicity goes up or down.

A number of dose-finding methods for dual-agent combination phase I trials relaxing the monotonicity assumption on the order of some of the combinations have

been proposed in the literature. They broadly belong to one of three categories; rule-based, model-based and model-free (also known as model-assisted) designs. Rule-based designs (e.g. 3+3+3 or extensions of this^{5,8,10}) rely on a number of pre-specified rules to determine when a dose is escalated, de-escalated and chosen as the MTD. Model-based designs (e.g. BLRM¹⁸, six-parameter model²⁷, Partial Ordering Continual Reassessment Method²⁸ and the modified logistic model²⁴) model the relationship between dose and probability of toxicity through a parametric function. Through the course of a trial, parameter estimates are updated to better describe this relationship. The model-free designs^{1,14} do not pre-specify any relationship between dose and toxicity, thus do not rely on any parametric assumptions in their search for the MTD. However, unlike rule-based designs, the decision process in which the dose can be escalated or de-escalated is assisted with a statistical model.

Despite numerous papers demonstrating flaws in rule-based designs and their performance in drug combination trials^{2,22,26}, it was reported that less than 5% of combination trials in oncology between 2011 and 2013 deviated from rule-based designs²³. It is perhaps the restrictions associated with model-based designs, such as difficulty of implementation or communication to clinicians, that have made these less commonly used in real trials. Recently, model-free designs have attracted attention due to their practicality²⁹, although these have not yet been fully evaluated in the literature.

The objective of this work is to review five recently proposed model-free dose-finding designs for phase I dual-agent combination studies, namely, the Bayesian Optimal Interval design¹¹ BOIN, the Keyboard design²⁰ KEY, the surface-free design¹³ SFD, the product of independent beta probabilities design¹² PIPE, and the Waterfall design³⁴. We evaluate their performance in an extensive simulation study. We note that some comparison has already

been investigated by previous authors, for example the comparison of the Keyboard design to two other designs²⁰. The novelty in our approach is to compare the methods on equal grounds. We hence propose a calibration procedure that selects the parameters of each of the designs that maximise the proportion of correct selections (subject to a safety constraint). We compare the performance of these designs to three model-based designs; the Bayesian Logistic Regression Model (BLRM), a model-based approach that uses a two-parameter logistic model for each compound¹⁸; the Partial Ordering Continual Reassessment Method (POCRM)²⁸; and the modified logistic model²⁴, as well as a non-parametric optimal benchmark¹⁷. We also evaluate the performance of each of the designs in a case study of neratinib and temsirolimus⁶, to highlight the differences between approaches in a real trial setting of a dose finding trial from combination therapies.

The rest of the paper continues as follows. We first provide a review of model-free designs, before using a novel method to calibrate the parameters of each design leading to good performance. We then present detailed results from our simulation study across a wide range of toxicity scenarios, including the model-based designs for comparison. Each design is also applied to the real case study of neratinib and temsirolimus. We finish with a discussion of our results.

Methodological Review

In this section, we describe the dose escalation procedure for each of the five model-free approaches in a general dose-finding trial. It is assumed patients enter the trial in cohorts, and the dose combination for the next cohort is assigned once the previous cohort's responses are available. We first define the admissible combinations for each design. These are the dose combinations that are allowable for assignment for the next cohort of patients based on the last tested combination. We then describe the details of

the escalation procedure in each of the designs in the following setting. Consider a dual-agent trial with I doses of drug A, denoted $d_1^A < \dots < d_I^A$ and J doses of drug B, denoted $d_1^B < \dots < d_J^B$. Let d_{ij} represent the combination of doses d_i^A and d_j^B for $i = 1, \dots, I$ and $j = 1, \dots, J$. The total number of patients who receive combination d_{ij} and the number of those who experience a toxic response on d_{ij} during the trial are denoted n_{ij} and y_{ij} respectively. The probability of toxic response at d_{ij} is written as π_{ij} and the target toxicity is denoted ϕ .

Admissible Combinations

Before deciding on a dose for the next cohort, each design defines a set of combinations that are admissible; i.e. combinations that the next cohort could be allocated to. These are best illustrated with a diagram, Figure 1. Suppose we are at d_{22} in Figure 1, indicated by the ‘●’ symbol. Admissible combinations for the BOIN, Waterfall and KEY designs are the same combination or adjacent combinations to the current one ($d_{ij} \rightarrow \{d_{i-1,j}, d_{i+1,j}, d_{i,j-1}, d_{i,j+1}\}$), represented by the ‘○’ symbols.

In addition to these combinations, the SFD and PIPE also allow for diagonal de-escalation, where the next cohort is administered a combination that is one dose level lower in each drug ($d_{ij} \rightarrow d_{i-1,j-1}$), and also allow for anti-diagonal escalation, meaning the next cohort receives a combination that is one dose level higher in one drug and one dose level lower in the other ($d_{ij} \rightarrow \{d_{i-1,j+1}, d_{i+1,j-1}\}$). These are depicted by the ‘*’ symbols in Figure 1, where reaching d_{11} requires diagonal de-escalation and reaching d_{31} or d_{13} requires anti-diagonal escalation. The rationale is that by enabling faster movement across the combination grid, the design can move to the MTC quickly, and de-escalate quickly if patients are treated at highly toxic combinations.

All designs prohibit diagonal escalation, where the next cohort receives a combination one dose level higher in each drug ($d_{ij} \rightarrow d_{i+1,j+1}$) and no dose levels can be skipped.

These non-admissible doses are shown by the '×' in the red cells in Figure 1.

[Figure 1 about here.]

The BOIN Design

The Bayesian Optimal Interval (BOIN) design¹¹ uses the intuitive estimator $\hat{\pi}_{ij} = y_{ij}/n_{ij}$ for the probability of toxicity at combination d_{ij} , so that $\hat{\pi}_{ij}$ is the proportion of observed toxic responses on d_{ij} across the whole trial. The estimator $\hat{\pi}_{ij}$ only updates after patient responses are observed on d_{ij} , and is then used to guide dose escalation. This escalation process is defined by pre-specified values of λ_e, λ_d to which $\hat{\pi}_{ij}$ is compared after each cohort. Whilst ϕ is the target toxicity, ϕ_1 is the highest toxicity probability deemed sub-therapeutic and ϕ_2 is the lowest toxicity probability deemed overly toxic. These can be specified by the clinicians. Values of $\phi_1 < \lambda_e < \phi$ and $\phi_2 > \lambda_d > \phi$ are chosen to locally minimise the chance of incorrect escalation and de-escalation decisions during a trial, and are calculated using constants ϕ_1 and ϕ_2 . λ_e and λ_d are defined in Equation (1):

$$\lambda_e = \frac{\log\left(\frac{1-\phi_1}{1-\phi}\right)}{\log\left(\frac{\phi(1-\phi_1)}{\phi_1(1-\phi)}\right)} \quad \text{and} \quad \lambda_d = \frac{\log\left(\frac{1-\phi}{1-\phi_2}\right)}{\log\left(\frac{\phi_2(1-\phi)}{\phi(1-\phi_2)}\right)}. \quad (1)$$

Both λ_e and λ_d are invariant to d_{ij} , n_{ij} and y_{ij} , so that optimising these parameters depends only on constants ϕ , ϕ_1 and ϕ_2 . After defining λ_e and λ_d , the rules for the dose-finding procedure are as follows:

- If $\hat{\pi}_{ij} \leq \lambda_e$, the next combination is chosen from $\mathcal{A}_E = \{d_{(i+1)j}, d_{i(j+1)}\}$.
- If $\hat{\pi}_{ij} > \lambda_d$, the next combination is chosen from $\mathcal{A}_D = \{d_{(i-1)j}, d_{i(j-1)}\}$.
- Otherwise, $\lambda_e < \hat{\pi}_{ij} \leq \lambda_d$ and the next combination is the same.

In this way, dose skipping, diagonal escalation and diagonal de-escalation are prohibited – see the ‘Admissible Combinations’ section for more details. If the next combination is to be chosen from an empty \mathcal{A}_E or \mathcal{A}_D (for example the current combination is the highest in both doses and the design chooses to escalate), then the next cohort receives the same combination. The design assumes each patient response is independent, $y_{ij} \sim \text{Binomial}(n_{ij}, \pi_{ij})$ and assigns a vague Beta(1,1) prior distribution to each π_{ij} , giving the posterior distribution for π_{ij} as

$$\pi_{ij} | n_{ij}, y_{ij} \sim \text{Beta}(y_{ij} + 1, n_{ij} - y_{ij} + 1). \quad (2)$$

To choose between combinations in the chosen set, the BOIN design computes the posterior probability $\mathbb{P}(\pi_{ij} \in (\lambda_e, \lambda_d) | n_{ij}, y_{ij})$. The combination maximising this probability is administered to the next cohort. For combinations yet to be tested, calculating this probability is based on the vague prior distribution only. In the event of ties, which is always the case when multiple potential combinations are yet to be administered, the next combination is selected at random from the chosen set. Note that no toxicity information is borrowed between the combinations under this model as the combinations are treated independently.

The design uses an overdosing criterion stating that a combination, and any that are more toxic under monotonicity, satisfying $\mathbb{P}(\pi_{ij} > \phi | n_{ij}, y_{ij}) \geq \epsilon_{BOIN}$ for some overdosing probability threshold $0 < \epsilon_{BOIN} \leq 1$, cannot be administered to the next cohort. For the BOIN design, if d_{ij} satisfies this condition, dose d_{ij} and higher combinations are eliminated from the trial, and the dose maximizing $\mathbb{P}(\pi_{ij} \in (\lambda_e, \lambda_d) | n_{ij}, y_{ij})$ within \mathcal{A}_D is chosen for the next cohort. If combination d_{11} satisfies the overdosing criterion, the trial is terminated earlier for safety.

After all patients are treated, estimates of each π_{ij} are calculated via matrix isotonic regression⁴. The simple

technique guarantees that estimates of π_{ij} at higher combinations are at least as high as estimates of π_{ij} at lower combinations, which follows the assumption of monotonicity. The MTC is selected as the combination with estimated π_{ij} closest to ϕ via isotonic regression⁴.

The Keyboard Design

The Keyboard design (KEY)³² is very similar to the BOIN design, defining an interval about the target toxicity ϕ , denoted $I_{target} = (\phi - \Delta_1, \phi + \Delta_2)$, for constants $\Delta_1, \Delta_2 > 0$, which can be chosen by the clinicians. A combination with estimated toxicity probability within this interval is said to have acceptable toxicity. The design then divides the $(0,1)$ space into “keys”, defined as intervals I_t of equal length $\Delta_1 + \Delta_2$ (allowing for shorter keys at either end of $(0,1)$) for $t = 1, \dots, T$, where T is the number of keys. The interval I_{target} is fixed pre-trial, chosen to minimise the chance of incorrect escalation and de-escalation decisions.

The KEY design assigns a vague Beta(1,1) prior distribution to each π_{ij} , and assumes that the number of toxic responses follows a binomial distribution, $y_{ij} \sim \text{Binomial}(n_{ij}, \pi_{ij})$. The posterior distribution for each π_{ij} is computed as in Equation 2. Again, this means there is no borrowing of toxicity information across combinations. The design then identifies the key I_t that is most likely to contain π_{ij} , labelled I_{max} ,

$$I_{max} = \underset{I_t: t \in (1, \dots, T)}{\operatorname{argmax}} \mathbb{P}(\pi_{ij} \in I_t | n_{ij}, y_{ij}). \quad (3)$$

Once the key I_{max} is identified, escalation and de-escalation decisions happen as follows:

- If $I_{max} < I_{target}$, the next combination is chosen from $\mathcal{A}_E = \{d_{(i+1)j}, d_{i(j+1)}\}$.
- If $I_{max} > I_{target}$, the next combination is chosen from $\mathcal{A}_D = \{d_{(i-1)j}, d_{i(j-1)}\}$.
- If $I_{max} = I_{target}$, the next combination is the same.

To choose between combinations in \mathcal{A}_E (or \mathcal{A}_D), the design computes the posterior probability $\mathbb{P}(\pi_{ij} \in I_{target} | n_{ij}, y_{ij})$ for all combinations in \mathcal{A}_E (or \mathcal{A}_D). The combination maximising this probability is administered to the next cohort. The remainder of the escalation process and the selection of the MTC is analogous with the BOIN design, with an identical overdosing rule using ϵ_{KEY} and the MTC chosen via isotonic regression⁴.

Both the BOIN and KEY designs model dose combinations independently, however in the following two designs, the connections between the dose combinations are also taken into account.

The Surface-Free Design

The surface-free design (SFD)¹³ does not restrict the MTC search to a parametric surface and does not require the order of toxicity between combinations to be known. The main idea is to parametrise ratios between toxicity probabilities for different combinations, defining $\theta = 1 - \pi_{11}$, and $\theta_i = \frac{1 - \pi_{i,j}}{1 - \pi_{i-1,j}}$. Then θ is the probability of a patient having no toxic response on the lowest dose combination and θ_i denotes the ratio between the probability of a patient having no toxic response on dose combinations d_{ij} and $d_{(i-1)j}$ for $j = 2, \dots, J$ and $i = 2, \dots, I$. Similarly, $\tau_j = \frac{1 - \pi_{i,j}}{1 - \pi_{i,j-1}}$ is defined as the ratio between the probability of a patient having no toxic response on d_{ij} and $d_{i(j-1)}$ for $j = 2, \dots, J$ and $i = 2, \dots, I$. Thus, the probability of toxicity for each combination d_{ij} is

$$\pi_{ij} = 1 - \theta\theta_2 \dots \theta_i \tau_2 \dots \tau_j. \quad (4)$$

Due to monotonicity, each ratio $\theta_i, \tau_j \in (0, 1)$ and the SFD assigns each of these ratios an independent Beta prior distribution. The hyper-parameters of the prior distributions can be chosen to match the clinicians' prior mean estimates of toxicity probability on each combination and effective sample sizes.

After each cohort, the SFD updates the posterior means for ratios $\theta, \theta_2, \dots, \theta_I, \tau_2, \dots, \tau_J$ using Bayes theorem, which can be related back to π_{ij} through Equation (4) to give estimates of the toxicity probabilities. In this way, the SFD is borrowing information across various drug combinations previously collected in the trial to make an informed decision on escalation. Additionally, the continual multiplication of Beta random variables implies that π_{ij} for higher combinations has higher variance, allowing for more cautious escalation at higher combinations. Considering all neighbouring combinations apart from the one higher in both doses, the next combination is chosen as the one with estimated π_{ij} closest to ϕ . An overdosing criterion prohibits any combination from being administered if $\mathbb{P}(\pi_{ij} > \phi | n_{ij}, y_{ij}) \geq \epsilon_{SFD}$ for some $\epsilon_{SFD} > 0$, and the trial is terminated if this is satisfied for d_{11} .

Once all patients have been treated, the MTC is selected as the combination with toxicity probability closest to ϕ . Note that the SFD design is more computationally intensive than the other model-free designs as MCMC methods are required to sample from the posterior distribution.

The PIPE Design

The PIPE design¹² differs from the model-free designs discussed so far in that it was originally proposed to find the MTC contour, labelled MTC_ϕ . This is a line partitioning the combination space into safe and overly toxic combinations. Those below the contour are believed to have toxicity probability less than target toxicity ϕ , whilst those above are believed to have toxicity probability greater than ϕ .

Assuming the π_{ij} are independent, they are assigned a Beta prior distribution, $\pi_{ij} \sim \text{Beta}(a_{ij}, b_{ij})$ for hyperparameters a_{ij} and b_{ij} , for $i = 1, \dots, I$ and $j = 1, \dots, J$. Priors can be pre-specified if knowledge on the toxicity of combinations is available. Assuming each patient is

independent such that $y_{ij} \sim \text{Binomial}(n_{ij}, \pi_{ij}) \forall i, j$, the posterior for π_{ij} can be written as

$$\pi_{ij} | n_{ij}, y_{ij} \sim \text{Beta}(y_{ij} + a_{ij}, n_{ij} - y_{ij} + b_{ij}). \quad (5)$$

The posterior distribution is only updated after a cohort of patients is treated on the corresponding combination, but the MTC_ϕ is re-estimated regardless of which combination was tested. The monotonicity assumption means that the PIPE design needs only to consider contours satisfying this property, limiting the number of possible contours to $\binom{I+J}{I}$.

Each contour can be represented by a binary matrix, where entries are 0 or 1 depending on whether estimates of the toxicity probability for a combination are below or above the contour respectively. Let ϑ be the set of all monotonic contours for an $I \times J$ dose combination space and define $C_s \in \vartheta$ as the binary matrix representing the contour $s = 1, \dots, \binom{I+J}{I}$.

To estimate the MTC_ϕ given the current data, the design calculates the posterior probability of each toxicity probability being less than or equal to ϕ , that is

$$p_{ij}(\phi | n_{ij}, y_{ij}) = \mathbb{P}(\pi_{ij} \leq \phi | n_{ij}, y_{ij}, a_{ij}, b_{ij}), \quad (6)$$

where the right-hand side of Equation (6) is equal to the cumulative distribution function of a Beta distribution. Equation (8) gives the general formula for calculating the probability that the MTC_ϕ is defined by the matrix C_s :

$$\begin{aligned} & \mathbb{P}(MTC_\phi = C_s | n_{ij}, y_{ij}) \quad (7) \\ &= \prod_{i=1}^I \prod_{j=1}^J \{1 - p_{ij}(\phi | n_{ij}, y_{ij})\}^{C_s^{[i,j]}} p_{ij}(\phi | n_{ij}, y_{ij})^{1 - C_s^{[i,j]}}, \quad (8) \end{aligned}$$

where $[i, j]$ represents the entry in the i th row and j th column of the binary matrix C_s . The contour maximising Equation (8) is the contour most likely to be the MTC_ϕ given the current data. This contour then assists the

escalation process by identifying the combinations closest to it, before the design selects one of these for the next cohort based on a weighted randomisation procedure. This involves weighting each combination by the inverse of their sample size, with the rationale being varied experimentation around the MTC_ϕ . Escalation continues in this way until all patients are treated, at which point all combinations closest from below the MTC_ϕ are recommended for phase II.

The design uses an overdosing rule which considers the expected probability of d_{ij} being above the most probable MTC_ϕ averaged over all monotonic contours. This is written as

$$q_{ij} = \sum_{C_s \in \vartheta} C_s[i, j] \mathbb{P}(MTC_\phi = C_s | \mathbf{Y}^{(m)}),$$

and d_{ij} cannot be administered to the next cohort if $q_{ij} \geq \epsilon_{PIPE}$ for some $\epsilon_{PIPE} > 0$. A trial is terminated if combination d_{11} satisfies this condition.

The PIPE design can recommend multiple combinations for phase II, as it recommends all combinations closest from below its MTC_ϕ . For consistency across designs, in our implementation we ensure only one combination is recommended as the MTC. Therefore for each recommended combination, we find the posterior mean probability of toxicity, which can be calculated using the posterior distributions in Equation (6). The combination with posterior mean closest to ϕ is selected as the MTC, choosing a combination at random in the event of a tie.

The Waterfall Design

The Waterfall design³⁴ also aims at finding the MTC contour. This design breaks down the two-dimensional dosing grid into a series of one-dimensional sub-trials. For the $I \times J$ dosing grid, the I sub-trials are as follows:

$$\begin{aligned}
S_I &= \{d_{I1}, \dots, d_{I1}, d_{I2}, \dots, d_{IJ}\}, \\
S_{I-1} &= \{d_{I-1,2}, \dots, d_{I-1,J}\}, \\
S_{I-2} &= \{d_{I-2,2}, \dots, d_{I-2,J}\}, \\
&\dots \\
S_1 &= \{d_{12}, \dots, d_{1J}\}.
\end{aligned}$$

Sub-trials are conducted sequentially using a single-agent dose-finding method, with the single-agent BOIN design recommended by the authors.

Firstly, sub-trial S_I is conducted, starting at dose d_{11} , and the first so-called ‘candidate MTD’ for the sub-trial is found using the single-agent design. The next sub-trial is chosen based on this candidate MTD. For a candidate MTD of d_{i^*,j^*} , the next sub-trial to be conducted is S_{i^*-1} . The process of selecting a sub-trial and candidate MTD is repeated until sub-trial S_1 is completed. All responses are then collated, and matrix isotonic regression is used to select the dose in each row of the dose grid with the estimate of toxicity probability closest to the target, unless all doses in that row are overly toxic. These dose combinations make up the MTC contour that is recommended.

Similarly to the PIPE design, since many combinations are recommended, for consistency we select one MTC based on the posterior mean probability of each dose combination that is recommended.

Calibration of Designs

Model-based and model-free designs based on a Bayesian framework give clinicians more control over their performance. The PIPE design, the SFD and most model-based designs allow for knowledge on the toxicity of each drug from monotherapy trials to be incorporated into the design through their prior distributions. As the BOIN and KEY designs assign vague priors to the toxicity

probabilities, their behaviour is primarily determined by the pre-defined intervals guiding escalation. Although it is in theory possible to incorporate historical data through the prior in the BOIN and KEY designs³³, for the purpose of this comparison, it would defeat the purpose of a design with all escalation boundaries pre-specified at the design stage for ease of implementation.

Since for all designs, the hyper-parameter values of the prior and the values of the intervals have a substantial effect on the escalation procedure, any attempt to compare designs objectively must ensure that these values are specified in a fair way. In this comparison study, the aim of the novel calibration procedure is to give all designs a set-up which leads to consistently high proportions of selections of combinations with toxicity probability close to ϕ in all scenarios, whilst keeping the number of patients allocated to unsafe doses low. Therefore, an a priori pre-specified criterion for selection of the prior parameters is used instead of the originally subjective proposed parameters for each model. Please see the Supplementary Materials for the comparison of calibrated and originally proposed hyper-parameters.

Each design considered in the comparison is calibrated using a novel two-stage approach. The first stage of the calibration is concerned with choosing values for hyper-parameters that give a good performance in selecting the MTC without considering safety. The second stage then focusses on safety, calibrating the overdose rule taking into account not only good performance in terms of recommending no combinations when considering an overly toxic scenario, but also the number of patients who are treated at unsafe doses. Although using similar principles to a standard fine-tuning approach, the novelty of this two-stage calibration is in this lesser subjective but still intuitive choice of hyper-parameters.

This two-stage calibration procedure based on high performance and safety is applied for all designs, employing

a grid search over hyper-parameter or interval values (depending on the design). At each stage this involves running simulations over four clinically plausible scenarios and determining which values lead to superior performance when averaged across the scenarios. We refer to the priors resulting in superior performance across the four scenarios as operational priors. Whilst this procedure gives the Bayesian designs an opportunity to be compared fairly against each other through their performance, these priors can also be applicable in the practical case where no reliable prior information about the compounds is available.

The first stage of the calibration procedure evaluates which design inputs lead to superior performance in recommending the MTC across the four scenarios. The proportion of correct selections (PCS) is examined in each scenario. This is the proportion of trials in which a design selects any combination with a true toxicity probability of exactly 0.30. To summarise the overall performance across these four scenarios, the geometric mean PCS is used. For the remainder of this section, the mean will refer to the geometric mean. Suppose x_1, \dots, x_N represent the PCS in N scenarios. The geometric mean, $\left(\prod_{i=1}^N x_i\right)^{1/N}$, is used instead of the arithmetic mean because it has the useful property of penalising cases in which PCS are more dispersed across scenarios. The combination of hyper-parameter or interval values that result in the highest geometric mean PCS across the scenarios is chosen. We note that during the first stage of the calibration procedure, no overdosing rules are included, meaning no trials are to be stopped before all patients have been recruited, because we choose to calibrate the parameter controlling the overdosing rule in the separate second stage. Once the first stage of calibration is complete, this will lead to the selection of intervals for the BOIN, Waterfall and KEY designs, and operational priors for the PIPE and SFD designs. The same approach is used to determine the

operational priors for the BLRM, the POCRM and the logistic model, which are used as model-based comparators in the next section.

The second stage of the calibration procedure is for ϵ , the parameter regulating the overdosing rule in each design. Calibration of ϵ involves decreasing its value starting from 1, and observing the number of patients treated at overly toxic doses and the proportion of correct outcomes in the chosen scenarios. As selecting overly toxic combinations is more of an ethical concern, as a general rule we take as a starting point the highest value of ϵ resulting in at least 85% of trials recommending no combinations when considering an overly toxic scenario. We acknowledge this proportion may differ in practice depending on the clinicians' judgement. It is important to note that the interpretation of ϵ differs between designs because of the construction of each overdosing rule, and should be accounted for when communicating with clinicians. This is reflected by subscripts for the individual designs in the following specifications. The second stage of the calibration procedure for each design is illustrated in Figure 2. Note that the proportion of correct selections for all four scenarios is shown for completeness, although the second stage of the procedure itself uses this metric only for scenario 14, where the correct selection is the selection of no dose combination since all are overly toxic.

Setting

Each design is calibrated in the same setting that is then explored in the simulation study, representative of a phase I trial in oncology. There are two drugs with three dose levels each, which results in nine combinations, and the first cohort is treated at the lowest combination. The objective is to select a single combination as the MTC with true toxicity probability $\phi = 0.30$. The sample size is 36 patients for which are recruited in cohorts of three patients. All combination-toxicity scenarios are presented in Table 1.

However, four scenarios are chosen to explore noticeably different clinical cases, in which the number and location of the MTCs vary, whilst restricting the number of scenarios makes the procedure computationally feasible.

[Table 1 about here.]

In stage 1 of the calibration procedure, Scenarios 1, 8, 10 and 13 are chosen to give a diverse range of plausible scenarios. We have found the procedure to be robust to changes in these scenarios provided the qualitative features (e.g. if there is at least one scenario where the optimal combination is off diagonal) are maintained. The number of scenarios is chosen to balance computational feasibility and breadth of differing scenarios. Scenarios 1 and 13 are chosen to represent the extremes: when the highest combination is the only true MTC and all others are safe, and when the lowest combination is the only true MTC and all others are overly toxic, respectively. Scenario 8 covers situations in which most combinations are safe but true MTCs do not lie on the same diagonal. Scenario 10 captures the case where most combinations are overly toxic and true MTCs lie on the same diagonal. Note that we often refer to the set of combinations in a scenario as the combination grid.

In stage 2, simulations are run for each design over Scenarios 8, 10, 13 and 14 for different values of ϵ . In Scenarios 8, 10 and 13, the PCS is as previously defined, whilst in the unsafe Scenario 14 we consider the PCS as the proportion of trials in which no combinations are selected. We refer to selecting no combinations in Scenario 14 as the ‘correct outcome’.

We summarize the optimal choice of hyper-parameters alongside the recommendations from the original proposal of each design in the online supplementary materials.

[Figure 2 about here.]

Calibrating the BOIN Design

To guide dose escalation, the BOIN design relies on the interval (λ_e, λ_d) around the target toxicity. Interval boundaries λ_e and λ_d are a function of ϕ , ϕ_1 and ϕ_2 , where $\phi_1 = a_1\phi$ and $\phi_2 = a_2\phi$ for constants $a_1 < 1, a_2 > 1$. To calibrate the design, we run 4000 simulations for each scenario for pairs (a_1, a_2) from the sets $a_1 = \{0.85, 0.80, \dots, 0.40\}$ and $a_2 = \{1.15, 1.20, \dots, 1.60\}$, resulting in a total of 100 pairs. As constants a_1 and a_2 deviate further from 1, the interval becomes wider, thus the design will choose to escalate and de-escalate on fewer occasions.

The optimal values are found to be $a_1 = 0.65$ and $a_2 = 1.4$, which substituting into Equation (1), we generate the interval boundaries λ_e and λ_d to give the interval $(0.245, 0.359)$ to guide dose escalation. This interval implies that escalation occurs if π_{ij} is below 0.245, de-escalation occurs if π_{ij} is above 0.359, else the combination remains the same.

In the second stage of calibration, we find that as ϵ_{BOIN} decreases, the design benefits more in Scenario 14, where the proportion of trials in which no combinations are recommended increases (see Figure 2). For $\epsilon_{\text{BOIN}} \leq 0.84$, over 85% of trials recommend no combinations in Scenario 14. The trade-off in the other scenarios with this ϵ_{BOIN} value is that PCS increases steeply when ϵ_{BOIN} increases, as well as the number of patients treated on overly toxic doses increasing. Therefore $\epsilon_{\text{BOIN}} = 0.84$ is chosen.

Calibrating the Keyboard Design

Using a similar method to BOIN, we first calibrate the parameters which define the interval for KEY. The interval $I_{\text{target}} = (b_1, b_2)$ guides escalation entirely so is an important component of the design. We run 4000 simulations across each scenario for pairs (b_1, b_2) from the sets $b_1 = \{0.27, 0.25, \dots, 0.19\}$ and $b_2 = \{0.33, 0.35, \dots, 0.41\}$, resulting in a total of 25 pairs. Mean PCS are displayed Figure 2 in the online supplementary

materials, indicating that interval (0.21, 0.39) yields the highest mean PCS, which differs from the recommendation of (0.25, 0.35) in the original paper³². As explained in the Methodological Review, this means escalation occurs only if the posterior probability $\mathbb{P}(\pi_{ij} \in (0.03, 0.21)|n_{ij}, y_{ij})$ is higher than $\mathbb{P}(\pi_{ij} \in (0.21, 0.39)|n_{ij}, y_{ij})$.

In the second stage of calibration, we find that as ϵ_{KEY} decreases, the design benefits more in Scenario 14, where the proportion of trials in which no combinations are recommended increases. Choosing $\epsilon_{\text{KEY}} = 0.84$ leads to approximately 85% of trials correctly selecting no combinations in Scenario 14, as shown in Figure 2, in line with the value obtained for the BOIN design.

Calibrating the Surface-Free Design

The SFD assigns Beta priors to each of its parameters; the ratios between toxicity probabilities. In this setting, there are five ratios (θ , θ_2 , θ_3 , τ_2 & τ_3 defined in the Methodological Review) to parametrise, meaning a total of 10 hyper-parameters for the beta priors must be defined for the operational priors. Instead of specifying these directly, we specify a prior mean and prior effective sample size for each ratio, which can be used to calculate the corresponding hyper-parameters. To make the calibration task computationally feasible, we assume that all prior mean ratios, m , are equal (meaning the increase in dose corresponds to the same proportion increase in toxicity) and all effective sample sizes for each ratio, s_{SFD} , are equal. This drastically reduces the dimensionality of the grid search. We note, however, that if one has reliable prior information about the ratio not being equal, one can calibrate over various ratios being mindful of the computational complexity of such a calibration. Thus we only need to calibrate pairs of m and s , which we choose from sets $m = \{0.95, 0.925, \dots, 0.85\}$ and $s_{\text{SFD}} = \{1, 2, \dots, 5\}$.

For each pair, we run 1000 simulations (which is lower than other model-free designs due to the computational demands of the design) and examine the mean PCS across the four scenarios. Our results in Figure 4 in the online supplementary materials show that the mean PCS is highest for $m = 0.875$ and $s_{\text{SFD}} = 4$. This is equivalent to every ratio being assigned the prior distribution $\text{Beta}(3.5, 0.5)$, and corresponds to mean prior toxicity probabilities on d_{11} and d_{33} of 0.125 and 0.487 respectively. We also note that the SFD has been found to be robust to small to moderate deviations in these parameters.

For the calibration of ϵ_{SFD} , in Figure 2, we found that $\epsilon_{\text{SFD}} = 0.65$ resulted in at least 85% of trials selecting no combinations in Scenario 14. There is evidence to suggest that increasing or decreasing ϵ_{SFD} not only has a sizeable effect on the PCS in Scenario 13, but also the number of patients treated at unsafe doses, demonstrating the design is highly sensitive to changes in its overdosing rule.

Calibrating the PIPE Design

Similar to the SFD, the PIPE designs assigns beta priors to each π_{ij} . A prior mean and prior sample size for each π_{ij} are specified, giving a total of 18 values to specify from which the hyper-parameters for the beta priors can be calculated. To make calibration feasible, we assume that prior sample size s_{PIPE} is equal for each combination and to set the prior means, we divide the grid of combinations into five diagonal segments, with toxicity increasing as we move through each segment. In this way, the design follows the monotonicity assumption. Like for the SFD, this drastically reduces the dimensionality of the grid search. To assign a toxicity to each combination, we specify the toxicity of the lowest combination, ρ , and the size of the increments in toxicity between each segment, δ . In the illustration in Figure 3, we have chosen $\rho = 0.05$ and $\delta = 0.025$ to construct the grid.

[Figure 3 about here.]

Our approach involves calibrating three parameters simultaneously to create operational priors, and are chosen from the sets $\rho = \{0.025, 0.05, 0.075, 0.10\}$, $\delta = \{0.025, 0.05, 0.075, 0.10\}$ and $s_{\text{PIPE}} = \{1/72, 1/36, 1/18, 1/9\}$. For each triple, we run 2000 simulations in each of the four scenarios, which is fewer than for the BOIN and KEY designs due to the minor increase in computational expense. We provide one grid in Figure 5 in the online supplementary materials to account for mean PCS on each s_{PIPE} value. The triple $s_{\text{PIPE}} = 1/18$, $\rho = 0.05$ and $\delta = 0.025$ leads to the highest mean PCS, although we observe that there were many triples that resulted in similar values. We note our choice of prior sample size, $s_{\text{PIPE}} = 1/18$, only differs to the recommendation of $1/9$ in the original paper¹². For prior sample sizes $s_{\text{PIPE}} \leq 1/18$, the design is found to be robust. Mean PCS only varies between 37% and 40%, suggesting that a number of operational priors could lead to consistently high PCS.

For the second stage of the calibration, the value of ϵ_{PIPE} is varied as shown in Figure 2, and $\epsilon_{\text{PIPE}} = 0.50$ is chosen as it provides at least an 85% chance of correctly recommending no combinations in Scenario 14, as well as balancing the number of patients treated at unsafe doses in the four considered scenarios.

Calibrating the Waterfall Design

The calibration of the Waterfall Design is in line with the calibration of the BOIN design, since the parameters of escalation are the same. Therefore the values of the hyper-parameters used for the interval (λ_e, λ_d) are the same: (0.245, 0.359) and the parameter for the safety constraint is also the same, $\epsilon_W = 0.84$. For the 3×3 dosing grid, there are three sub-trials, the first has 5 dose combinations and the following two have 2 dose combinations. Therefore to split the 12 cohorts between the sub-trials, six are allocated

to the first sub-trial and three to each of the other sub-trials.

Simulation Study: 3×3 Dosing Grid

In this section we describe the setting for the simulation study of 3×3 dosing grids before presenting the results, including a comparison to a model-based approach and a non-parametric optimal benchmark.

Setting

In order to compare the discussed designs, we conduct a simulation study, performing 2000 simulations of each of the 15 scenarios depicted in Table 1 for all five designs. As before, the objective is to select a single combination as the MTC with true toxicity probability $\phi = 0.30$. Any combination with probability of toxicity greater than 0.33 is labelled as overly toxic, and any combination with probability of toxicity in the interval $[0.16, 0.33]$ is labelled as acceptable. In this section, the mean refers to the arithmetic mean unless specified otherwise. All simulations are performed in R²¹, with code provided in the online supplementary materials.

In general, the number of overly toxic combinations available for selection increases as we move through Scenarios 1 to 14. Scenario 1 has a single MTC which is the highest combination available. Scenarios 3 and 4 contain very few overly toxic combinations and have MTCs on the edge of the grid. Scenario 5 is similar to these, except its only MTC is located in the centre of the grid. In Scenarios 2, 6, 7, 8, 9 and 10, there are multiple combinations to explore which have toxicity probability ϕ . In particular, Scenarios 8 and 9 aim to investigate design behaviour when underlying MTCs are not on the same diagonal. Scenarios 11, 12 and 13 represent settings in which most combinations are overly toxic, meaning designs should avoid combinations away from d_{11} . Scenario 14 is of importance because all of its combinations are

overly toxic, making the trial very unethical. In this instance, the only correct outcome is to recommend no combination for phase II. Scenario 15 represents a situation where all combinations are true MTCs, and is used to monitor escalation behaviour when combinations are safe and increasing the dose of either drug does not affect toxicity.

In order to accentuate the differences in the designs, we do not implement any accuracy or sufficient information rules, as these may mask some key elements of the designs. We focus on the operating characteristics of proportion of correct selections (PCS) and proportion of acceptable selections (PAS) as measures of accuracy, and proportion of overly toxic selections and the number of patients treated on unsafe dose combinations as measures of safety.

Model-Based Comparators

To provide a comparison between model-free and model-based designs, we also consider conventional model-based approaches in our simulation study, the two-dimensional Bayesian Logistic Regression Model (BLRM)¹⁸, the Partial Ordering Continual Reassessment Method (POCRM)²⁸, and a modified design based on the logistic model (referred to here as the Riviere design)²⁴.

The same proposed calibration procedure as is applied to the model-free designs is applied to the model-based, with details provided in the online supplementary materials. Note that the form of the overdosing rule may be different for the model-based designs, compared the model-free designs, as described in their respective original proposals.

Bayesian Logistic Regression Model In this approach, the toxicity probability for each combination, π_{ij} , are modelled as in Equation 10 for $i = 1, \dots, I$ and $j = 1, \dots, J$, where doses d_i^A and d_j^B are scaled by reference doses. Let d_{ij} be combination of d_i^A and d_j^B , while n_{ij} and y_{ij} are the number of patients and toxic responses on each combination respectively. Parameters α_1 and β_1 describe

the toxicity of drug A, α_2 and β_2 describe the toxicity of drug B, and η models the interaction between drugs. The five parameters are assigned normal prior distributions, and the likelihood is a product of Bernoulli densities, proportional to $\prod_{i=1}^I \prod_{j=1}^J \pi_{ij}^{y_{ij}} (1 - \pi_{ij})^{n_{ij} - y_{ij}}$. After each cohort is observed, the joint posterior distribution is approximated using MCMC methods, and samples of each parameter are drawn from their full conditional distributions. Estimates of π_{ij} are made by sampling parameters from their posteriors and substituting these along with the corresponding doses into Equation 10. Note that all parameters except η are sampled on the log scale and then exponentiated since they must be positive.

$$\pi_{ij}(\alpha_1, \alpha_2, \beta_1, \beta_2, \eta | d_i^A, d_j^B) = \frac{[\alpha_1 (d_i^A)^{\beta_1} + \alpha_2 (d_j^B)^{\beta_2} + \alpha_1 \alpha_2 (d_i^A)^{\beta_1} (d_j^B)^{\beta_2}] \exp(\eta d_i^A d_j^B)}{1 + [\alpha_1 (d_i^A)^{\beta_1} + \alpha_2 (d_j^B)^{\beta_2} + \alpha_1 \alpha_2 (d_i^A)^{\beta_1} (d_j^B)^{\beta_2}] \exp(\eta d_i^A d_j^B)} \quad (9)$$

The BLRM can only escalate to combinations satisfying the neighbourhood constraint and Escalation With Overdose Control (EWOC) principle. The neighbourhood constraint prevents escalation or de-escalation to any combination that is more than one dose level of either drug away, and also prevents escalation to a combination in which both dose levels are higher. For a trial with target toxicity $\phi = 0.30$, the EWOC principle states that d_{ij} can only be administered if $\mathbb{P}(\pi_{ij} > 0.33) < \epsilon_{\text{BLRM}}$. The combination maximising the probabilistic statement $\mathbb{P}(0.16 < \pi_{ij} < 0.33)$ is administered to the next cohort. If no combinations satisfy the two constraints, the trial is terminated. Once the sample size has been exhausted, the MTC is selected from combinations which have been experimented on with at least six patients, and is the one maximising $\mathbb{P}(0.16 < \pi_{ij} < 0.33)$. The BLRM requires dosing quantities for each drug to be specified, in all of the implementations of the BLRM, these doses are 100, 200 and 300mg for each drug.

Partial Ordering Continual Reassessment Method The Bayesian Partial Ordering Continual Reassessment Method (POCRM) ²⁸ generalises the original CRM design to the setting of combination trials. The POCRM design assumes that there are R feasible orderings of the combinations satisfying the monotonicity assumption within each agent. Let k be the index of the combination, $k = 1, \dots, K$, r be the index of ordering, $r = 1, \dots, R$, q_{kr} be the standardised regimen level at combination k under the ordering r , and let π_{kr} be the corresponding probability of a DLT. Then, the combination-toxicity model takes the form

$$\pi_{kr} = q_{kr}^{\exp(\alpha_r)}. \quad (11)$$

where α_r is the (scalar) model parameter under the ordering r that has a normal prior distribution, $\mathcal{N}(\mu, \sigma^2)$. The working models q_{kr} are constructed from standardised values, \tilde{q}_k (also known as skeleton) by re-ordering them according to the order r . The original design proposal included the possibility of the early stopping if the lower bound of the η -% approximated confidence interval ¹⁹ is above the target toxicity level, ϕ . We refer the reader to the original publication for further technical details.

Logistic Model by Riviere et.al (2014) Another model-based approach considered is the modified logistic model ²⁴. Following the recommendation on reducing the dimensionality of the parametric models in Phase I trials ⁷, we consider the 3-parameter logistic model (rather than the original 4-parameter one ¹⁶) and was found to result in the same or better, on average, operating characteristics of the design for small to moderate sample sizes.

Specifically, the combination-toxicity is modelled using the 3-parameter logistic model

$$\text{logit}(\pi_{j,k}) = \beta_0 + \beta_1 u_j + \beta_2 v_k \quad (12)$$

where $\beta_0 \in \mathbb{R}$, $\beta_1 > 0$, $\beta_2 > 0$ are unknown parameters that denote the intercept (β_0) the toxicity effect of agent

1 (β_1) and agent 2 (β_2), and $u_j = \log\left(\frac{p_j}{1-p_j}\right)$ and $v_k = \log\left(\frac{q_k}{1-q_k}\right)$ are the standardised doses of the two agents. The parameters p_1, \dots, p_j , and q_1, \dots, q_k , being the prior estimates of the toxicity probabilities for the dose levels of agent 1 and 2, respectively, when administered as monotherapies. The terms u_j and v_k are also known as the skeleton and are unchanged throughout the trial. The escalation rules follow the ones originally used²⁴. The safety constraint for the early stopping is taking the form of if $\mathbb{P}(\pi_{1,1} > \phi) > \epsilon_{\text{logistic}}$ is satisfied then the trial is stopped earlier.

A Non-Parametric Optimal Benchmark Comparator

While the primary goal of this work is to compare the performance of different model-free designs to each other, there is a risk that all methods might perform equally poorly on some scenarios. In this case, the comparison of the designs to each other would not identify why the poor performance is observed – due to the challenging scenario or due to all designs having difficulties identifying a particular MTC. To provide context for the comparison of operating characteristics, we include the performance of the non-parametric benchmark for combination studies, a tool that provides an estimate for the upper bound on the PCS under the given combination-toxicity scenario^{15,17}. The benchmark takes into account the “difficulty” of a scenario in terms of how close the toxicity risks for the combinations (under this scenario) are to the target level of 30%, and also accounts for the unknown monotonic ordering in the combination setting. We refer the reader to the recent work by Mozgunov et al.¹⁷ for further technical details on the benchmark for combinations implementation.

Results

Accuracy Index and Proportions of Correct and Acceptable Selections The results are presented here in terms of

proportions of correct and acceptable selections (PCS and PAS) as defined in the ‘Calibration of Designs’ section. We also calculate an Accuracy Index, defined in equation (13), where ρ_{ij} is the proportion of simulations that select dose d_{ij} as the MTD (see Hirakawa et al.⁹ for further details). Table 2 presents these results.

$$A_n = 1 - I \times J \times \frac{\sum_{i=1}^I \sum_{j=1}^J |\pi_{ij} - \phi| \times \rho_{ij}}{\sum_{i=1}^I \sum_{j=1}^J |\pi_{ij} - \phi|} \quad (13)$$

Figure 4 presents the summary of the operating characteristics of the considered designs in terms of the PCS and PAS (with the full set of results given in the online supplementary materials). Model free designs are shown in blue, model-based designs in purple, and the non-parametric benchmark is in black. Scenarios 14 and 15 have been excluded as these have no true MTCs for the design to select. For scenarios in which the only acceptable combinations are also correct combinations (Scenarios 6, 9, 10, 11 and 13), the PCS and PAS are equal. The mean PCS across Scenarios 1-13 for the BOIN, KEY, SFD, PIPE, Waterfall, BLRM, POCRM and Riviere designs is 39.8%, 42.4%, 41.6%, 31.2%, 32.3%, 40.0%, 32.2% and 48.3% respectively, whilst the mean PAS are 58.7%, 62.1%, 59.0%, 56.0%, 53.4%, 58.4%, 44.4% and 64.6% respectively.

[Figure 4 about here.]

[Table 2 about here.]

First of all, the benchmark reveals the differences in how challenging it is to identify the MTC in the considered scenarios: the PCS for the benchmark varies between approximately 35% under Scenario 7 to more than 80% under Scenario 13. As expected, the benchmark corresponds to the highest average PCS and PAS - 55% and nearly 70%, respectively. Similarly, under the majority of scenarios the benchmark corresponds to the highest

PCS and PAS as it employs the concept of the complete information. The largest difference between the benchmark and other designs can be seen under Scenario 13. At the same time, there are scenarios under which the benchmark is outperformed by a competing design - this can be a sign of the design favouring particular combinations under the calibrated priors - for example under Scenario 7. Since the aim of the calibration procedure was to obtain a prior with good operating characteristics across many plausible scenarios, in the simulation study some scenarios will have better or worse performance than the benchmark.

The variety of performances across the scenarios demonstrates the variability between the different designs in different settings. Considering the model-free designs, on average the KEY design has the highest proportion of both correct and acceptable selections, but is vastly outperformed in some scenarios by the SFD design. In five of the scenarios, the KEY has the highest PCS out of all the model-free designs, being superior in scenarios with few overly toxic combinations. However, for example in Scenario 11, where the MTC is the middle dose of drug A and lowest dose of drug B, the SFD outperforms the next best performing design by 19.2%. The PIPE design shows poor performance in many scenarios, most notably in Scenario 1 where the PCS is 5.5% and PAS is 54.0%. A likely reason is that for the PIPE design, the choice of MTC must be below the MTC contour, and a scenario where the true MTC is the highest dose combination gives rise to underestimation since we cannot explore above the true MTC contour. In addition, the procedure discussed in the Methodological Review of the PIPE design to choose one MTC from the recommended set will make our results differ from those originally reported by Mander and Sweeting¹², where a ‘correct selection’ was defined as the MTC being in the set of recommended doses. Although the Waterfall design also recommends a set of doses, the performance in the simulations is better than that of the PIPE design.

Although it has poor performance in Scenario 1, where the MTC is the highest dose combination, it is the best performing design in Scenarios 9 and 13.

When considering the model-based designs as a comparator, we see that in many scenarios these outperform the KEY. For example, in Scenario 1 where the MTC is the highest combination, all of the model based designs achieve a PCS at least 20% higher than the next best performing design, the KEY. In fact, when including these designs in the comparison, the KEY is only the best performing design in one scenario, Scenario 8. The SFD does however outperform the model-based designs in some cases, with the model-based designs having the highest PCS in Scenarios 1, 2, 3, 5, 7 and 11 and the SFD is the best performing in Scenarios 6, 10 and 12.

In terms of the Accuracy Index, the SFD, BLRM and Riviere designs show the highest mean value, supporting the collective evidence that these are the most accurate designs in selecting the MTC.

Proportions of Overly Toxic Selections Figure 5 illustrates the proportion of overly toxic selections for each design. Scenarios 1 and 15 have no overly toxic combinations, so the proportion is zero for these cases. Model-free designs showed lower proportions of overly toxic selection than model-based designs in many scenarios. We observe that the POCRM recommends the most overly toxic combinations on average by far, in 36.8% of trials. It stands out in multiple scenarios with a very high percentage of simulated trials recommending overly toxic doses. In 10 scenarios this is over 30%, highlighting how aggressive this approach is.

[Figure 5 about here.]

Of the model-free designs, the SFD has the highest percentage, in 20.4% of trials. In scenarios 5, 9 and 12, it recommends overly toxic combinations in over 25% of the

simulated trials and in 9 of the scenarios, it is the model-free design with the highest proportion of overly toxic recommendations. This is evidence of the trade-off between selecting combinations close to ϕ and the willingness to recommend more overly toxic combinations.

The PIPE design demonstrates a very low proportion of overly toxic selections with a mean of 9.2% across the 13 scenarios, 6.2% below any of the other designs. It has the lowest in all but four scenarios. This is a further illustration of the feature of the design to recommend combinations near but lower than the estimated MTC contour.

A focus on Scenario 14, where all dose combinations are overly toxic, shows the BLRM is the most efficient at stopping for safety, with 93.7% of simulations not recommending any dose combination. Noticeably, the Waterfall design is the least efficient at stopping for safety in this scenario.

Number of Patients Treated at Overly Toxic Combinations

Figure 6 outlines the mean number of patients treated at overly toxic combinations in Scenarios 1-15 for each design. Note that we report the number rather than proportion of patients, as this will also give insight into how effectively each design stops for safety.

[Figure 6 about here.]

The most notable feature of these results are the large number of patients treated at overly toxic combinations by the model-based designs. This aggressive escalation is driven by the informative prior, calibrated to give high values of PCS. We refer the reader to the online supplementary materials where an alternative BLRM prior leading to more conservative escalation (but considerably lower PCS and PAS) is explored.

The SFD, KEY and BOIN have reasonable performance, with the Waterfall design showing a strong performance with the lowest overall mean number of patients treated on overly toxic doses of five patients.

Careful attention must again be paid to Scenario 14, where all dose combinations are overly toxic. The PIPE design treats an average of 20 patients per trial, over six cohorts, which is an unacceptable level of exploration in such a scenario. In this scenario, we also consider that although the BLRM showed good performance in stopping early for safety in the highest number of simulated trials, it also has a high number of patients treated on average before stopping. The Waterfall design has the fewest patients treated by far, with an average of five patients, but this leads to the erroneous recommendation of a dose in a large proportion of simulated trials.

We see that overall the model-free approaches are more conservative in their escalation than the model-based designs, with fewer patients treated on unsafe doses, with no noticeable increase in PCS. Of the model-free approaches, the SFD shows the most promising PCS over the different scenarios, at the cost of somewhat higher overly toxic selections. It is also worth noting that the SFD has a substantially higher computational cost than the other model-free designs.

Simulation Study: Alternative Dosing Grids

All of the simulations so far have concerned a dosing grid whereby there are three levels of each drug. However, it may alternatively be the case where there are differing numbers of dose levels for the two drugs. The most common being two levels of one drug and three or four of the other. Therefore in this section, we consider two alternative dosing grid sizes: 2×3 and 2×4 .

Setting

Six additional scenarios are considered for the alternative dosing grids, outlined in Table 3. Scenarios 16-18 have a 2×3 dosing grid and are chosen to be similar to scenarios 10, 9 and 6, but without the highest dose of drug A in order to explore the behaviour of the smaller grid. Scenarios 19-21

have a 2×4 dosing grid, with Scenarios 19 and 21 having an MTD at the fourth dose level of drug B and Scenario 20 similar to Scenario 18, but with the fourth dose level of drug B unsafe.

The priors and intervals used in each design are equivalent to those used for the 3×3 dosing grid, meaning that hyper-parameter values are the same. The cohort size remains three patients, with a maximum of 12 cohorts. For the Waterfall design, 8 cohorts are allocated to the first sub-trial and 4 cohorts to the second sub-trial.

[Table 3 about here.]

Results

We again compare the designs' operating characteristics in terms of performance of selecting correct and acceptable dose combinations, and safety in terms of patients allocated to overly toxic doses and selection of overly toxic doses.

Table 4 gives the accuracy index for the designs across the six scenarios, calculated using Equation 13. It shows a similar trend to the 3×3 dosing grid, with the SFD giving the best overall performance of the model-assisted designs, and PIPE giving the poorest performance, let down by scenarios where the MTC is high in the dosing grid.

[Table 4 about here.]

Figure 7 shows the percentage of correct and acceptable selections for the alternative dosing grid scenarios. In scenarios 16 and 17, the benchmark illustrates that these 2×3 dosing grid scenarios are less difficult than their 3×3 comparators. In Scenario 16, all designs perform at least as well as in the 3×3 case with the SFD performing the best in Scenario 16 and the Keyboard and Waterfall performing the best in Scenario 17, however in Scenario 17, the PIPE design performs slightly worse. PIPE performs considerably worse in Scenario 18, where it is evident that in the 3×3 grid, dose d_{31} was frequently chosen, and when the third

dose level of drug *A* is removed, the other MTCs are not chosen as frequently. The benchmark shows that the 2×4 dosing grid scenarios are more challenging than the other sizes of grid, however most designs give a high level of PCS, apart from the PIPE design. However, the PIPE design does give the highest levels of PAS in the 2×4 dosing grid scenarios.

Figures 8 and 9 show the percentage of overly toxic selections and the number of patients treated at overly toxic combinations for scenarios 16-21 respectively. There is a clear trend that the model-based designs have a higher percentage of overly toxic selections, with the PIPE design showing a very low percentage, substantially lower than all designs apart from in Scenario 21, where it has the same level as the Waterfall design. There is a similar trend in the mean number of patients treated on overly toxic doses, with the PIPE design showing the smallest average overall, and POCRM and Riviere showing the highest. Scenario 20 has many more patients treated on unsafe doses than Scenario 18 in all but the SFD, PIPE and BLRM designs. This indicates that for the other five designs, escalation to the fourth dose level of drug *B* was restricted by the grid rather than the design itself.

[Figure 7 about here.]

[Figure 8 about here.]

[Figure 9 about here.]

Case Study

The simulation studies gave insight into the operating characteristics of each design, however for further insight into the escalation behaviour, we apply each method to an example case study. We consider a phase I oncology (breast and lung cancer) study enrolling patients to dosing combinations of four dose levels of neratinib and temsirolimus⁶. A total sample size of 60 patients (cohorts

of size 2 or 3) were treated on 12 of 16 possible dosing combinations. Results from 52 patients were included and 10 DLTs were observed, with full results of the trial displayed in Table 5.

In order to strengthen the case that our conclusions are generalizable with less risk of selective reporting, the purpose of this case study is not to investigate whether each design chooses the same MTC as the real study did. The purpose is to give an illustration of how each design explores the dosing grid, given identical patient responses.

In order to use the calibrated prior specifications, and in line with the simulation study, we restrict the dosing grid to three doses of each drug, removing the lowest dose of temsirolimus and the highest dose of neratinib. We also fix the cohort size to three patients and maximum total sample size to 36.

To ensure a fair comparison between designs, we define a fixed set of 36 ordered patient responses for each dose combination. The first patient responses in this set are the true y_{ij} DLT responses and $n_{ij} - y_{ij}$ non-DLT responses, in a random permutation (note that this is the same random permutation for each of the methods). The remaining $36 - n_{ij}$ responses are generated in the following way. Each patient has an individual probability of DLT, generated from $\text{Beta}(1 + y_{ij}, 1 + n_{ij} - y_{ij})$. Then a binary response is generated with this probability. Where there were no patients assigned to the dose combination in the real study, the individual $P(\text{DLT})$ is generated from a $\text{Beta}(3,3)$ distribution, to indicate the dose combination is unsafe, since this is the reason the combination was not escalated to. This process uses the information from the real study, but also introduces enough variability in the subsequent responses to account for the small sample size.

[Table 5 about here.]

Table 5 displays the results of each of the methods, with the number of patients treated at each combination,

the number of DLTs observed, and the concluded MTC highlighted in bold.

The BOIN and KEY designs show very similar exploration, first escalating in neratinib, then temsirolimus. The highest combination is not explored, as the combinations with the next lowest dose of each drug were considered unsafe. The only difference is that the KEY assigns one more cohort to the 200mg/50mg combination, even when the previous cohort had 2/3 observed DLT responses.

The PIPE design explores differently, not escalating to the highest dose of temsirolimus at all, even though only 1/12 DLT responses were observed on the 160mg/50mg combination. The SFD explores more of the highest dose of temsirolimus, although still not the highest combination. An interesting observation here is that the final recommended dose has observed 4/9 DLT responses, a level that would generally be an unsafe standard. This is in line with the simulation results that showed this design to have the highest level of overly toxic selections. The Waterfall design explores the entire dosing grid apart from the highest dose combination 200mg/75mg, a more even spread of patients across the grid than the other model-free designs.

The BLRM is executed with two prior distributions, the calibrated prior and the alternative, more realistic based on safety concerns, prior. Surprisingly, both show a more aggressive escalation than the model free designs, with patients allocated to the highest combination. The calibrated prior gives the most aggressive approach with a second cohort assigned to a dose, even when the first observed 2/3 DLT responses. This also means that for this prior, the dosing grid is not as well explored as some of the model-free designs, as the lowest dose of temsirolimus is only explored in combination with the lowest dose of neratinib. For the calibrated prior, these results are in line with the simulation study, where the BLRM had on average

the most patients treated on overly toxic doses and also a high proportion of overly toxic recommendations. However, even the alternative prior shows more aggressive escalation than the model-free designs in this case study. Both the POCRM and the Riviere design had a balanced exploration of the dose grid.

The case study highlights some key differences in the approaches, illustrating how both the escalation schemes and final recommendation differ. Particularly of note is the somewhat aggressive behaviour of not de-escalating when observing 2/3 observed DLT responses, and recommending a final dose combination with 4/9 observed DLT responses from both the SFD and calibrated prior BLRM. This behaviour, that could be considered unsafe, is not necessarily obvious from simulation results and underlines the importance of studying the individual escalations in an example case study. It is also important to consider that in practice, such a statistical approach is a guidance for dose recommendation that should be supported by an overall evaluation of the safety, pharmacokinetics and clinical rationale.

Discussion

This paper provides a review of a wide range of combination designs in phase I oncology, exploring the more recently proposed model-free designs in detail, as well as providing a novel approach for the calibration of such designs. The comprehensive simulation study we conduct suggests that model-free designs are competitive with the model-based designs in terms of the proportion of correct combinations selected. The operating characteristics of model-free designs in a number of scenarios suggest they offer a safer alternative. The case study example highlighted the key differences in how the methods explore the dosing grid given the same patient responses, with more aggressive approaches missing the lower doses, and conservative approaches missing the higher ones.

The discussed results depend upon the specification of the intervals for the BOIN, KEY and Waterfall designs, and the operational priors for the PIPE, SFD, POCRM, Riviere and BLRM designs, which were calibrated using a novel approach. This included calibrating the overdosing rules in each design to reduce the risk of recommending overly toxic combinations for phase II. Naturally, our work does not allow for comparison between designs when complete and reliable prior information on the toxicity of each drug is available. In practice, the PIPE, SFD, POCRM, Riviere and BLRM designs can exploit this prior knowledge to help the escalation process.

The calibration procedure, although novel in approach, is relatively straightforward to implement. It does however highlight the computational intensity of the different methods. Both the BLRM and SFD are very computationally intensive, with the calibration procedure taking substantially longer than for any of the other designs. It has shown great promise in specifying prior distributions that yield high PCS values, removing the subjectivity from the specification.

Moreover, our simulations do not allow for the early selection of an MTC. For example, if at least 9 patients are treated at a combination and the next cohort is recommended to be treated at this combination, then a trial could be stopped and this combination selected as the MTC. We acknowledge this rule is useful to reduce sample sizes, especially in scenarios where the true MTC is a low dose combination. One of the advantages of model-based approaches is that they allow for selection of unplanned intermediate doses. This is an advantage that was not used in the simulation study, but must be considered in practice.

An additional area of interest for such dose-finding studies is the sample size and cohort size. Conducting a sensitivity analysis on both of these for each design would be an excellent opportunity to investigate whether designs can still achieve high PCS with fewer patients, or

significantly higher PCS with extra patients, and whether a larger or smaller cohort size would lead to better exploration of the dosing grid.

Finally, we conclude this comparison with an overview of recommendations for the use of each design in the context of this work. The BOIN and KEY designs give a balanced approach, with a good level of PCS and PAS across a range of scenarios. Overly toxic explorations and selections are also well balanced across scenarios. The PIPE design is more cautious in its selection, with a consistently low proportion of overly toxic selections, although at the cost of also recommending correct combinations a lower proportion of the time. The Surface Free design offers a high PCS and PAS and a generally low number of patients treated at overly toxic selections, but this must be balanced with the high proportion of overly toxic selections. The Waterfall design is most cautious in its allocation of patients, with a similar level of overly toxic recommendations as KEY and BOIN. However, the overall PCS and PAS are somewhat lower than the other designs. The model-based designs provide the most aggressive approach with a calibrated prior, with a large number of patients treated on overly toxic doses, however a good level of PCS and PAS. For the BLRM, with an alternative, intuitive prior, the number of overly toxic explorations is reduced, but at the cost of the high PCS values.

Data Availability Statement

The data that supports the findings of this research are available in Table 5 of this article, originally from Gandhi et al.⁶, with all other data simulated according to the specifications described.

Acknowledgements

This report is independent research supported by the National Institute for Health Research (NIHR Advanced Fellowship, Dr Pavel Mozgunov,

NIHR300576; and Prof Jaki's Senior Research Fellowship, NIHR-SRF-2015-08-001) and by the NIHR Cambridge Biomedical Research Centre (BRC-1215-20014). The views expressed in this publication are those of the authors and not necessarily those of the NHS, the National Institute for Health Research or the Department of Health and Social Care (DHSC). T Jaki and P Mozgunov received funding from UK Medical Research Council (MC_UU_00002/14).

References

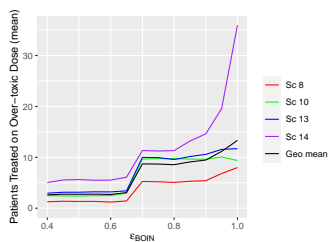
1. Rachid Abbas, Caroline Rossoni, Thomas Jaki, Xavier Paoletti, and Pavel Mozgunov. A comparison of phase I dose-finding designs in clinical trials with monotonicity assumption violation. *Clinical Trials*, 17(5):522–534, 2020.
2. Mark R Conaway and Gina R Petroni. The impact of early-phase trial design in the drug development process. *Clinical Cancer Research*, 25(2):819–827, 2019.
3. Jennifer Couzin-Frankel. Cancer Immunotherapy. *Science*, 342(6165):1432–1433, 2013.
4. Richard L Dykstra and Tim Robertson. An algorithm for isotonic regression for two or more independent variables. *The Annals of Statistics*, 10(3):708–716, 1982.
5. Shenghua Kelly Fan, Alan P Venook, and Ying Lu. Design issues in dose-finding phase I trials for combinations of two agents. *Journal of biopharmaceutical statistics*, 19(3):509–523, 2009.
6. Leena Gandhi, Rastislav Bahleda, Sara M. Tolaney, Eunice L. Kwak, James M. Cleary, Shuchi S. Pandya, Antoine Hollebecque, Richat Abbas, Revathi Ananthakrishnan, Anna Berkenblit, Mizue Krygowski, Yali Liang, Kathleen W. Turnbull, Geoffrey I. Shapiro, and Jean Charles Soria. Phase I study of neratinib in combination with temsirolimus in patients with human epidermal growth factor receptor 2-dependent and other solid tumors. *Journal of Clinical Oncology*, 32(2):68–75, 2014.
7. Mauro Gasparini. General classes of multiple binary regression models in dose finding problems for combination therapies. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 62(1):115–133, 2013.
8. Paul Hamberg and Jaap Verweij. Phase I Drug Combination Trial Design: Walking the Tightrope. *Journal of Clinical Oncology*, 27(27):4441–4443, 2009. PMID: 19704054.
9. Akihiro Hirakawa, Nolan A. Wages, Hiroyuki Sato, and Shigeyuki Matsui. A comparative study of adaptive dose-finding designs for

- phase i oncology trials of combination therapies. *Statistics in Medicine*, 34(24):3194–3213, 2015.
10. Anastasia Ivanova and Kai Wang. A non-parametric approach to the design and analysis of two-dimensional dose-finding trials. *Statistics in Medicine*, 23(12):1861–1870, 2004.
 11. Ruitao Lin and Guosheng Yin. Bayesian optimal interval design for dose finding in drug-combination trials. *Statistical Methods in Medical Research*, 26(5):2155–2167, 2017.
 12. Adrian P Mander and Michael J Sweeting. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Statistics in medicine*, 34(8):1261–1276, 2015.
 13. Pavel Mozgunov, Mauro Gasparini, and Thomas Jaki. A surface-free design for phase I dual-agent combination trials. *Statistical Methods in Medical Research*, 29(10):3093–3109, 2020.
 14. Pavel Mozgunov and Thomas Jaki. 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(2):347, 2019.
 15. Pavel Mozgunov, Thomas Jaki, and Xavier Paoletti. A benchmark for dose finding studies with continuous outcomes. *Biostatistics*, 21(2):189–201, 2020.
 16. Pavel Mozgunov, Rochelle Knight, Helen Barnett, and Thomas Jaki. Using an interaction parameter in model-based phase i trials for combination treatments? a simulation study. *International journal of environmental research and public health*, 18(1):345, 2021.
 17. Pavel Mozgunov, Xavier Paoletti, and Thomas Jaki. A benchmark for dose-finding studies with unknown ordering. *Biostatistics*, 2021.
 18. Beat Neuenschwander, Alessandro Matano, Zhongwen Tang, Satrajit Roychoudhury, Simon Wandel, and Stuart Bailey. A Bayesian industry approach to phase I combination trials in oncology. In *Statistical methods in drug combination studies*, pages 95–135. Chapman & Hall/CRC Press: Boca Raton, FL, 2015.
 19. John O’Quigley and Mark Conaway. Continual reassessment and related dose-finding designs. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 25(2):202, 2010.
 20. Haitao Pan, Ruitao Lin, Yanhong Zhou, and Ying Yuan. Keyboard design for phase i drug-combination trials. *Contemporary Clinical Trials*, 92:105972, 2020.
 21. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2020.
 22. M-K Riviere, F Dubois, and S Zohar. Competing designs for drug combination in phase I dose-finding clinical trials. *Statistics in*


- medicine*, 34(1):1–12, 2015.
23. M-K Riviere, C Le Tourneau, X Paoletti, F Dubois, and Sarah Zohar. Designs of drug-combination phase I trials in oncology: a systematic review of the literature. *Annals of Oncology*, 26(4):669–674, 2015.
 24. Marie-Karelle Riviere, Ying Yuan, Frédéric Dubois, and Sarah Zohar. A bayesian dose-finding design for drug combination clinical trials based on the logistic model. *Pharmaceutical statistics*, 13(4):247–257, 2014.
 25. Padmanee Sharma and James P Allison. Immune checkpoint targeting in cancer therapy: toward combination strategies with curative potential. *Cell*, 161(2):205–214, 2015.
 26. Peter F Thall and S-J Lee. Practical model-based dose-finding in phase I clinical trials: methods based on toxicity. *International Journal of Gynecologic Cancer*, 13(3):251–261, 2003.
 27. Peter F Thall, Randall E Millikan, Peter Mueller, and Sang-Joon Lee. Dose-finding with two agents in phase I oncology trials. *Biometrics*, 59(3):487–496, 2003.
 28. Nolan A Wages, Mark R Conaway, and John O’Quigley. Continual reassessment method for partial ordering. *Biometrics*, 67(4):1555–1563, 2011.
 29. Nolan A Wages, Anastasia Ivanova, and Olga Marchenko. Practical designs for phase I combination studies in oncology. *Journal of Biopharmaceutical Statistics*, 26(1):150–166, 2016.
 30. Chi Heem Wong, Kien Wei Siah, and Andrew W Lo. Estimation of clinical trial success rates and related parameters. *Biostatistics*, 20(2):273–286, 2019.
 31. Kit Man Wong, Anna Capasso, and S Gail Eckhardt. The changing landscape of phase I trials in oncology. *Nature Reviews Clinical oncology*, 13(2):106–117, 2016.
 32. Fangrong Yan, Sumithra J Mandrekar, and Ying Yuan. Keyboard: a novel Bayesian toxicity probability interval design for phase I clinical trials. *Clinical Cancer Research*, 23(15):3994–4003, 2017.
 33. Yanhong Zhou, J Jack Lee, Shunguang Wang, Stuart Bailey and Ying Yuan. Incorporating historical information to improve phase I clinical trials. *Pharmaceutical Statistics*, 20(6):1017-1034, 2021.
 34. Liangcai Zhang and Ying Yuan. A practical Bayesian design to identify the maximum tolerated dose contour for drug combination trials. *Statistics in Medicine*, 35(27):4924-4936, 2016.

d_4^A	×	×	×	×
d_3^A	*	○	×	×
d_2^A	○	●	○	×
d_1^A	*	○	*	×
	d_1^B	d_2^B	d_3^B	d_4^B


Figure 1. Illustration for $I = J = 4$ of the admissible combinations for each design. The '●' symbol illustrates the current dose combination at d_{22} , the symbol '○' represents dose combinations achieved through escalation/de-escalation in one drug only and the symbol '*' represents dose combinations achieved through diagonal de-escalation or anti-diagonal escalation.




(a) BOIN & Waterfall: Patients treated at overly toxic doses




(b) BOIN & Waterfall: Proportion of correct selections




(c) KEY: Patients treated at overly toxic doses




(d) KEY: Proportion of correct selections




(e) Surface-Free: Patients treated at overly toxic doses



(f) Surface-Free: Proportion of correct selections



(g) PIPE: Patients treated at overly toxic doses



(h) PIPE: Proportion of correct selections

Figure 2. Calibration of ϵ for the five designs

		Drug B		
		d_1^B	d_2^B	d_3^B
Drug A	d_1^A	0.05	0.075	0.10
	d_2^A	0.075	0.10	0.125
	d_3^A	0.10	0.125	0.15

Figure 3. Constructing prior mean toxicity probabilities when calibrating the PIPE design.

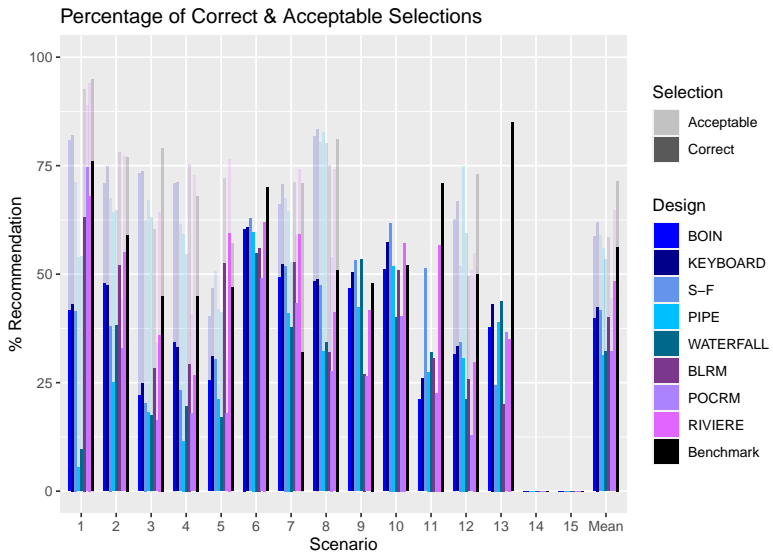


Figure 4. An illustration of the PCS and PAS for Scenarios 1-13 for each design. The solid bars measure the PCS and the more transparent bars measure the PAS. The rightmost group of bars show the means.

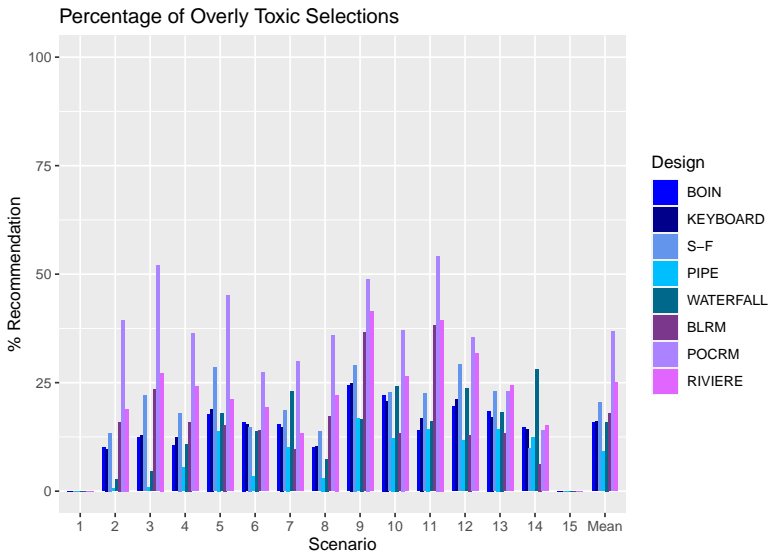


Figure 5. An illustration of the proportion of overly toxic selections across Scenarios 1-15 for each design. The rightmost group of bars show the means.




Figure 6. An illustration of the number of patients treated at overly toxic combinations during trials in Scenarios 1-15 for each design. The rightmost group of bars show the means.

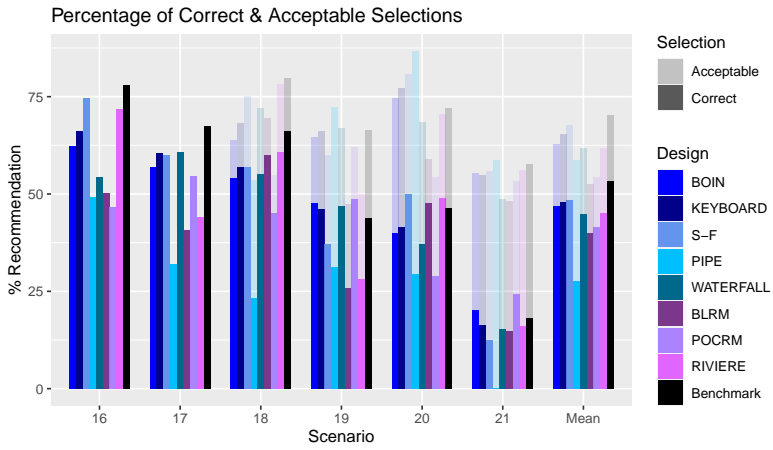


Figure 7. An illustration of the PCS and PAS for Scenarios 16-21 for each design. The solid bars measure the PCS and the more transparent bars measure the PAS. The rightmost group of bars show the means.

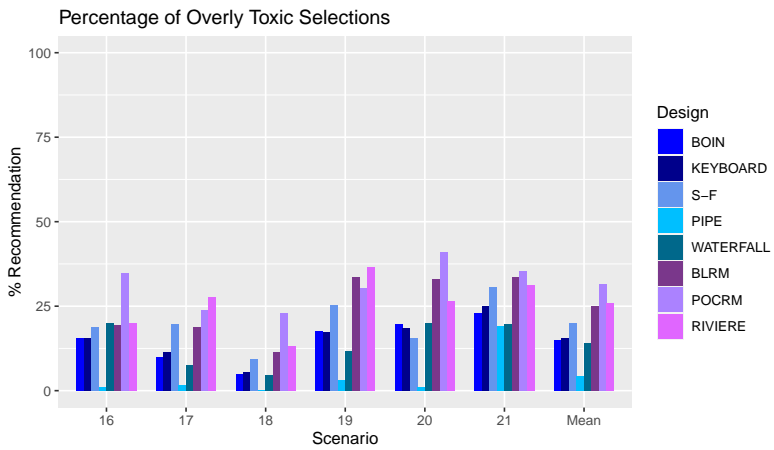


Figure 8. An illustration of the percentage of overly toxic selections across Scenarios 16-21 for each design. The rightmost group of bars show the means.

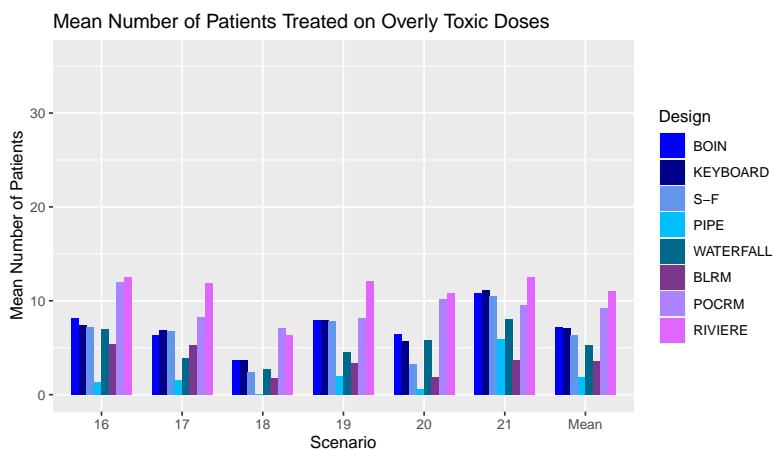


Figure 9. An illustration of the number of patients treated at overly toxic combinations during trials in Scenarios 16-21 for each design. The rightmost group of bars show the means.

Scenario 1				Scenario 2				Scenario 3			
	d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B
d_1^A	0.05	0.10	0.15	d_1^A	0.05	0.10	0.15	d_1^A	0.02	0.05	0.10
d_2^A	0.10	0.15	<u>0.20</u>	d_2^A	0.10	<u>0.20</u>	0.30	d_2^A	0.10	0.15	<u>0.20</u>
d_3^A	0.15	<u>0.20</u>	0.30	d_3^A	<u>0.20</u>	0.30	0.45	d_3^A	<u>0.20</u>	0.30	0.45
Scenario 4				Scenario 5				Scenario 6			
	d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B
d_1^A	0.05	0.10	0.15	d_1^A	0.02	0.05	0.15	d_1^A	0.10	0.15	0.30
d_2^A	0.10	<u>0.20</u>	0.30	d_2^A	<u>0.20</u>	0.30	0.45	d_2^A	0.15	0.30	0.45
d_3^A	<u>0.20</u>	0.45	0.60	d_3^A	0.45	0.55	0.65	d_3^A	0.30	0.45	0.60
Scenario 7				Scenario 8				Scenario 9			
	d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B
d_1^A	0.10	<u>0.20</u>	0.45	d_1^A	0.05	0.10	<u>0.20</u>	d_1^A	0.10	0.15	<u>0.20</u>
d_2^A	0.15	0.30	0.50	d_2^A	0.10	<u>0.20</u>	0.30	d_2^A	0.30	0.40	0.50
d_3^A	0.30	0.50	0.60	d_3^A	0.30	0.45	0.55	d_3^A	0.40	0.50	0.60
Scenario 10				Scenario 11				Scenario 12			
	d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B
d_1^A	0.15	0.30	0.45	d_1^A	0.02	0.05	0.10	d_1^A	<u>0.20</u>	0.30	0.45
d_2^A	0.30	0.45	0.55	d_2^A	0.30	0.45	0.60	d_2^A	0.45	0.50	0.55
d_3^A	0.45	0.55	0.65	d_3^A	0.45	0.60	0.75	d_3^A	0.65	0.70	0.75
Scenario 13				Scenario 14				Scenario 15			
	d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B		d_1^B	d_2^B	d_3^B
d_1^A	0.30	0.45	0.50	d_1^A	0.45	0.50	0.55	d_1^A	0.10	0.10	0.10
d_2^A	0.45	0.50	0.55	d_2^A	0.50	0.55	0.60	d_2^A	0.10	0.10	0.10
d_3^A	0.50	0.55	0.60	d_3^A	0.55	0.60	0.65	d_3^A	0.10	0.10	0.10

Table 1. Toxicity scenarios to evaluate the combination designs. Rows and columns refer to the dose of drug A and B respectively. True MTCs are in bold and 'acceptable' combinations are underlined.

Scenario	BOIN	KEYBOARD	S-F	PIPE	Waterfall	BLRM	POCRM	Riviere
1	0.538	0.541	0.489	0.003	0.205	0.728	0.790	0.767
2	0.484	0.482	0.375	0.042	0.352	0.538	0.251	0.568
3	0.394	0.422	0.345	0.148	0.287	0.435	0.225	0.482
4	0.487	0.477	0.362	0.207	0.326	0.462	0.156	0.437
5	0.416	0.461	0.472	0.362	0.358	0.653	0.180	0.699
6	0.558	0.528	0.578	0.343	0.461	0.493	0.364	0.562
7	0.535	0.549	0.556	0.384	0.378	0.579	0.420	0.611
8	0.539	0.541	0.515	0.258	0.415	0.390	0.178	0.449
9	0.490	0.525	0.603	0.339	0.533	0.350	0.303	0.498
10	0.619	0.646	0.723	0.634	0.517	0.690	0.478	0.663
11	0.329	0.378	0.604	0.380	0.418	0.475	0.330	0.697
12	0.722	0.704	0.723	0.744	0.671	0.798	0.480	0.668
13	0.842	0.854	0.812	0.911	0.842	0.887	0.796	0.794
14	0.903	0.909	0.932	0.938	0.822	0.957	0.908	0.898
15	0.040	0.013	0.032	0.018	0.040	0.008	0.029	0.029
Mean	0.527	0.535	0.541	0.381	0.442	0.563	0.393	0.588

Table 2. Values of the Accuracy Index (13) for each of the designs across Scenarios 1-15.

Scenario 16

	d_1^B	d_2^B	d_3^B
d_1^A	0.10	0.30	0.45
d_2^A	0.30	0.45	0.60

Scenario 17

	d_1^B	d_2^B	d_3^B
d_1^A	0.10	0.15	0.30
d_2^A	0.30	0.45	0.60

Scenario 18

	d_1^B	d_2^B	d_3^B
d_1^A	0.10	0.20	0.30
d_2^A	0.15	0.30	0.50

Scenario 19

	d_1^B	d_2^B	d_3^B	d_4^B
d_1^A	0.10	0.15	0.20	0.30
d_2^A	0.30	0.40	0.50	0.60

Scenario 20

	d_1^B	d_2^B	d_3^B	d_4^B
d_1^A	0.10	0.20	0.30	0.40
d_2^A	0.20	0.30	0.40	0.50

Scenario 21

	d_1^B	d_2^B	d_3^B	d_4^B
d_1^A	0.15	0.20	0.25	0.30
d_2^A	0.40	0.45	0.50	0.55

Table 3. Alternative dosing grid toxicity scenarios to evaluate the combination designs. Rows and columns refer to the dose of drug A and B respectively. True MTCs are in bold and 'acceptable' combinations are underlined.

Scenario	BOIN	KEYBOARD	S-F	PIPE	Waterfall	BLRM	POCRM	Riviere
16	0.545	0.577	0.745	0.289	0.447	0.657	0.431	0.689
17	0.535	0.549	0.568	0.241	0.567	0.533	0.206	0.375
18	0.418	0.444	0.495	0.079	0.459	0.612	0.192	0.520
19	0.524	0.497	0.451	0.371	0.505	0.428	0.111	0.345
20	0.366	0.392	0.520	0.229	0.326	0.478	0.307	0.475
21	0.404	0.330	0.397	0.294	0.372	0.386	0.154	0.349
Mean	0.465	0.465	0.529	0.250	0.446	0.516	0.233	0.459

Table 4. Values of the Accuracy Index (13) for each of the designs across the alternative dosing grid scenarios.

		Raw Trial Data				Temsirolimus			
				15mg			25mg	50mg	75mg
Neratinib		120mg	0/2	0/4	1/5	0/4			
		160mg	1/4	1/4	0/5	3/6			
		200mg	0/4	1/8	1/2				
		240mg	2/4						

		BOIN				KEY				
		Temsirolimus				Temsirolimus				
		25mg	50mg	75mg			25mg	50mg	75mg	
Neratinib		120mg	0/3	0/0	0/0	Neratinib	120mg	0/3	0/0	0/0
		160mg	1/6	0/6	3/6		160mg	1/6	0/6	3/6
		200mg	2/12	2/3	0/0		200mg	2/9	3/6	0/0

		SFD				PIPE				
		Temsirolimus				Temsirolimus				
		25mg	50mg	75mg			25mg	50mg	75mg	
Neratinib		120mg	0/3	0/0	0/6	Neratinib	120mg	0/3	0/0	0/0
		160mg	1/6	0/6	4/9		160mg	1/3	1/12	0/0
		200mg	0/3	2/3	0/0		200mg	2/15	2/3	0/0

		Waterfall			
		Temsirolimus			
		25mg	50mg	75mg	
Neratinib		120mg	0/3	0/3	0/6
		160mg	1/6	0/3	3/6
		200mg	1/6	2/3	0/0

		BLRM (c)				BLRM (a)				
		Temsirolimus				Temsirolimus				
		25mg	50mg	75mg			25mg	50mg	75mg	
Neratinib		120mg	0/3	0/3	0/6	Neratinib	120mg	0/3	0/3	0/6
		160mg	0/0	0/9	4/9		160mg	0/0	0/9	4/12
		200mg	0/0	0/0	5/6		200mg	0/0	0/0	2/3

		POCRM				Riviere				
		Temsirolimus				Temsirolimus				
		25mg	50mg	75mg			25mg	50mg	75mg	
Neratinib		120mg	0/3	1/6	4/12	Neratinib	120mg	0/3	0/0	2/6
		160mg	0/3	1/6	0/0		160mg	0/3	1/9	2/3
		200mg	0/0	2/6	0/0		200mg	0/3	5/9	0/0

Table 5. Results for each of the designs applied to the case study, including the raw trial data of the study by Gandhi et al.⁶. Each entry represents y_{ij}/n_{ij} . The MTC as chosen by each design is highlighted in bold. In the case of the BLRM, (c) indicates the calibrated prior hyper-parameters were used and (a) indicates the alternative values were used.