Received ; | Revised ; | Accepted

DOI: xxx/xxxx

RESEARCH ARTICLE

Combined criteria for dose optimisation in early phase clinical trials

M. Iftakhar Alam[1] | D. Stephen Coad[2] | Barbara Bogacka[2]

[1]Institute of Statistical Research and Training, University of Dhaka, Dhaka 1000, Bangladesh

[2]School of Mathematical Sciences, Queen Mary University of London, London E1 4NS, U.K.

Correspondence
M. Iftakhar Alam, Institute of Statistical Research and Training, University of Dhaka, Dhaka 1000, Bangladesh.
Email: iftakhar@isrt.ac.bd

The paper aims to investigate whether any bridge is possible between so-called best intention and *D*-optimum designs. It introduces combined criteria for dose optimisation in seamless phase I/II adaptive clinical trials. Each of the optimality criteria considers efficacy and toxicity as endpoints and is based on the probability of a successful outcome and on the determinant of the Fisher information matrix for estimation of the dose-response parameters. In addition, one of the criteria incorporates penalties for choosing a toxic or inefficacious dose. Starting with the lowest dose, the adaptive design selects the dose for each subsequent cohort that maximises the respective defined criterion. The methodology is illustrated with a dose-response model that assumes trinomial responses. Simulation studies show that the method is capable of identifying the optimal dose accurately without exposing many patients to toxic doses.

KEYWORDS:
adaptive design, continuation ratio model, *D*-optimum design, penalty function, phase I/II trial

1 | INTRODUCTION

A phase I clinical trial is the first step in applying a new drug to humans. The designs for phase I trials, particularly those in cancer, focus on the maximum tolerated dose (MTD).[1] This MTD is based on toxicity and it ignores efficacy. Phase II designs aim to determine the efficacy level of an experimental drug, assuming that a dose range has been established in phase I. For cytotoxic agents that are used in treating cancer, it is believed that the higher doses are more likely to be toxic, and also more likely to kill the cancer cells and therefore be efficacious. As a consequence, such designs for cancer clinical trials assume that the dose-toxicity and dose-efficacy relationships are both monotone non-decreasing functions of dose. Therefore, the MTD determined from the phase I trial will provide a dose with a desirable level of efficacy.

In recent years, targeted therapies have drawn attention for treating cancer patients. These are drugs or other substances that block the growth and spread of cancer by interfering with specific molecules that are involved in the growth, progression and spread of cancer. Wages and Tait[2] discussed the emergence of targeted agents in oncology drug development. For these targeted agents, the dose-efficacy curve may not increase with dose. Therefore, conducting the phase II trials based on the MTD from phase I for those agents may not be an efficient way of running the trial. In such situations, it will be useful to consider toxicity and efficacy simultaneously to determine an optimum dose for further investigation in the next phase.

Also, for drugs other than those for cancer, a dose will be acceptable only if it is efficacious and safe. Since it is quite difficult to find a dose which is both efficacious and non-toxic, efforts are necessary to achieve a trade-off between the two. The specific goal is to find a dose which has a high probability of efficacy and a low probability of toxicity. Regardless of the method, phase I trials are small in size and consequently the dose-toxicity curves may not be well estimated. They often determine a dose which

can be found to be either unacceptably toxic or ineffective in a later phase. Seamless phase I/II trials are larger in comparison, and so the associated methods could potentially lead to more efficient dose selection.

Different methods are being developed to increase the popularity of seamless phase I/II clinical trials. There are designs proposed by Thall and Russell,[3] Thall and Cook,[4] and Zhang et al.[5] These designs have the intention of allocating the best dose to the cohort of patients based on the current knowledge and are known as the "best intention designs". They may lead to poor learning of the dose-response relationship. In contrast, there are methods which rely on optimal design criteria for estimation of model parameters. Heise and Myers[6] constructed the D-optimal design using the Gumbel model for bivariate binary data. Fan and Chaloner[7] described the D-optimal design for trinomial responses using a continuation ratio model. Dragalin and Fedorov[8] considered binary outcomes for each endpoint and used either Gumbel bivariate binary logistic regression or the Cox bivariate binary model. They proposed a penalised D-optimal design to improve dose allocation to the cohorts. The design maximises the information under the control of a penalty function for treating patients at doses which are too low or too high. Dragalin et al[9] proposed a design for the combination of two drugs when both efficacy and toxicity are endpoints. This can be viewed as an extension of their previous work, as it utilises the multivariate probit model and generalises the results to combinations of two drugs. Pronzato[10] developed a penalised D-optimum design that is flexible in setting a compromise between the information gained and the cost of the experiment. All of these designs focus on the efficient estimation of the dose-response parameters, so that the optimum dose can be obtained more accurately at the end of the trial. However, they may often expose patients to subtherapeutic or toxic doses on their way to efficient parameter estimation. Alam et al[11] developed a design that, in addition to efficacy and toxicity endpoints, also considered pharmacokinetic information in dose escalation.

Pocock[12] pointed out that a trial should involve a balance between individual and collective ethics and achieving such a balance is not that simple. Individual ethics are concerned with the quality of treatment that each successive cohort in a trial receives. On the other hand, collective ethics are concerned about the future benefits of the patients from the collected information in the trials. This paper investigates whether any trade-off between the best intention designs taking care of individual ethics and D-optimal designs focusing more on collective ethics is possible. The underlying idea is to develop a design that exposes few cohorts in a trial to either subtherapeutic or toxic doses and that can also find the optimum dose accurately. In other words, the paper aims to achieve a balance between individual and collective ethics.

The paper is organised as follows. Section 2 introduces the proposed dose-optimisation criteria. The stopping rules are described in Section 3. Some measures to evaluate the designs are given in Section 4. To illustrate the methodology, an example is introduced in Section 5. Finally, a discussion appears in Section 6.

2 | METHODS

In this section, we present optimisation criteria to find the best dose in a seamless phase I/II adaptive clinical trial. The method is general and can be applied to any dose-response model that allows for efficacy and toxicity endpoints.

2.1 | General idea

We assume that a set of candidate doses $\mathcal{X} = \{x^{(1)}, \ldots, x^{(d)}\}$, developed from preclinical studies or from very early phase Ia clinical studies, is known. We develop a model-based adaptive procedure to investigate the doses for recommendation in later phases of clinical trials. At each stage of a trial, the optimum dose is selected based on one of the newly-defined optimisation criteria. We use a linear combination of the probability of success and D-optimality, both with and without a penalty function applied, to achieve the best of the two approaches. The procedure initially assigns the lowest dose to a cohort of patients. The successive cohorts then receive doses based on the chosen criterion until the trial stops according to the stopping rules. The following sections give details of the proposed method.

2.2 | Maximisation of probability of success

Since our design is for seamless phase I/II trials, there are two endpoints, efficacy and toxicity. Consideration of binary outcomes for each of the endpoints will lead to four possible responses. The corresponding probabilities of the outcomes can be represented as ψ_{00}, ψ_{01}, ψ_{10} and ψ_{11}, where the double subscript indicates presence and absence of efficacy and toxicity in the response. Trinomial responses are also common in such dose-finding studies, where an outcome which is neither efficacious nor toxic is

treated as a neutral response. The outcome which is efficacious but non-toxic is regarded as a success. If an outcome is toxic, irrespective of the efficacy, it is considered as a toxic outcome. The probabilities of these outcomes are represented as ψ_0, ψ_1 and ψ_2, respectively. For the bivariate binary case, ψ_{10} is regarded as the probability of a success. To keep the following discussion general, we denote the probability of a success outcome by ψ_S. In the trinomial case, ψ_1 is taken to be ψ_S. This is a function of dose and it usually depends on some unknown model parameters $\boldsymbol{\vartheta}$.

In what follows, Bayesian parameter estimation is used. Assume that we are at the kth stage of the adaptive procedure, and, based on the current data, we have the posterior estimates $\hat{\boldsymbol{\vartheta}}_k$ of the dose-response model parameters. Then we select the dose $x_{k+1}^{\psi_S}$ for the next cohort of patients that maximises the probability of success, that is,

$$x_{k+1}^{\psi_S} = \arg\max_{x\in\mathcal{X}} \psi_S(x, \hat{\boldsymbol{\vartheta}}_k). \tag{1}$$

The above criterion is defined so that it allocates the doses with the highest probability of efficacy assessed from the current knowledge at stage k. The issue of allocating the most efficacious doses during a trial is very important from an ethical point of view. Note that ψ_S is often called a *utility function* in optimisation theory.

The criterion, by definition, does not take into account the efficiency of parameter estimation. Rather it allows the current best dose to be allocated with the hope that some valuable information can be collected. The designs that allocate the best dose have become very popular in dose-finding studies. Some notable results in this direction are by Zhang et al,[5] Thall and Russell,[3] and Thall and Cook.[4] But these designs have the limitation that they may not converge to the right dose, as indicated in Pronzato[13] and Fedorov et al.[14]

2.3 | Maximisation of determinant of Fisher information matrix

This approach allocates those doses to the cohorts which contribute most to the efficient estimation of the dose-response parameters. Assume that we are at the kth stage in an adaptive clinical trial and that the doses allocated to the cohorts so far are $\boldsymbol{\xi}_k = \{x_1, x_2, \ldots, x_k\}$. Also, based on the data collected from the k cohorts, we have the posterior estimates $\hat{\boldsymbol{\vartheta}}_k$ of the dose-response parameters. Let us define the augmented Fisher information matrix (FIM) as

$$\boldsymbol{M}(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) = \frac{k}{k+1}\boldsymbol{M}(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) + \frac{1}{k+1}\boldsymbol{I}(x, \hat{\boldsymbol{\vartheta}}_k), \tag{2}$$

where $\boldsymbol{M}(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) = \sum_{i=1}^{k} \boldsymbol{I}(x_i, \hat{\boldsymbol{\vartheta}}_k)$ and $\boldsymbol{I}(x, \hat{\boldsymbol{\vartheta}}_k)$ is the estimated Fisher information matrix for a cohort which received dose x. For our dose-response model, the expression for $\boldsymbol{I}(x, \boldsymbol{\vartheta})$ is given in (11). In all of the designs that we consider in this paper, the cohort size is assumed to be the same. Atkinson et al[15] discussed the construction of designs based on augmented FIMs. Bogacka et al[16] used the same approach for the construction of adaptive designs.

We select the dose x_{k+1}^{D} for the next cohort of patients which maximises the criterion Φ_D, that is,

$$x_{k+1}^{D} = \arg\max_{x\in\mathcal{X}} \Phi_D\{\boldsymbol{M}(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}, \tag{3}$$

where $\Phi_{\mathrm{D}}\{\boldsymbol{M}\} = |\boldsymbol{M}|$. This is well known in the statistical literature as the D-optimality criterion for precise estimation of the parameters: see, for example, Atkinson et al[17] and Fedorov and Hackl.[18] The Cramér-Rao inequality tells us that the covariance matrix of the parameter estimators is at least, and asymptotically approaches, the inverse of the FIM. Therefore, by maximising the determinant of the FIM, we minimise the asymptotic lower bound for the generalised variance of the estimated model parameters. Since a dose-response model is usually non-linear, the FIM depends on the model parameters. As indicated earlier, we utilise the current estimates of the parameters to optimise the criterion and consequently the design is locally D-optimal.[17] Because the doses allocated to the cohorts in this design are often found at the extremes of the design region $\mathcal{X}$, such a design has the limitation that patients in a trial may be exposed to non-efficacious or highly toxic doses.

2.4 | Maximisation of determinant of penalised FIM

To reduce the possibility of exposing patients to extreme doses, we can introduce a penalty function in the search for D-optimum doses and consider a penalised D-optimum design, as in Dragalin and Fedorov.[8] This design tries to improve the individual ethics of a trial.

We denote the penalty function for an observation taken at the design point x by $\varphi(x, \boldsymbol{\vartheta})$. Following Dragalin and Fedorov,[8] we choose

$$\varphi(x, \boldsymbol{\vartheta}) = \{\psi_S(x, \boldsymbol{\vartheta})\}^{-C_S}\{1 - \psi_T(x, \boldsymbol{\vartheta})\}^{-C_T}, \tag{4}$$

where ψ_S and ψ_T are the probabilities of success and toxicity at a given dose, respectively, and C_S and C_T are control parameters used to construct an appropriate penalty function at the low and high dose levels. The function is defined so that the lower the probability of success and/or the higher the probability of toxicity at the assigned dose, the higher the penalty for the observation taken. Note that, in the bivariate binary model case, $\psi_{01} + \psi_{11}$ is considered as ψ_T, while in the trinomial model it is ψ_2.

For a given dose x and values of the parameter vector $\boldsymbol{\vartheta}$, the larger the values of C_S and C_T are, the higher the value of the penalty function will be, and thus the design will avoid allocating doses with a low probability of efficacy or a high probability of toxicity to the patients. The case $\varphi(x, \boldsymbol{\vartheta}) = \text{constant}$ means the same penalty across the doses and consequently will lead to the original D-optimum design. The larger the values of the control parameters are, the further the design is from the D-optimum design and the less efficient is the parameter estimation.

The penalised D-criterion takes into account the penalty function in (4). In this version of the D-optimum design, we select the dose at the kth stage as

$$x_{k+1}^{PD} = \arg\max_{x \in \mathcal{X}} \Phi_{\text{PD}}\{\boldsymbol{M}(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}, \tag{5}$$

where $\Phi_{\text{PD}}\{\boldsymbol{M}\} = |\boldsymbol{M}/\boldsymbol{\varphi}|$. The matrix $\boldsymbol{M}(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)$ is defined in (2) and

$$\boldsymbol{\varphi}(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) = \frac{k}{k+1}\boldsymbol{\varphi}(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) + \frac{1}{k+1}\varphi(x, \hat{\boldsymbol{\vartheta}}_k),$$

where $\boldsymbol{\varphi}(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) = \sum_{i=1}^{k} \varphi(x_i, \hat{\boldsymbol{\vartheta}}_k)$. According to Dragalin and Fedorov,[8] the penalised adaptive D-optimum design is more efficient than the adaptive D-optimum design in that both the precision of estimation and the information per cost are higher in the former case. The authors also found that the convergence property holds for the penalised D-optimum design.

The measures based on the information for the evaluation of a design are now described. In the theory of optimal design, the normalised FIM is defined as $\boldsymbol{\mathcal{M}}(\boldsymbol{\xi}_N, \boldsymbol{\vartheta}) = \boldsymbol{M}(\boldsymbol{\xi}_N, \boldsymbol{\vartheta})/N$, where $\boldsymbol{M}(\boldsymbol{\xi}_N, \boldsymbol{\vartheta}) = \sum_{i=1}^{N} \boldsymbol{I}(x_i, \boldsymbol{\vartheta})$ and N is the number of cohorts in a trial. Now, $|\boldsymbol{\mathcal{M}}(\boldsymbol{\xi}_N, \boldsymbol{\vartheta})|$ gives the information per observation in the trial. The higher the value of the information per observation, the higher the estimation precision is. Similarly, the information per cost can be obtained as $|\boldsymbol{M}(\boldsymbol{\xi}_N, \boldsymbol{\vartheta})|^{1/p}/\boldsymbol{\varphi}(\boldsymbol{\xi}_N, \boldsymbol{\vartheta})$, where $\boldsymbol{\varphi}(\boldsymbol{\xi}_N, \boldsymbol{\vartheta}) = \sum_{i=1}^{N} \varphi(x_i, \boldsymbol{\vartheta})$ and p is the number of model parameters.[8] A design with a higher information per cost is regarded as more efficient.

2.5 | Combined criteria

Although the penalised D-criterion introduces a penalty function to improve the quality of treatment during dose escalation, we have found it not to be improved as expected. Therefore, further effort has been taken with the combined criteria defined below. Also, clinicians may be interested in achieving several objectives, such as efficient estimation of the model parameters and allocation of the most efficacious doses to the cohorts during a clinical trial. The combined criteria in (6) and (7) balance these two objectives. Each of the criteria is defined so that one can obtain the design which ensures either efficient parameter estimation or efficacious dose allocation or a combination of both. The two criteria are named as the penalised combined criterion and the simple combined criterion, respectively.

The penalised combined criterion is a linear combination of the determinant of the Fisher information matrix for the dose-response model, penalised for inefficacy and toxicity, and the probability of success. On the other hand, the simple combined criterion does not penalise for inefficacy and toxicity. At each stage of the adaptive trial, we select that dose for which the criterion is maximised.

To implement the penalised criterion, we initially determine the doses $x_{k+1}^{\psi_S}$ and x_{k+1}^{PD} in (1) and (5) that maximise the probability of success and the determinant of the penalised FIM, respectively. Since the determinant and the probability of success may have quite different magnitudes, we scale them at the dose x as

$$E_{PD}(x) = \frac{\Phi_{PD}\{\boldsymbol{M}(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}}{\Phi_{PD}\{\boldsymbol{M}(x_{k+1}^{PD}|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}}$$

and

$$E_{\psi_S}(x) = \frac{\psi_S(x, \hat{\boldsymbol{\vartheta}}_k)}{\psi_S(x_{k+1}^{\psi_S}, \hat{\boldsymbol{\vartheta}}_k)}.$$

The penalised combined criterion then selects the dose x_{k+1} for the next cohort of patients so that

$$x_{k+1} = \arg\max_{x\in\mathcal{X}} \left\{ aE_{PD}(x) + (1-a)E_{\psi_S}(x) \right\}, \tag{6}$$

where a is some weight such that $0 \leq a \leq 1$.

If the D-optimum dose x_{k+1}^{D} and $E_D(x)$ are chosen instead, then the simple combined criterion takes the form

$$x_{k+1} = \arg\max_{x\in\mathcal{X}} \left\{ aE_{D}(x) + (1-a)E_{\psi_S}(x) \right\}. \tag{7}$$

This is a special case of the penalised combined criterion if $\varphi(x, \boldsymbol{\vartheta}) = \text{constant}$. This is due to the fact that, when $\varphi(x, \boldsymbol{\vartheta}) = \text{constant}$, the penalised D-criterion reduces to the D-criterion. Pronzato[10] considered a similar compound criterion to balance between gaining information and the cost of the experiment (cf. Remark 9.2).[19] Compound criteria were introduced by Läuter.[20]

Obviously, the results will depend on the choice of a. It is clear that, when $a = 1$, the combined criterion is simply the penalised D-criterion or D-criterion. Similarly, for $a = 0$, we have dose selection based on the probability of success only. For any other choice of a, the design is expected to allocate the most efficacious doses to the cohorts and to give precise estimates of the parameters leading to the recommendation of the best dose for further study. As we show later, the choice of a is an important component of the method.

Although the probability of success appears twice in the penalised combined criterion, through $E_{PD}(x)$ and $E_{\psi_S}(x)$, its rôle is different in each case. In the penalised D-criterion, the probability of success is used to scale the FIM so that doses with low efficacy or high toxicity result in small values for the determinant of the FIM. The E_{PD} part of the criterion guides us to choose the dose that will provide maximum information regarding parameter estimation, taking into account the efficacy levels of all of the available doses. The maximisation of the probability of success on the other hand tends to choose a dose from the efficacy viewpoint, without accounting for parameter estimation. A combination of them, as presented above, is expected to facilitate a balance between the two approaches.

3 | STOPPING RULES AND OPTIMUM DOSE

We check the possibility of stopping early for futility and/or toxicity at each stage of a trial. More specifically, at the kth stage, we check whether

$$\max_{x\in\mathcal{X}} \{\psi_S(x, \hat{\boldsymbol{\vartheta}}_k) - \psi_T(x, \hat{\boldsymbol{\vartheta}}_k)\} < \lambda.$$

If it is found to be true, we also check whether $\psi_T(x, \hat{\boldsymbol{\vartheta}}_k) > \gamma$ at the dose where the above maximum difference occurs. If both the conditions are found to be satisfied, we stop the trial for futility and/or toxicity and recommend no dose as the optimum dose (OD). Otherwise, we continue the trial. The idea behind such stopping is that, if the maximum difference between the estimated success and toxicity rates is smaller than a desired value λ and also the estimated toxicity at that dose is greater than the acceptable value, then there is no hope of obtaining an OD by continuing this trial. If $\psi_S(x, \hat{\boldsymbol{\vartheta}}_k)$ and $\psi_T(x, \hat{\boldsymbol{\vartheta}}_k)$ takes the same value, then their difference is 0. If $\psi_S(x, \hat{\boldsymbol{\vartheta}}_k)$ takes the value 1, then $\psi_T(x, \hat{\boldsymbol{\vartheta}}_k)$ will have value 0, and, in that case, the difference will be 1. As a consequence, we assume that $0 \leq \lambda \leq 1$. A small value of λ means a flexible checking condition, in the sense that it will allow less frequent checking of $\psi_T(x, \hat{\boldsymbol{\vartheta}}_k)$ than that at a large value of λ. We can set λ as the difference between δ and γ. Note that, starting at the first stage, we keep checking for futility and/or toxicity until the penultimate stage of a trial.

Apart from stopping for futility and/or toxicity, a trial is stopped when the same dose is repeated for r cohorts or when the trial reaches the maximum number of available cohorts m, whichever comes first. Terminating early when the same dose is being repeated saves resources while no more information is obtained.

The optimum dose is defined as the dose that has been repeated r times for the early-stopped trials. For the trials that utilise the maximum number of cohorts m, the optimum dose is defined as the one for which the estimated probability of success is maximum, subject to the constraints that the estimated probability of toxicity at that dose is no more than an acceptable level for toxicity γ and that the estimated probability of success is at least δ.

4 | EVALUATION OF THE DESIGNS

To assess the quality of the proposed adaptive design, we introduce two measures of performance: the *decision efficiency* and the *sampling efficiency*. Both are based on the simulation results, where the responses are obtained for some assumed true values of the model parameters denoted by $\boldsymbol{\vartheta}_{\text{true}}$. Along with these, we also compute some other measures as presented below.

DECISION EFFICIENCY

The *decision efficiency* (DE) is based on the distribution $\tilde{\xi}$ of the OD obtained in a simulation study and written as

$$\tilde{\xi} = \left\{ \begin{array}{c} x^{(1)} \ , \ldots , \ x^{(d)} \\ \tilde{w}^{(1)} \ , \ldots , \ \tilde{w}^{(d)} \end{array} \right\},$$

where $\tilde{w}^{(j)}$ denotes the proportion of times that dose $x^{(j)}$ was recommended for the next phase. In other words,

$$\tilde{w}^{(j)} = \frac{q_j}{U}, \tag{8}$$

where q_j is the number of times that dose $x^{(j)}$ was recommended as optimum in U simulated trials.

To assess how good the recommended dose is, we introduce the ratio ρ of the probability of success at dose x to the probability of success at the true optimum dose x_{OD}, both calculated at the value $\boldsymbol{\vartheta}_{\text{true}}$, that is,

$$\rho(x, \boldsymbol{\vartheta}_{\text{true}}) = \frac{\psi_S(x, \boldsymbol{\vartheta}_{\text{true}})}{\psi_S(x_{\text{OD}}, \boldsymbol{\vartheta}_{\text{true}})}. \tag{9}$$

Dose x_{OD} maximises ψ_S for $\boldsymbol{\vartheta}_{\text{true}}$, and hence $\rho(x, \boldsymbol{\vartheta}_{\text{true}}) \leq 1$. To measure the global efficiency of the decision regarding the dose recommended for further study, we calculate the weighted sum

$$\text{DE} = \sum_{j=1}^{d} \tilde{w}^{(j)} \, \rho(x^{(j)}, \boldsymbol{\vartheta}_{\text{true}}) \, I\{x^{(j)} \in \mathcal{A}\},$$

where the $\tilde{w}^{(j)}$ are defined in (8), $I(B)$ is the indicator function of the event B and $\mathcal{A}$ is the set of doses for which the probability of toxicity is below or equal to the acceptable level. If $x_{\text{OD}} \in \mathcal{A}$ and it is recommended in each of the simulated trials, then the weight $\tilde{w}^{(j)}$ corresponding to x_{OD} is 1 and 0 for all other dose levels, and hence DE = 1. Similarly, if all of the recommended doses are from the toxic region, then the indicator function takes the value 0 for each of them, and hence DE = 0. In general, $0 \leq \text{DE} \leq 1$. If x_{OD} does not exist for a dose-response scenario, then we cannot compute $\rho(x, \boldsymbol{\vartheta}_{\text{true}})$ in (9). However, we can treat ρ as 1 if no dose is recommended in a trial under such a scenario. If a dose from $\mathcal{X}$ is recommended as the OD, then we can treat ρ as 0. These considerations facilitate the calculation of the DE for that scenario. Similarly, a trial may recommend no dose as the OD when there exists a true OD in the scenario. The proportion of such recommendations obtained in the simulations can be ignored in the calculation of the DE.

SAMPLING EFFICIENCY

We define the *sampling efficiency* (SE) based on the information on the allocated doses to the cohorts over the simulations. The distribution ξ of the allocated doses to the cohorts over all simulations is denoted as

$$\xi = \left\{ \begin{array}{c} x^{(1)} \ , \ldots , \ x^{(d)} \\ w^{(1)} \ , \ldots , \ w^{(d)} \end{array} \right\},$$

where $w^{(j)}$ is the proportion of times that dose $x^{(j)}$ was allocated to the cohorts over all simulated trials. In other words,

$$w^{(j)} = \frac{1}{n_{\text{total}}} \sum_{i=1}^{U} n_{ji}, \tag{10}$$

where $n_{\text{total}} = \sum_{i=1}^{U} n_{(i)}$ and $n_{(i)}$ denotes the number of cohorts used in the ith simulated trial. Thus, n_{total} is the total number of cohorts used in all simulations and n_{ji} is the number of cohorts given dose $x^{(j)}$ in the ith simulation.

With the $w^{(j)}$ as in (10), we define the SE as

$$\text{SE} = \sum_{j=1}^{d} w^{(j)}\, \rho(x^{(j)}, \boldsymbol{\vartheta}_{\text{true}})\, I\{x^{(j)} \in \mathcal{A}\}.$$

This criterion assesses the quality of treatment received by the patients during the trials. If the cohorts in the trials only receive x_{OD}, then SE = 1. In that case, there will be no experimental variation in the dose level. In general, we can expect the SE to be appreciably smaller than the DE, unless the variation in ψ_S over $\mathcal{X}$ is small. Here also, $0 \leq \text{SE} \leq 1$.

Hardwick et al[21] consider similar efficiency measures. However, they do not penalise for doses which have a probability of toxicity above the acceptable level. There may be some dose-response scenarios where the maximum value of $\psi_S(x, \boldsymbol{\vartheta}_{\text{true}})$ is obtained at a dose for which the probability of toxicity is above the acceptable level. As a result, $\rho(x, \boldsymbol{\vartheta}_{\text{true}})$ may be greater than 1. Consequently, a design choosing these doses more frequently will have larger values for their efficiency measures than one that limits such toxic doses. This should not be the case, since these are the doses which a good design should not choose because of their toxicity. Therefore, we have presented measures that penalise for toxic doses. The proposed measures are capable of judging a design selecting low efficacious or high toxic doses as a poor one.

RISKS

A drug development process includes the following: targeted population, sampled patients, a specific patient, a sponsor and various regulatory agencies. Since in our design we allocate doses to cohorts of patients, by patients here we would mean cohorts. For the targeted population, the potential loss can be determined by the uncertainty of the recommended optimum dose x^*. Like Lai and Robbins,[22] we consider the quadratic penalty $\phi(x) = (x - x^*)^2$ to assess the risks for various groups. For the targeted population, the potential loss that is due to the uncertainty in the recommended optimal dose x^* is $E\{(x_N^* - x^*)^2\}$, where x_N^* is the optimal dose recommended in a trial using N cohorts. Note that the value of N would be between 1 and m, as we stop a trial early either for futility and/or toxicity or for repetition of a dose r times. Also, the risk for the targeted population is the same as the mean square error of the estimate of the optimal dose. For the sampled patients, we can calculate the risk as $E\left\{\sum_{i=1}^{N}(x_i - x^*)^2\right\}$, where x_i is the dose allocated to the ith patient in a trial. For the nth in-trial patient, the risk is $E\{(x_n - x^*)^2\}$.

NONPARAMETRIC BENCHMARK

We now describe the construction of the nonparametric benchmark design proposed by Cheung.[23] The purpose is to compare our design with the benchmark produced by that design. In actual trials, the outcome for a patient is observed at the administered dose only. However, the essence of the nonparametric design is to find responses for a patient at all d dose levels.

Let $Y = Y(g)$ denote the patient's multinomial outcome at dose level g for $g = 1, \ldots, d$ and let $\pi(g)$ be the associated probability distribution. Assume that Y can take $L + 1$ possible values $\omega_0 < \omega_1 \ldots < \omega_L$. If we consider trinomial responses, then $L = 2$, and the possible values are ω_0, ω_1 and ω_2. The method defines $\pi_l(g) = P\{Y(g) \geq \omega_l\}$ for $l = 1, \ldots, L$, so that we obtain the outcome distribution as $\pi(g) = \{\pi_1(g), \ldots, \pi_L(g)\}$. If we consider trinomial responses, then $\omega_l = l$ for $l = 0, 1, 2$. Now, $\pi_1(g) = P\{Y(g) \geq \omega_1\} = P\{Y(g) \geq 1\} = P\{Y(g) = 1\} + P\{Y(g) = 2\} = \psi_1(g) + \psi_2(g)$. Similarly, it can be shown that $\pi_2(g) = \psi_2(g)$ and $\pi_0(g) = 1$. Note that $\pi_1(g) - \pi_2(g) = \psi_1(g)$.

Now we can draw an outcome Y from the distribution π via latent variables. More specifically, for each patient i, we randomly draw a tolerance profile U_{i1} and U_{i2} from $U(0, 1)$. For a given dose level g, we generate $Y_i(g) = \omega_l$ if $U_{i,l+1} > \pi_{l+1}(g)/\pi_l(g)$ and $U_{ij} \leq \pi_j(g)/\pi_{j-1}(g)$ for all $j = 1, \ldots, l$. The generated $Y_i(g)$ is a simulated outcome for patient i at dose level g. A complete outcome profile for patient i is $\{Y_i(1), \ldots, Y_i(d)\}$. Thus, we can obtain the complete outcome profiles for all of the mc patients in a trial, where c is the cohort size. Then, at each g, we obtain the estimate of π as

$$\hat{\pi}_l(g) = \frac{1}{mc} \sum_{i=1}^{mc} I\{Y_i(g) \geq \omega_l\}$$

for $l = 1, 2$. Then we determine the optimum dose level g^* such that

$$b(\hat{\pi}) = g^* = \arg\max_g \hat{\psi}_1(g),$$

subject to the conditions that $\hat{\psi}_2(g^*) \leq \gamma$ and $\hat{\psi}_1(g^*) \geq \delta$. Note that ψ_1 and ψ_2 correspond to ψ_S and ψ_T, respectively. The dose corresponding to the level g^* is the OD. Thus, we obtain the distribution of the OD from the simulations for a dose-response scenario and compare with it the results for our proposed design.

5 | AN EXAMPLE

To explore the proposed methodology, we introduce an example which is based on the continuation ratio dose-response model. As indicated in Section 2.5, a combined criterion is intended to serve two purposes: allocation of the most efficacious doses to the cohorts and improvement in the estimation of the dose-response parameters. The latter should help to identify the OD more accurately. Simulation studies, detailed in Section 5.2, are conducted to investigate the properties of the design.

5.1 | Dose-response model

Assume that, for each patient, the outcomes can be categorised as neutral, successful or toxic. As mentioned in Section 2.2, success is referred to as an outcome which is efficacious but non-toxic. The occurrence of these outcomes depends on dose. We assume an experimental drug for which the probability of a neutral response decreases monotonically with dose and the probability of toxicity increases monotonically with dose. However, the probability of success may be non-monotonic, increasing or decreasing. As indicated earlier, the corresponding probabilities are represented by $\psi_0(x, \boldsymbol{\vartheta}), \psi_2(x, \boldsymbol{\vartheta})$ and $\psi_1(x, \boldsymbol{\vartheta})$, so that $\psi_0(x, \boldsymbol{\vartheta}) + \psi_1(x, \boldsymbol{\vartheta}) + \psi_2(x, \boldsymbol{\vartheta}) = 1$, where $\boldsymbol{\vartheta}$ is the vector of parameters.

The proportional odds model may be chosen in such a situation, but it requires the dose to have the same effect across the cumulative logits. As it is difficult to satisfy this assumption, we use a more flexible continuation ratio model,[24] which is given by

$$\log\left\{\frac{\psi_1(x, \boldsymbol{\vartheta})}{\psi_0(x, \boldsymbol{\vartheta})}\right\} = \vartheta_1 + \vartheta_2 x \text{ and } \log\left\{\frac{\psi_2(x, \boldsymbol{\vartheta})}{1 - \psi_2(x, \boldsymbol{\vartheta})}\right\} = \vartheta_3 + \vartheta_4 x.$$

The parameter ϑ_1 represents the baseline log-relative probability, ϑ_2 reflects the contribution of dose in the log-relative probability of having a success relative to a neutral outcome, ϑ_3 is the baseline log-odds and ϑ_4 is the contribution of dose in the log-odds of a toxic outcome relative to a neutral or successful one. The above equations have the solutions

$$\psi_0(x, \boldsymbol{\vartheta}) = \frac{1}{(1 + e^{\vartheta_1 + \vartheta_2 x})(1 + e^{\vartheta_3 + \vartheta_4 x})},$$

$$\psi_1(x, \boldsymbol{\vartheta}) = \frac{e^{\vartheta_1 + \vartheta_2 x}}{(1 + e^{\vartheta_1 + \vartheta_2 x})(1 + e^{\vartheta_3 + \vartheta_4 x})}$$

and

$$\psi_2(x, \boldsymbol{\vartheta}) = \frac{e^{\vartheta_3 + \vartheta_4 x}}{1 + e^{\vartheta_3 + \vartheta_4 x}}.$$

The parameter space Θ is restricted so that the above non-linear functions of dose exhibit the assumed behaviour of the responses:

$$\Theta = \left\{\boldsymbol{\vartheta} = (\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4)^T : \vartheta_1 \geq \vartheta_3,\ \vartheta_3 < 0 \text{ and } \vartheta_2, \vartheta_4 > 0\right\}.$$

If we are at the kth stage in a trial, then k cohorts have been treated with doses selected from the set of ordered doses $\mathcal{X}$. Let $\boldsymbol{x}$ be the $k \times 1$ vector of doses with components x_l and let $\boldsymbol{R}$ be the $k \times 3$ outcome matrix with $\boldsymbol{R}_l = (R_{l0}, R_{l1}, R_{l2})$ as the lth row, $l = 1, 2, \ldots, k$. Note that $R_{l0} + R_{l1} + R_{l2} = c$, where c is the number of subjects in a cohort treated with dose x_l. The successive components of $\boldsymbol{R}_l$ are the counts of neutral, successful and toxic responses for the lth cohort. Thus, the likelihood function is

$$L_k(\boldsymbol{\vartheta} \mid \boldsymbol{x}, \boldsymbol{R}) \propto \prod_{l=1}^{k} \{\psi_0(x_l, \boldsymbol{\vartheta})\}^{R_{l0}} \{\psi_1(x_l, \boldsymbol{\vartheta})\}^{R_{l1}} \{\psi_2(x_l, \boldsymbol{\vartheta})\}^{R_{l2}}.$$

Since maximum likelihood estimation is unsuitable because of small sample sizes at the early stages of a trial, we employ a Bayesian approach to estimate the parameters $\boldsymbol{\vartheta}$. The posterior estimate of $\boldsymbol{\vartheta}$ at the kth stage is

$$\hat{\boldsymbol{\vartheta}}_k = \frac{\int_\Theta \boldsymbol{\vartheta}\, p(\boldsymbol{\vartheta}) L_k(\boldsymbol{\vartheta} \mid \boldsymbol{x}, \boldsymbol{R}) d\boldsymbol{\vartheta}}{\int_\Theta p(\boldsymbol{\vartheta})\, L_k(\boldsymbol{\vartheta} \mid \boldsymbol{x}, \boldsymbol{R})\, d\boldsymbol{\vartheta}},$$

where $p(\boldsymbol{\vartheta})$ is the joint prior distribution of the parameters. Let us assume that $0 < \vartheta_2 < u_1$, $0 < \vartheta_4 < u_2$, $v_1 < \vartheta_1 < v_2$ and $v_3 < \vartheta_3 < v_4$, and that the joint prior distribution is uniform. Then we obtain

$$p(\boldsymbol{\vartheta}) = \frac{2}{u_1 u_2 (v_2 - v_3)^2}, \quad \boldsymbol{\vartheta} \in \tilde{\Theta},$$

where

$$\tilde{\Theta} = \left\{ \boldsymbol{\vartheta} : \ v_3 < \vartheta_3 \leq \vartheta_1 < v_2, \ 0 < \vartheta_2 < u_1, \ 0 < \vartheta_4 < u_2 \right\}.$$

The associated FIM for the model parameters is of the form

$$\boldsymbol{I}(x, \boldsymbol{\vartheta}) = c \begin{bmatrix} \frac{\psi_1(1-\psi_2)}{\psi_0(1+e^{\vartheta_1+\vartheta_2 x})^2} & \frac{x\psi_1(1-\psi_2)}{\psi_0(1+e^{\vartheta_1+\vartheta_2 x})^2} & 0 & 0 \\ \frac{x\psi_1(1-\psi_2)}{\psi_0(1+e^{\vartheta_1+\vartheta_2 x})^2} & \frac{x^2\psi_1(1-\psi_2)}{\psi_0(1+e^{\vartheta_1+\vartheta_2 x})^2} & 0 & 0 \\ 0 & 0 & \psi_2(1-\psi_2) & x\psi_2(1-\psi_2) \\ 0 & 0 & x\psi_2(1-\psi_2) & x^2\psi_2(1-\psi_2) \end{bmatrix}. \tag{11}$$

This can be easily obtained from (7) and (22) in Atkinson et al.[15] Matrix $\boldsymbol{I}(x, \boldsymbol{\vartheta})$ is block diagonal and has non-zero submatrices of rank 1 each. In this case, we need at least two different doses, say x_1 and x_2, to obtain a combined non-singular information matrix $\boldsymbol{I}(x_1, \boldsymbol{\vartheta}) + \boldsymbol{I}(x_2, \boldsymbol{\vartheta})$. The first cohort receives the lowest dose in our adaptive trial and the optimality criterion is applied to the second cohort onwards. As a result, the singularity will not be an issue in implementing the method, as the second cohort has the opportunity to receive any dose other than the lowest one to ensure non-singularity. Note that $\boldsymbol{I}(x, \boldsymbol{\vartheta})$ is used in (2) to find the optimum dose in (3). At each stage of a trial, the prior distribution $p(\boldsymbol{\vartheta})$ is utilised with the responses so far to obtain the posterior means $\hat{\boldsymbol{\vartheta}}_k$ of the parameters. These means are then used in a dose-optimisation criterion to obtain the dose for the next cohort. That is to say, the prior that we assume is being used in both the parameter estimation stage and in the design stage, though not in the same way as the standard Bayesian approach.

5.2 | Simulation settings

Six dose-response scenarios are considered for the simulation study, as shown in Figure A1 . The parameter values for these scenarios are chosen to obtain various shapes for the dose-response curves. For each scenario, it is assumed that 20 doses are available in the set $\mathcal{X} = \{0.5, 1.0, \ldots, 10.0\}$ and that the acceptable level for the probability of toxicity is $\gamma = 0.2$. Also, the minimum success probability that an OD should have is assumed to be $\delta = 0.5$. To check stopping for futility and/or toxicity, we set $\lambda = \delta - \gamma = 0.3$. We take the dense grid of dose levels to examine the methodology. In practice, there may be fewer available dose levels, but the procedure should not be affected by this. The true OD in Scenario 1 is 0.5, since it has a probability of success greater than 0.5 and a probability of toxicity less than 0.2. Any dose of 1.5 or greater is toxic for this scenario, as the probability of toxicity at these doses is above γ. Doses 5.5 and 6 are the true ODs in Scenario 2, where doses 6.5 and greater are toxic. The true OD in Scenario 3 is 5.5-7.5, as the success curve is flat here. Doses beyond 8 are toxic in this scenario. Dose 10 is the true OD in Scenario 4, which has many lower non-efficacious doses. None of the doses are toxic in this scenario. Scenario 5 has a very flat success curve. Since the minimum success probability for a dose to be an OD is 0.5, this scenario does not have an OD. Since all of the doses have a probability of toxicity above the target toxicity rate of 0.2 in Scenario 6, it does not have a true OD.

FIGURE A1 ABOUT HERE

Each trial assigns the lowest available dose of 0.5 mg/kg body weight to a cohort of patients. Given a dose assigned under a scenario, the dose-response outcomes are simulated from a trinomial distribution using the corresponding true probabilities of the outcomes. After the first cohort, the dose escalation is based on the criterion in (6) or in (7), using the updated posterior means of the dose-response parameters at each stage. The posterior means are obtained through numerical integration using the *R* package *cubature*.[25]The package carries out adaptive multidimensional integration over hypercubes. As mentioned in Section 5.1, we use a joint uniform prior distribution for $\boldsymbol{\vartheta}$. The parameter space $\tilde{\Theta}$ is chosen for each scenario so that the true values of the parameters lie in the middle of the corresponding intervals. For instance, since Scenario 1 has the true parameters $\boldsymbol{\vartheta} = (1.44, 0.26, -1.70, 0.25)^T$, we consider $\tilde{\Theta} = \{\boldsymbol{\vartheta} \in \Theta : 0 < \vartheta_1 < 2.88, \ 0 < \vartheta_2 < 0.52, \ -3.40 < \vartheta_3 < 0, \ 0 < \vartheta_4 < 0.50\}$.

The same approach is followed for the other scenarios. Each trial is stopped following the rules described in Section 3, where it is assumed that $r = 6$ and $m = 20$. The number of patients in each cohort is 3, that is, $c = 3$.

OPTIMUM VALUES FOR C_S AND C_T

To motivate the values for the control parameters C_S and C_T in the penalty function (4), we conduct the following investigation. In Scenario 1 in Figure A1 , the true optimal dose lies at the beginning of the dose range. Since the probability of success is low and the probability of toxicity is high at the higher doses, the penalty resulting from non-zero values of C_S and C_T will restrict the dose selection from the end of the dose region: see Figure A1 .

In Scenario 2, the success rate is low at the beginning and the toxicity rate is high at the higher doses. Therefore, penalising doses from both ends will keep the dose selection in the middle, where the true optimal dose lies. The success curve in Scenario 3 is of a flat type. Even at the early doses, success is quite high, and toxicity at the higher doses is not as high as in the previous two scenarios. As a result, the penalty for taking observations from both ends is smaller than that for the previous scenario. However, it still keeps the dose selection in the middle of the dose region. The true optimal dose in Scenario 4 is at the end of the dose region. The probabilities of success and toxicity remain very low for more than one half of the available doses. Also, this scenario does not have any dose for which the probability of toxicity is very high. Penalising observations for a low success rate and a high toxicity rate gives very high penalty values for most of the early doses and thereby will keep the dose selection at the upper end of the dose region. The success probabilities at various doses are very low in Scenario 5, and, more importantly, the curve is flat. The toxicity probability reaches the acceptable level towards the upper end of the dose region. As a consequence, we also obtain a flat penalty curve. As seen from Figure A1 , the penalty function will force us to select doses from the beginning of the range. Scenario 6 is a scenario where the probability of toxicity is well above the target toxicity rate throughout the dose range. The probability of success mainly increases with dose and we are likely to select doses from the middle of the dose range because of the penalty function. Considering all of these, we plan to use (4) that penalises doses which have either a low probability of success or a high probability of toxicity, or both.

Having decided on the non-zero values for the control parameters, we assume that they are the same, that is, $C_S = C_T = C$. Since a non-zero value of C will keep more patients away from doses with a low chance of efficacy or from highly toxic doses, we need to use its optimum value. A strong penalty, that is, a larger value of C, will result in designs far from non-penalised D-optimum designs. Although the details are not presented here, we have found that values of C larger than 1 considerably compromised the precision of estimation. As the proposed dose-optimisation criterion is intended to improve the parameter estimates, we have decided to use $C = 1$. This value worked well in terms of preventing the frequent choice of undesirable doses.

5.3 | Numerical results

With the control parameter values set to 1, we run the penalised combined criterion for each of the scenarios. One thousand simulated trials are generated in each case for various values of the weight a. Note that, when $a = 0$, the dose selection is based on the probability of success only. On the other hand, the criterion reduces to the penalised D-criterion for $a = 1$: see (6). Table A1 gives the simulation results for the penalised combined criterion for the scenarios considered. The higher the values of %OD and %AD are, the better the design is. Similarly, %TD is expected to be as small as possible. The sampling and decision efficiency measures can be obtained once the distributions of dose allocation and optimum dose selection are available, as presented in Section 4. The information per observation and per cost and various risk measures are also obtained following Section 4. We expect the information per observation and per cost to be as high as possible. Similarly, the risks are expected to be as low as possible. Although many measures are presented, the best a is based on the DE, as this reflects the quality of the decision regarding the future of an experimental drug.

Consider the results for Scenario 1. The penalised D-criterion ($a = 1$) gives better results than maximisation of the probability of success only ($a = 0$) in terms of most of the measures. However, the SE is the highest when $a = 0$. This is expected, as the criterion based on the maximisation of the probability of success is defined so that more patients are treated during the trial with efficacious doses. The penalised D-optimum criterion does not take care of this. The design performs the best overall when $a = 0.8$, indicating the importance of the penalised D-criterion. Here, the design can identify the true optimum dose 0.5 most accurately. Only a few trials recommended toxic doses for further study in the next phase. About 48.5% of the cohorts are treated at the OD throughout the trials. The decision efficiency is very high in this case (0.918), while the sampling efficiency is

reasonable. For 6.4% of the time, no dose is recommended as the OD. Since this scenario has the very first dose as the true OD, many trials are likely to stop for toxicity in the simulations. The average number of cohorts in a trial is very similar, apart from the case when $a = 1.0$. Both the information per observation and the information per cost increase as a increases. The risk for the targeted population is minimum when $a = 0.8$. As a increases, the risks for both the patient sample and the nth patient increase. This is obvious as the design puts more weight on penalised D-optimality with the increase in a. Note that 10,000 simulated trials are generated for the benchmark design in each scenario. It is clear from the measures that our design outperforms the nonparametric benchmark design described in Section 4. Some of the measures are absent for the benchmark design because of its construction. Since here each patient is exposed to all of the dose levels, calculating %AD and SE will not make any sense. For the same reason, we do not compute the risks for the patient sample and the nth patient. The average cohort size is 20, as we use 60 patients in each simulated trial to be consistent with our design. The information per observation and per cost are not calculated, as the benchmark design does not use the D-optimality criterion at all. The results obtained from the simulated trials utilising the simple combined criterion in (7) are presented in Table A2 . Whereas the penalised combined criterion selects the true OD 91.7% of the time when $a = 0.8$, the simple combined criterion selects this dose only 88.4% of the time. On the whole, the combined criterion with $a = 0.8$ enhances the performance for Scenario 1, whether D-optimality is penalised or not.

TABLES A1 AND A2 ABOUT HERE

Now consider the results for Scenario 2 in Table A1 . The best result in terms of the DE is attained when $a = 0.4$. Of the other values, $a = 0.0$ and $a = 0.6$ also provide satisfactory results. In fact, there is little variation in the performance for the values of a between 0 and 0.6. The most noticeable difference is observed for $a = 1$. For this choice of weight, the design identifies the OD much less accurately compared with the other weights. Fewer cohorts are treated with the optimum dose as the weight increases and there is also a sharp drop when $a = 1$. The efficiency measures DE and SE are both considerably smaller in this case as well. The ratio of the probability of correct identification of the true OD at $a = 0.4$ relative to the benchmark is $70.4 \div 71.5 \times 100 = 98\%$. However, the DE at $a = 0.4$ is well above that for the benchmark design. As seen in Table A2 , the simple combined criterion produces very similar results for this scenario. However, the best a is observed at 0.0.

We obtain similar results for the weights ranging between 0 and 0.8 in Scenario 3 in Table A1 . The best a is observed at 0.4. As in the previous scenario, we observe a decreasing trend in the measures %AD for dose allocation and SE, with a sharp drop at $a = 1$. All of the performance values indicate that the penalised D-criterion on its own is performing poorly compared to the other cases. However, it might be worth combining the criteria in this case, with $a \in [0.2, 0.8]$. Also, the proposed design outperforms the benchmark design in this scenario. Similar conclusions can be drawn in the case of the simple combined criterion, as seen in Table A2 .

For the penalised combined criterion in Scenario 4, it can be argued that the design is performing similarly for weights between 0.4 and 0.8. But the penalised D-optimum design on its own is not performing well in this scenario either. The proposed design is more efficient than the benchmark design in identifying the true OD. However, the DE for the benchmark design is the maximum value in this scenario. This happens as the distribution of the estimated OD for the benchmark design is more centred around the true OD than for the other design: see Scenario 4 in Figures 2, 3 and 5 in the Supplementary Material. A good percentage of trials do not recommend any dose for further development in this scenario. As the highest dose is the true OD here, a trial has more chance to stop early for futility. The results obtained for the simple combined criterion are found to be quite competitive with those for the penalised combined criterion.

The design is most efficient when $a = 1$ in Scenario 5. The DE decreases until $a = 1$. The design is equally as efficient as the benchmark design when $a = 1$. Similar results are found when the simple combined criterion is used. In Scenario 6, the performance of the penalised combined criterion is very similar across the values of a. In addition, the DE of our design at these values is well above that of the benchmark design. The average number of cohorts utilised in each trial is very small. Also, the results are very consistent with those produced by the simple combined criterion in Table A2 . Since a very small number of cohorts is used in each trial, the penalised/simple combined criterion has little rôle in the identification of the OD. It is the stopping rule for futility and/or toxicity that plays the significant rôle and leads to the very similar results. For the last two scenarios, the DE has been computed as described in Section 4. Note that the DE in these scenarios is the same as %ND. We cannot compute the SE here as x_{OD} does not exist. For the same reason, we cannot find %AD for these scenarios. However, to facilitate the computation of various risks, we assume that the true OD is 0.

It is worth mentioning that the SE and %AD are large when $a = 0$ and small when $a = 1$ for the first four scenarios considered, whether the penalised D-criterion is used or not. This is because maximisation of the probability of success allocates most of the cohorts to the most efficacious doses compared to the D-criterion. When we penalise the D-criterion for low efficacious or

highly toxic doses, the improvement in the SE is, however, noticeable. The combined criterion performs well with respect to these measures. In Scenario 4, the SE values are close to each other, apart from the one corresponding to $a = 1$. This is a scenario where we do not have any dose for which the probability of toxicity exceeds the acceptable level, but it has many doses for which the probability of success is very low. Therefore, the differences among the measures are not very substantial for this scenario.

Interesting observations can also be made by comparing the results for penalised and non-penalised D-optimality, that is, when $a = 1$. Tables A1 and A2 clearly show the superiority of the penalised criterion. Most of the measures have better values. Also, the penalised D-optimum design ($a = 1$) always requires on average more cohorts than the designs for $a = 0$ and for the best a. The penalised combined criterion with the best a needs very similar numbers of cohorts to the design with $a = 0$. A similar trend is found when the simple combined criterion is utilised.

The information per observation and per cost increase as a increases across the scenarios. However, the trend in Scenario 6 is different, as the penalised combined criterion has little rôle here because of the small number of cohorts. The risk for the targeted population is minimum at the best a. Both the risk for the patient sample and that for the nth patient increase with the increase in a for the scenarios. A similar trend is found with the simple combined criterion. As this criterion does not employ a penalty function, the information per cost is not computed.

Figures 1 and 2 in the Supplementary Material show the distribution of the estimated OD obtained in the simulations at various values of a for the scenarios. For convenience, %ND is presented above 0. The penalised combined criterion ensures that the most appropriate doses are recommended more often than too toxic and non-efficacious ones. Also, %ND is usually smaller at the best a in Scenarios 1-4 compared to the other values.

The distribution of the estimated OD for the nonparametric benchmark design is presented in Figure 3 in the Supplementary Material. This design recommends many toxic doses as the OD in Scenario 1. The ODs recommended for Scenario 2 are centred around the true OD. Many toxic doses are recommended as the OD in Scenario 3 and not many non-efficacious ones are recommended in Scenario 4. Also, %ND is smaller compared to the penalised combined criterion for this scenario. The benchmark produced in Scenario 5 is quite satisfactory, whereas many toxic doses are identified as the OD in Scenario 6.

The distribution of the estimated OD obtained for the simple combined criterion is presented in Figures 4 and 5 in the Supplementary Material. Many trials recommend no dose as the OD in Scenario 1. The distribution is very similar to the penalised combined criterion for Scenario 2. Many toxic doses are recommended as the OD when $a = 1$ in Scenario 3. Like the penalised combined criterion, many trials recommend no dose as the OD when $a = 1$ in Scenario 4. The distributions in Scenarios 5 and 6 are very similar to their penalised combined criterion counterparts.

6 | DISCUSSION

We have looked at three different approaches for dose finding. In the first approach, the intention is to allocate doses to the patients that are most efficacious according to the current knowledge while searching for the best dose, an approach known in the literature as the “best intention". The second approach targets the most effective, from the parameter estimation point of view, gathering of information and consequently should give the best estimate of the optimum dose. The third approach tries to achieve a trade-off between the two. Best intention designs are ethically attractive, as they take care of the patients, but, unlike the one based on the D-optimality criterion, they have limitations in terms of convergence to the optimum dose. Pronzato[13] and Fedorov et al[14] report that “best intention" designs may converge to a sub-optimal dose. Their studies are based on the frequentist approach and use the least squares or the maximum likelihood estimates of the parameters.

We are using Bayesian parameter estimation, and, to our knowledge, the convergence properties are not known in this case. The best intention design may converge to a sub-optimal dose due to the possibility of poor estimates of the parameters. This possibility is embedded in the criterion (1), which chooses the optimum dose for efficacy, rather than for precise estimation of the parameters. It may happen that, when the responses are successes for the same dose given to a number of consecutive cohorts, the algorithm will not have explored many other doses. For a small number of cohorts, we have observed that, on some occasions, the algorithm chooses a sub-optimal dose. In the majority of cases, the algorithm stops at the true optimum dose, as shown in Tables A1 and A2 . However, as seen, combining the two approaches on the whole can improve the performance of the adaptive design.

Maximum likelihood estimates of the parameters are often used in finding D-optimum designs. But, to form a valid estimator, one needs sufficient information. That is to say, a likelihood equation has no solution until all types of outcomes are observed. On many occasions, we have found toxic outcomes not to be observed, even after applying doses to 10 of the cohorts. Therefore,

we use Bayesian parameter estimates. Note that, to be consistently Bayesian, a total information matrix should be used in (2) by adding to $\boldsymbol{M}$ the information matrix for the uniform prior distribution defined in Section 5.1. However, this possibility is not considered here.[26]

The gains in the combined criteria over the penalised D-criterion or D-criterion are evident from the presented results. All of the performance measures are found to be improved. Most importantly, we notice an appreciable improvement in the quality of treatment allocation, reflected through the sampling efficiency measure SE and the measure of OD allocation during the trial, %AD. The quality of optimum dose selection for the next phase, presented through the DE, is also found to be improved. The combined criteria also outperform the criterion based on the maximisation of the probability of success. In general, the combined criteria utilise a reasonable number of cohorts compared to the other two designs.

All of these results guide us to recommend the proposed combined criteria as dose-optimisation tools in early phase clinical trials. In terms of performance, the penalised combined criterion does slightly better than the simple combined criterion. The choice of values for a will solely depend on the objective. In extreme scenarios like 1, 4, 5 and 6, we have seen that high values of a perform surprisingly well. The middle values are also found to perform satisfactorily in the majority of the scenarios. Since, in reality, we may not know the shape of the dose-response relationship in advance of the trial, we suggest using the middle values. Alternatively, we can have some idea of the optimum value of a during the progress of a trial. After a reasonable number of stages, we can locate where the optimum dose may lie based on the current estimates of the probabilities of success and toxicity. If it is found to be either at the lower end or at the upper end of the dose region, we can use higher values of a afterwards. However, if the OD is likely to be in the middle of the dose region, we can use middle values of a.

ACKNOWLEDGEMENTS

The first author would like to thank Queen Mary University of London for offering him a Doctoral Training Account Studentship during which the work was carried out. In addition, the authors would like to thank the referees for their very helpful suggestions to improve the paper.

CONFLICT OF INTEREST

The authors declare no potential conflict of interests.

References

1. Berry S. M., Carlin B. P., Lee J. J., Müller P. *Bayesian Adaptive Methods for Clinical Trials*. CRC Press, New York; 2010.
2. Wages N. A., Tait C. Seamless phase I/II adaptive design for oncology trials of molecularly targeted agents. *Journal of Biopharmaceutical Statistics.* 2015;25(5):903–920.
3. Thall P. F., Russell K. E. A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics.* 1998;54(1):251–264.
4. Thall P. F., Cook J. D. Dose-finding based on efficacy–toxicity trade-offs. *Biometrics.* 2004;60(3):684–693.
5. Zhang W., Sargent D. J., Mandrekar S. An adaptive dose-finding design incorporating both toxicity and efficacy. *Statistics in Medicine.* 2006;25(14):2365-2383.
6. Heise M. A., Myers R. H. Optimal designs for bivariate logistic regression. *Biometrics.* 1996;52(2):613–624.
7. Fan S. K., Chaloner K.. Optimal designs and limiting optimal designs for a trinomial response. *Journal of Statistical Planning and Inference.* 2004;126(1):347-360.
8. Dragalin V., Fedorov V.. Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference.* 2006;136(6):1800-1823.

9. Dragalin V., Fedorov V., Wu Y. Adaptive designs for selecting drug combinations based on efficacy-toxicity response. *Journal of Statistical Planning and Inference.* 2008;138(2):352–373.

10. Pronzato L. Penalized optimal designs for dose-finding. *Journal of Statistical Planning and Inference.* 2010;140(1):283–296.

11. Alam M. I., Bogacka B., Coad D. S. Pharmacokinetically guided optimum adaptive dose selection in early phase clinical trials. *Computational Statistics and Data Analysis.* 2017;111:183–202.

12. Pocock S. J. *Clinical Trials*. Wiley, Chichester; 1983.

13. Pronzato L. Adaptive optimization and D-optimum experimental design. *Annals of Statistics.* 2000;28(6):1743–1761.

14. Fedorov V. V., Flournoy N., Wu Y., Zhang R. Best intention designs in dose-finding studies Preprint2011.

15. Atkinson A. C., Fedorov V. V., Herzberg A. M., Zhang R. Elemental information matrices and optimal experimental design for generalized regression models. *Journal of Statistical Planning and Inference.* 2014;144(1):81–91.

16. Bogacka B., Uciński D., Atkinson A. C.. Comparison of optimum adaptive experimental designs for dose-finding in early phase clinical trials Preprint2014.

17. Atkinson A. C., Donev A. N., Tobias R. D. *Optimum Experimental Designs, with SAS*. Oxford University Press, Oxford; 2007.

18. Fedorov V. V., Hackl P. *Model-Oriented Design of Experiments. Lecture Notes in Statistics*. Springer, New York; 1997.

19. Fedorov Valerii V, Leonov Sergei L. *Optimal Design for Nonlinear Response Models*. CRC Press; 2013.

20. Läuter Elisabeth. Experimental design in a class of models. *Mathematische Operationsforschung und Statistik.* 1974;5(4-5):379–398.

21. Hardwick J., Meyer M. C., Stout Q. F. Directed walk designs for dose-response problems with competing failure modes. *Biometrics.* 2003;59(2):229–236.

22. Lai TL, Robbins Herbert. Adaptive design in regression and control. *Proceedings of the National Academy of Sciences.* 1978;75(2):586–587.

23. Cheung Y. K. Simple benchmark for complex dose finding studies. *Biometrics.* 2014;70(2):389–397.

24. Agresti A. *Categorical Data Analysis*. Wiley, New York; 1990.

25. Johnson S. G., Narasimhan B. *Cubature: Adaptive Multivariate Integration Over Hypercubes.* R package version 1.1-2; 2013.

26. Fedorov Valerii, Wu Yuehui, Zhang Rongmei. Optimal dose-finding designs with correlated continuous and discrete responses. *Statistics in Medicine.* 2012;31(3):217–234.

□

APPENDIX

A TABLES AND FIGURES

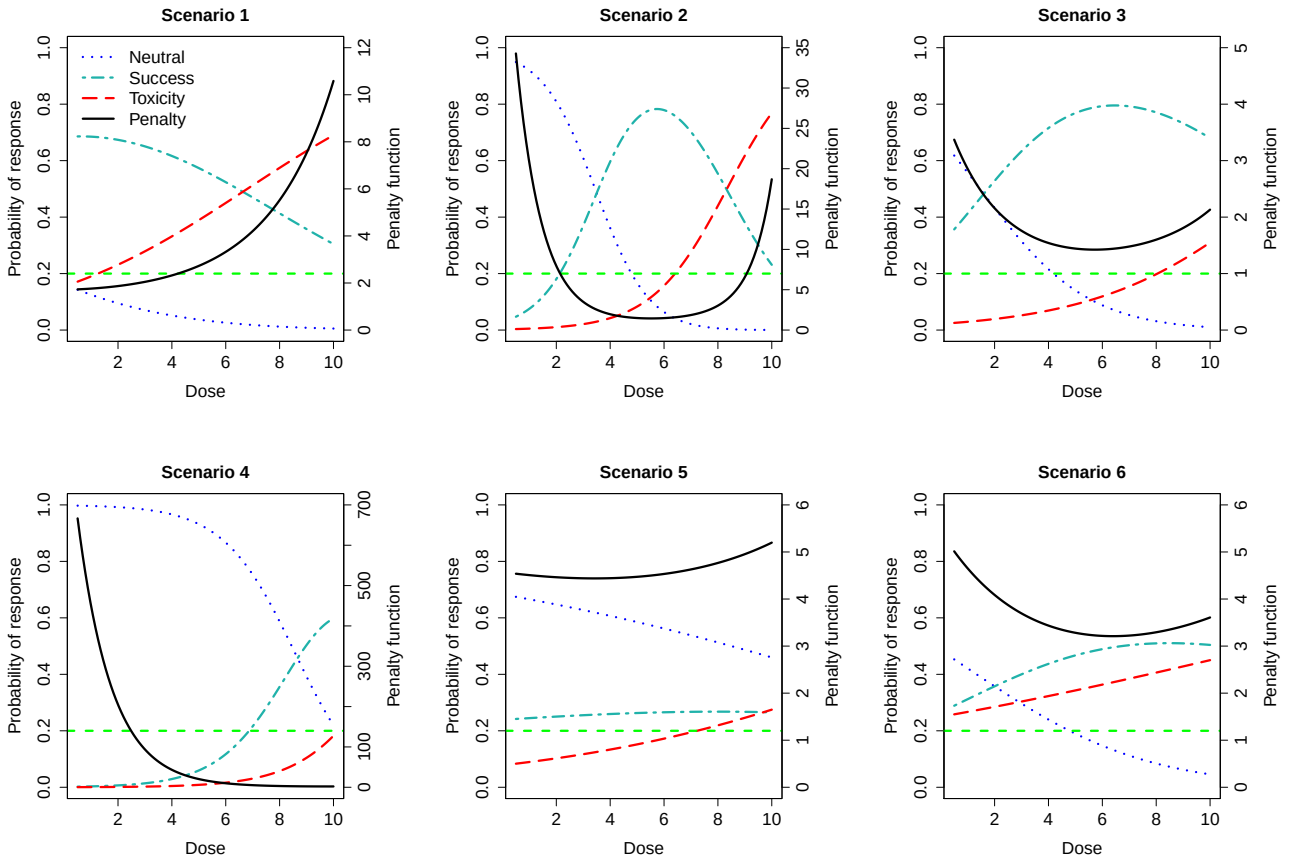

FIGURE A1 Dose-response scenarios for the continuation ratio model. The respective parameter values are: Scenario 1, $\boldsymbol{\vartheta} = (1.44, 0.26, -1.70, 0.25)^T$; Scenario 2, $\boldsymbol{\vartheta} = (-3.50, 1.00, -6.00, 0.72)^T$; Scenario 3, $\boldsymbol{\vartheta} = (-0.80, 0.50, -3.80, 0.30)^T$; Scenario 4, $\boldsymbol{\vartheta} = (-6.50, 0.75, -8.00, 0.65)^T$; Scenario 5, $\boldsymbol{\vartheta} = (-1.05, 0.05, -2.47, 0.15)^T$; and Scenario 6, $\boldsymbol{\vartheta} = (-0.6, 0.3, -1.1, 0.09)^T$. The horizontal dashed line indicates the acceptable level for toxicity and the penalty functions for the scenarios are obtained assuming that $C_S = C_T = 1$.

TABLE A1 Performance of the penalised combined criterion for the six scenarios. Percentage of optimum doses chosen as the true optimum one, recommended for further studies (%OD), percentage of no dose recommended (%ND), percentage of doses recommended as optimum, but carrying the probability of toxicity above the acceptable level (%TD), percentage of cohorts treated at the true optimum doses throughout the trials (%AD), decision efficiency (DE), sampling efficiency (SE) and some other measures. The value of a in bold is regarded as the best in terms of the performance measure DE. Note that $b(\hat{\pi})$ is the nonparametric benchmark.

Scenario	a	%OD	%ND	%TD	%AD	DE	SE	Average cohorts	Info. per obs.	Info. per cost	Risk		
											Targeted population	Patient sample	nth patient
	0.0	63.2	5.6	21.6	42.9	0.728	0.593	10.5	0.003	0.083	0.563	14.7	1.42
	0.2	72.6	6.5	17.8	48.2	0.757	0.547	10.3	0.009	0.125	0.455	17.9	1.73
1	0.4	75.6	7.7	15.2	45.8	0.771	0.481	11.2	0.046	0.212	0.469	39.7	3.54
	0.6	86.8	6.2	6.6	50.2	0.872	0.508	11.0	0.087	0.247	0.255	51.4	4.66
	0.8	91.7	6.4	1.7	48.5	0.918	0.484	11.7	0.136	0.266	0.084	72.0	6.13
	1.0	76.1	14.1	6.3	34.0	0.796	0.339	15.6	0.188	0.272	0.133	151.5	9.47
	$b(\hat{\pi})$	28.0	14.9	37.9	-	0.472	-	20.0	-	-	1.495	-	-
	0.0	68.8	3.8	7.7	39.4	0.872	0.665	16.8	0.011	0.093	1.423	62.9	3.74
	0.2	66.3	4.1	8.0	38.9	0.863	0.656	16.9	0.012	0.094	1.533	63.8	3.78
2	**0.4**	70.4	2.6	7.0	39.2	0.890	0.656	17.3	0.013	0.098	1.050	64.7	3.74
	0.6	68.4	3.4	7.5	37.0	0.876	0.625	17.7	0.019	0.107	1.308	68.4	3.86
	0.8	67.4	4.9	8.8	34.7	0.847	0.603	18.3	0.025	0.114	1.795	74.1	4.05
	1.0	41.0	7.7	5.1	17.2	0.777	0.498	19.7	0.043	0.118	3.508	104.9	5.31
	$b(\hat{\pi})$	71.5	0.0	5.2	-	0.863	-	20.0	-	-	0.288	-	-
	0.0	88.7	0.2	1.5	58.7	0.974	0.869	18.0	0.095	0.332	0.765	103.8	5.78
	0.2	87.1	0.2	2.1	57.7	0.967	0.851	18.1	0.104	0.339	0.816	105.9	5.84
3	**0.4**	89.4	0.1	1.2	57.7	0.978	0.839	18.2	0.116	0.348	0.742	107.5	5.89
	0.6	88.3	0.2	2.2	54.7	0.967	0.802	18.6	0.141	0.364	0.868	113.7	6.11
	0.8	87.2	0.1	2.2	51.0	0.968	0.766	18.9	0.176	0.382	0.874	121.7	6.44
	1.0	60.9	0.0	1.6	26.5	0.921	0.657	20.0	0.334	0.442	3.797	205.8	10.30
	$b(\hat{\pi})$	76.3	0.0	5.2	-	0.935	-	20.0	-	-	1.113	-	-
	0.0	71.2	15.1	0.0	28.6	0.816	0.544	17.9	0.017	0.007	16.025	355.1	19.78
	0.2	71.7	13.6	0.0	28.7	0.821	0.542	17.9	0.017	0.007	14.996	357.0	19.88
4	0.4	74.4	13.6	0.0	29.7	0.829	0.541	17.8	0.017	0.007	14.404	356.5	20.06
	0.6	73.1	14.3	0.0	29.4	0.814	0.538	17.8	0.017	0.007	15.335	358.0	20.10
	0.8	76.6	11.8	0.0	31.0	0.847	0.539	17.6	0.017	0.007	12.788	356.0	20.21
	1.0	47.9	23.7	0.0	26.4	0.647	0.506	19.2	0.018	0.007	26.495	381.7	19.84
	$b(\hat{\pi})$	60.1	10.4	0.0	-	0.873	-	20.0	-	-	10.480	-	-
	0.0	-	84.4	10.2	-	0.844	-	17.3	1.699	0.215	10.785	615.4	35.54
	0.2	-	79.4	16.9	-	0.794	-	16.6	2.069	0.228	17.115	628.5	37.82
5	0.4	-	73.8	22.7	-	0.738	-	16.4	2.499	0.245	22.898	661.8	40.3
	0.6	-	66.0	28.8	-	0.660	-	16.2	2.992	0.262	28.865	660.5	40.72
	0.8	-	57.3	32.7	-	0.573	-	15.6	3.833	0.281	32.746	617.3	39.41
	1.0	-	99.9	0.0	-	0.999	-	17.5	2.858	0.270	0.000	569.2	32.49
	$b(\hat{\pi})$	-	99.9	0.0	-	0.999	-	20.0	-	-	0.015	-	-
	0.0	-	99.1	0.9	-	0.991	-	1.9	0.097	0.037	0.337	21.4	11.23
	0.2	-	99.8	0.2	-	0.998	-	1.7	0.081	0.036	0.200	13.3	7.57
6	0.4	-	99.8	0.2	-	0.998	-	1.8	0.084	0.037	0.048	15.3	8.47
	0.6	-	99.7	0.3	-	0.997	-	1.9	0.115	0.044	0.242	19.3	9.89
	0.8	-	99.5	0.5	-	0.995	-	1.8	0.102	0.037	0.500	17.5	9.70
	1.0	-	99.9	0.1	-	0.999	-	1.9	0.137	0.038	0.016	17.7	9.28
	$b(\hat{\pi})$	-	93.5	6.5	-	0.935	-	20.0	-	-	2.054	-	-

TABLE A2 Performance of the simple combined criterion for the six scenarios. Percentage of optimum doses chosen as the true optimum one, recommended for further studies (%OD), percentage of no dose recommended (%ND), percentage of doses recommended as optimum, but carrying the probability of toxicity above the acceptable level (%TD), percentage of cohorts treated at the true optimum doses throughout the trials (%AD), decision efficiency (DE), sampling efficiency (SE) and some other measures. The value of a in bold is regarded as the best in terms of the performance measure DE. Note that $b(\hat{\pi})$ is the nonparametric benchmark.

Scenario	a	%OD	%ND	%TD	%AD	DE	SE	Average cohorts	Info. per obs.	Info. per cost	Risk		
											Targeted population	Patient sample	nth patient
1	0.0	64.3	6.2	20.0	44.6	0.738	0.606	10.3	0.003	-	0.464	13.4	1.31
	0.2	68.8	5.5	19.6	45.2	0.749	0.589	10.6	0.005	-	0.475	17.1	1.61
	0.4	73.0	7.7	16.1	43.6	0.762	0.469	11.6	0.043	-	0.454	41.6	3.57
	0.6	80.5	9.8	8.8	46.6	0.814	0.473	11.4	0.117	-	0.286	65.7	5.74
	0.8	88.4	7.3	3.7	46.0	0.889	0.460	12.1	0.214	-	0.143	104.4	8.62
	1.0	40.0	32.4	15.0	16.8	0.526	0.168	18.8	0.272	-	0.333	354.6	18.87
	$b(\hat{\pi})$	28.0	14.9	37.9	-	0.472	-	20.0	-	-	1.495	-	-
2	**0.0**	68.5	2.2	7.7	40.0	0.887	0.669	16.8	0.010	-	0.930	62.6	3.73
	0.2	69.6	3.8	7.9	39.9	0.871	0.658	17.0	0.011	-	1.405	63.5	3.73
	0.4	69.2	3.7	7.5	39.6	0.876	0.652	17.2	0.013	-	1.352	64.8	3.77
	0.6	68.2	4.5	7.3	36.6	0.867	0.622	17.7	0.019	-	1.665	69.1	3.90
	0.8	66.5	4.4	6.3	34.7	0.877	0.604	18.3	0.026	-	1.620	74.5	4.07
	1.0	42.7	7.4	6.8	13.1	0.747	0.387	19.8	0.058	-	3.782	141.4	7.14
	$b(\hat{\pi})$	71.5	0.0	5.2	-	0.863	-	20.0	-	-	0.288	-	-
3	0.0	86.3	0.0	2.3	57.9	0.967	0.866	17.8	0.094	-	0.776	104.3	5.8
	0.2	89.5	0.1	1.6	58.3	0.974	0.858	18.1	0.100	-	0.719	104.8	5.8
	0.4	88.9	0.1	1.3	56.8	0.977	0.836	18.2	0.117	-	0.723	108.2	5.9
	0.6	88.2	0.2	2.2	54.7	0.967	0.800	18.6	0.145	-	0.810	113.7	6.1
	0.8	88.2	0.1	2.2	50.6	0.968	0.758	18.9	0.182	-	0.883	123.0	6.5
	1.0	77.8	0.2	5.0	29.2	0.914	0.657	19.9	0.365	-	2.466	226.3	11.3
	$b(\hat{\pi})$	76.3	0.0	5.2	-	0.935	-	20.0	-	-	1.113	-	-
4	0.0	69.0	15.7	0.0	27.9	0.795	0.542	18.0	0.016	-	17.268	356.7	19.8
	0.2	71.6	13.2	0.0	28.4	0.818	0.542	17.9	0.017	-	14.680	355.0	19.8
	0.4	68.6	15.6	0.0	28.4	0.794	0.536	17.9	0.016	-	17.059	361.7	20.2
	0.6	75.5	12.1	0.0	30.2	0.836	0.540	17.7	0.017	-	13.552	356.5	20.2
	0.8	76.1	10.6	0.0	31.0	0.847	0.539	17.6	0.017	-	11.879	354.0	20.1
	1.0	32.2	30.0	0.0	24.8	0.523	0.491	19.7	0.018	-	34.431	404.0	20.5
	$b(\hat{\pi})$	60.1	10.4	0.0	-	0.873	-	20.0	-	-	10.480	-	-
5	0.0	-	81.3	12.4	-	0.813	-	17.3	1.733	-	13.300	618.8	35.7
	0.2	-	78.8	16.5	-	0.788	-	16.8	2.038	-	16.691	642.5	38.2
	0.4	-	76.7	20.7	-	0.767	-	16.7	2.392	-	20.820	659.2	39.5
	0.6	-	62.9	32.3	-	0.629	-	16.1	3.170	-	32.371	671.5	41.7
	0.8	-	53.3	39.3	-	0.533	-	15.7	4.175	-	39.342	660.1	42.0
	1.0	-	100.0	0.0	-	1.000	-	17.5	3.008	-	0.000	591.1	33.8
	$b(\hat{\pi})$	-	99.9	0.0	-	0.999	-	20.0	-	-	0.015	-	-
6	0.0	-	99.5	0.5	-	0.995	-	1.8	0.087	-	0.251	14.8	8.2
	0.2	-	99.8	0.2	-	0.998	-	1.8	0.090	-	0.142	13.9	7.7
	0.4	-	99.7	0.3	-	0.997	-	1.8	0.100	-	0.137	19.1	10.5
	0.6	-	99.4	0.6	-	0.994	-	1.8	0.105	-	0.317	17.3	9.4
	0.8	-	99.7	0.3	-	0.997	-	1.8	0.093	-	0.212	18.9	10.3
	1.0	-	99.7	0.3	-	0.997	-	1.8	0.103	-	0.115	13.0	7.3
	$b(\hat{\pi})$	-	93.5	6.5	-	0.935	-	20.0	-	-	2.054	-	-