



A dynamic stopping rule for phase I clinical trials

M. Iftakhar Alam, Mohaimen Mansur

Institute of Statistical Research and Training, University of Dhaka,
Dhaka-1000, Bangladesh, e-mail: iftakhar@isrt.ac.bd

SUMMARY

This paper investigates a stopping rule to be utilised in phase I clinical trials. The motivation is to develop a dynamic rule so that a trial stops early if the maximum tolerated dose lies towards the beginning of a dose region. Also, it will employ many patients if the maximum tolerated dose lies towards the end of a dose region. A two-parameter logistic model is assumed for the dose-response data. A trial is stopped early before reaching the maximum number of patients when the width of the Bayesian posterior probability interval of the slope parameter meets a desired value. Instead of setting a pre-specified width to stop at, we determine it based on the parameter estimate obtained after a reasonable number of steps in a trial. Simulation studies of six plausible dose-response scenarios show that the proposed stopping rule is capable of limiting the number of patients to be recruited depending on the underlying scenario. Although the rule is applied to a D -optimum design here, it will be equally applicable to other model-based designs.

Key words: phase I clinical trial, maximum tolerated dose, sequential design, D -optimum design, early stopping rule

1. Introduction

Along with identification of the pharmacokinetics of a drug, a phase I clinical trial establishes the maximum tolerated dose (MTD). This is the dose at which an acceptable level of probability of toxicity is attained. The accurate determination of an MTD depends on its location in the dose region. If the MTD lies towards the beginning of a dose region, then a small number of cohorts may be enough to identify it correctly. If the dose lies towards the end of a dose region, a small number of cohorts may be insufficient to find it. This issue may be solved by engaging large numbers of cohorts for all dose-response scenarios. But this will be costly in terms of both

resources and time, as unnecessarily large numbers of cohorts will be used for scenarios with MTDs at the beginning of the region. That is to say, running a trial until it reaches the maximum number of available cohorts is not always a good choice. An alternative could be stopping trials early if some stopping rule is met. Many such stopping rules are available in the dose-finding literature.

Brunier and Whitehead (1994) proposed a Bayesian decision-theoretic approach to determine the number of patients that should be treated in phase II trials. The method determines the optimal sample size by maximising a utility function, which incorporates both the number of successes obtained and the costs of treatment. Korn et al. (1994) proposed a rule for a continual reassessment method (CRM) that stops after c patients are assigned to the same dose, and via simulation they found the value for c to be 6. Stallard (1998) also determined sample sizes for phase II clinical trials under a Bayesian decision-theoretic framework. O'Quigley and Reiner (1998) proposed a stopping rule for the CRM. The rule stops a trial early on the basis that continuing the trial would not lead to a change in the dose recommendation, with high probability. Heyd and Carlin (1999) stopped a trial when either the width of the Bayesian posterior probability interval for the model parameter was less than a pre-specified value or when the trial reached the maximum number of 24 patients, whichever came first. But it is hard and even impossible to recommend a single value for the width which would serve as a stopping criteria in all dose-finding studies. The paper does not provide any insight on this issue. Thall and Russell (1998) developed a design for phase I/II clinical trials in which a trial stopped early when the posterior probability that the toxicity rate and the response rate of all doses were higher or lower, respectively, than the target rates. Zohar and Chevret (2001) proposed Bayesian stopping rules for the CRM, based on either posterior or predictive probability distributions that can be applied sequentially during the trial. These rules aim at early detection of either the mis-choice of dose range or a prefixed gain in the point estimate or accuracy of estimated probability of response associated with the MTD. Zohar et al. (2003) developed software to implement the stopping rules in Zohar and Chevret (2001) for the CRM.

O'Quigley (2002) pointed out two potential problems with the stopping rule of Heyd and Carlin (1999). Firstly, to achieve the most commonly accepted level of precision, many patients are likely to be recruited, and therefore, in practice the trial is not likely to halt before the maximum

number of patients has been included under a fixed-sample scheme. Secondly, obtaining some fixed level of precision for the probability of toxicity at the recommended dose is a major concern per se. O’Quigley (2002) suggested a simple rule, stopping a trial when the same dose is repeated m times. This is similar to the idea in which a trial stops if it is found that the future doses will not be different from the current dose. In both cases, a trial is stopped when convergence to a dose level is achieved. Tighiouart and Rogatko (2012) determined the sample size for escalation with overdose control in a way that the posterior variance of the MTD, on the average over all possible trials, is no more than a specified margin. Early stopping in a clinical trial is implemented in East software, version 6.3 (Cytel). Its sequential feature assists in conducting cost-effective trials. It helps at a midway point to decide whether or not a trial should be continued. Validated stopping rules for efficacy and futility are available in this software.

Alam (2016) developed a D -optimum design for phase I clinical trials and compared it with the CRM. It was learnt from the comparison that the D -optimum design can outperform the CRM in many cases. In this article, we present a modified version of the proposal for early stoppage given by Heyd and Carlin (1999). In particular, instead of setting a pre-specified width for stopping, we set it based on the parameter estimate after the accrual of some data in a trial. The new stopping rule, which we call a dynamic rule, is applied to the D -optimum design developed by Alam (2016). Section 2 gives the motivation of the paper. Section 3 outlines the dose-finding design and the proposed stopping rule. We outline how the method can be implemented in section 4. The simulation settings are considered in section 5, to investigate a variety of circumstances. Section 6 gives an illustration of how the method would work in practice through large-scale simulations. Finally, conclusions appear in section 7.

2. Motivation

The paper is intended to propose a simple stopping rule that will engage either many or few patients to recommend the MTD for further testing in the next phase. Assume that a drug was found promising in preclinical testing and hence deserves testing in humans. Assume that d ordered doses $\mathcal{X} = \{x^{(1)}, \dots, x^{(d)}\}$ of that experimental drug are available for testing in humans, and we would like to determine the MTD to be carried to the next phase. Ideally, a dose-finding design should be able to identify the MTD accurately without exposing many patients. Continuing a trial until

the entire sample is exhausted is often problematic, as the location of the MTD is not the same across scenarios. Sometimes it may lie towards the beginning of a dose region, and therefore, a small number of patients may be enough to identify it correctly. In other cases, it lies towards the end of a dose region, and the correct identification of the MTD requires engaging a large number of patients. Hence determining the optimal number of patients to be recruited in a trial is crucial to avoid wastage of resources. Although a number of papers on this subject have been published, it remains an interesting research problem for which new solutions can be developed.

It is likely that as the number of patients increases, the identification of the MTD improves. Alam (2016) also found improvements in identification with an increased number of patients for both of their designs. We are motivated by those findings to consider a stopping rule that should engage an appropriate number of patients depending on the underlying shape of the dose-response relationship. Here we present a rule that is utilised in finding the MTD when dose allocation is based on a D -optimum design only. However, the method should work equally in the CRM and other model-based adaptive designs. In the following section, we describe the D -optimum design for dose allocation along with the proposed dynamic stopping rule.

3. Methods

3.1. Design

The theory of optimal design can be used in designing clinical trials, as it provides some advantages over non-optimal designs. A D -optimum design depends on the Fisher information matrix (FIM), and it ensures efficient parameter estimation. The FIM for a linear model depends on the design variable, not on the model parameters. However, for a non-linear dose-response model like ours, the FIM also depends on the model parameters along with the design variable. As a consequence, to find the D -optimum dose, some values of the parameters need to be assumed, and such a design is known as locally optimal (Chernoff, 1953). Model parameter estimates obtained at each stage of an adaptive trial can be utilised in the FIM to determine the optimum dose for the next patient.

We utilise the following two-parameter logistic model to depict the dose-response relationship accurately:

$$\psi(x, \boldsymbol{\vartheta}) = \frac{\exp(\vartheta_1 + \vartheta_2 x)}{1 + \exp(\vartheta_1 + \vartheta_2 x)}, \quad (1)$$

where x is the dose given to a patient and $\boldsymbol{\vartheta} = (\vartheta_1, \vartheta_2)$ is the vector of dose-response parameters. Assume that we are conducting an adaptive trial, and at each stage thereof, a dose from \mathcal{X} is given to a patient based on the design criterion. Let $\mathbf{x} = \{x_1, \dots, x_l, \dots, x_k\}^T$ be the vector of doses assigned to k successive patients, and $\mathbf{r} = \{r_1, \dots, r_l, \dots, r_k\}^T$ be the corresponding vector of binary toxic outcomes. Then we can form the likelihood function at stage k as

$$L_k(\boldsymbol{\vartheta} \mid \mathbf{x}, \mathbf{r}) = \prod_{l=1}^k \{\psi(x_l, \boldsymbol{\vartheta})\}^{r_l} \{1 - \psi(x_l, \boldsymbol{\vartheta})\}^{1-r_l}.$$

Until we get both toxic and non-toxic responses from patients receiving different doses, it is not possible to estimate the parameters through maximum likelihood estimation. If we run a rule-based design until a toxic outcome occurs in a trial, only then can the maximum likelihood procedure be applied, and the design can follow a model-based procedure afterwards. However, to make the design fully model-based, here we estimate the dose-response parameters $\boldsymbol{\vartheta}$ utilising the Bayesian approach. At the k th stage, the posterior means of the parameters are obtained as

$$\hat{\vartheta}_{ik} = \frac{\int_{\Theta} \vartheta_i g(\boldsymbol{\vartheta}) L_k(\boldsymbol{\vartheta} \mid \mathbf{x}, \mathbf{r}) d\boldsymbol{\vartheta}}{\int_{\Theta} g(\boldsymbol{\vartheta}) L_k(\boldsymbol{\vartheta} \mid \mathbf{x}, \mathbf{r}) d\boldsymbol{\vartheta}}, \quad i = 1, 2, \quad (2)$$

where Θ is the parameter space and a choice of $\tilde{\Theta} = \{\boldsymbol{\vartheta} : u_1 < \vartheta_1 < u_2, u_3 < \vartheta_2 < u_4\}$ gives the prior distribution as

$$g(\boldsymbol{\vartheta}) = \frac{1}{(u_2 - u_1)(u_4 - u_3)}, \quad \boldsymbol{\vartheta} \in \tilde{\Theta}.$$

To compute the variance of the estimate of ϑ_2 in section 3.2, we also compute

$$\hat{\vartheta}_{2k}^2 = \frac{\int_{\Theta} \vartheta_2^2 g(\boldsymbol{\vartheta}) L_k(\boldsymbol{\vartheta} \mid \mathbf{x}, \mathbf{r}) d\boldsymbol{\vartheta}}{\int_{\Theta} g(\boldsymbol{\vartheta}) L_k(\boldsymbol{\vartheta} \mid \mathbf{x}, \mathbf{r}) d\boldsymbol{\vartheta}}. \quad (3)$$

It can be shown that the Fisher information matrix for a patient is

$$\mathbf{I}(x_l, \boldsymbol{\vartheta}) = \begin{bmatrix} \psi_l(1 - \psi_l) & x_l \psi_l(1 - \psi_l) \\ x_l \psi_l(1 - \psi_l) & x_l^2 \psi_l(1 - \psi_l) \end{bmatrix}.$$

Since the D -optimum design at stage k depends on the accumulated information and the new information coming from a dose, the weighted Fisher information matrix is defined as

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

where $\mathbf{M}(\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k) = \sum_{l=1}^k \mathbf{I}(x_l, \hat{\boldsymbol{\vartheta}}_k)$ is the accumulated information and $\mathbf{I}(x_l, \hat{\boldsymbol{\vartheta}}_k)$ is the Fisher information matrix for a patient who received the dose x_l . A more detailed description of the construction of an optimum design is available in Atkinson et al. (2014). The dose for the next patient is selected so that

$$x_{k+1} = \arg \max_{x \in \mathcal{X}} \Phi_D\{\mathbf{M}(x|\boldsymbol{\xi}_k, \hat{\boldsymbol{\vartheta}}_k)\}, \quad (4)$$

where $\Phi_D\{\mathbf{M}\} = |\mathbf{M}|$. A trial is continued until it meets the stopping rule proposed in section 3.2 or it reaches the maximum number of n patients, whichever comes first. Once the trial stops, the MTD is determined as the dose that would be allocated to patient $n+1$ if he/she were in the trial. This is the dose for which the absolute difference between the estimated probability of toxicity and the target toxicity rate γ is minimum. That is,

$$x_{\text{MTD}} = \arg \min_{x \in \mathcal{X}} |\psi(x, \hat{\boldsymbol{\vartheta}}_k) - \gamma|. \quad (5)$$

3.2. Dynamic stopping rule

The method involves constructing a Bayesian posterior probability interval for the dose-response parameter ϑ_2 . An approximate 95% interval at stage k is obtained as $\hat{\vartheta}_{2k} \pm 1.96\sqrt{\widehat{\text{var}}(\vartheta_{2k})}$, where $\text{var}(\vartheta_{2k}) = E(\vartheta_{2k}^2) - [E(\vartheta_{2k})]^2$. This variance is computed using the estimates in (2) and (3). The width of such an interval is $2 \times 1.96\sqrt{\widehat{\text{var}}(\vartheta_{2k})}$. A trial can be stopped early if the width of the interval is less than or equal to a specified value. If we specify a large value, the trial will stop quickly, but if we specify a small value then the trial will take a relatively long time to stop. Also, it is hard to find a width that will work in all situations. That is, choosing an appropriate width in advance of a trial is a significant challenge. Therefore, we propose setting the value based on the parameter estimate $\hat{\vartheta}_{2k}$ obtained at stage $k = 15$. Stage 15 is chosen to ensure that the dose-response curve has been learned sufficiently. It is clear that a trial does not stop before 15 patients have been engaged. The width for stopping is taken as $w\hat{\vartheta}_{2k}$, where $w = 2/3$. Other

choices of w such as $1/2$ or $1/3$ could be made as well, but that would make the stopping value even smaller, resulting in longer trials. Also, we avoid using just $\hat{\vartheta}_{2k}$, as it would give a relatively large value, causing a trial to stop too early. A trial stops when the width of the approximate Bayesian posterior probability interval is less than or equal to the specified width (as explained) or when the trial reaches a maximum number of n patients, whichever comes first

4. Computational algorithm

The associated computational algorithm to implement the D -optimum design is given below, along with the proposed stopping rule. As already indicated, the stage of the trial is represented by k , its initial value being taken as 1.

Step 1: Assign the current best dose to a patient and observe the response.

Step 2: Estimate the dose-response parameters ϑ .

Step 3: Determine the dose for the next patient. Do not jump more than one dose level at a time during escalation.

Step 4: If $k < 15$, set $k = k + 1$ and repeat Steps 1–4.

Step 5: If $k = 15$ then determine the width $w\hat{\vartheta}_{2k}$ for stopping.

Step 6: Obtain the width of the interval for ϑ_2 . If the obtained width is either smaller than or equal to the stopping width or if the trial reaches $k = n$ then stop, and go to Step 7.

Otherwise, set $k = k + 1$ and repeat Steps 1–6.

Step 7: Determine the MTD as the dose for which the absolute difference between the estimated probability of toxicity and the target toxicity rate is minimum.

5. Simulation settings

To investigate the operating characteristics of the design under the new stopping rule, we consider six plausible dose-response scenarios as shown in Figure 1. Note that the scenarios are the same as in Alam (2016), and this serves to assess the new stopping rule under the same settings. In each scenario, we have six dose levels: $\mathcal{X} = \{1, 3, 5, 7, 9, 11\}$. The steepness of

the dose-toxicity curve decreases as we move from Scenario 1 to 4. The target toxicity probability γ is assumed as 0.33, so that 3, 5, 7, and 11 are the respective true MTDs in the first four scenarios. These doses have true toxic probabilities of 0.32, 0.32, 0.33, and 0.32, respectively. Scenarios 5 and 6 are different from the others in that the target toxicity probability lies midway between two available doses. The true MTD in Scenario 5 is dose 6. Since dose 6 is not available in \mathcal{X} , we assume 5 as the true MTD. It is important to note that the probabilities of toxicity at dose 5 and 7 are 0.24 and 0.43, respectively. Similarly, dose 10 is the true MTD in Scenario 6. Since it is not available in the dose vector \mathcal{X} , we assume 9 as the true MTD. Doses 9 and 11 have probabilities of toxicity 0.28 and 0.39, respectively. The last two cases are considered as they may arise in the actual conduct of trials. To facilitate the Bayesian estimation of dose-response parameters $\boldsymbol{\vartheta}$, we assume a bivariate uniform distribution. The parameter space $\tilde{\Theta} = \{\boldsymbol{\vartheta} : -4.3 < \vartheta_1 < -2.3, 0 < \vartheta_2 < 1\}$ is considered for all the scenarios to see what would happen if a common prior were utilised. It has been found that the above parameter space can lead to many extreme scenarios, including those assumed for the simulation study. Hence, the assumed prior can be regarded as non-informative enough to work with.

Doses are allocated to the patients sequentially in a trial with the lowest dose of 1 mg/kg-body weight applied to the first patient. After receiving a dose, the binary response for that patient is generated from the Bernoulli distribution, and the true probability of toxicity in this case comes from the respective dose-response scenario. After obtaining the dose-response outcomes, the parameters are estimated as described in section 3.1. The dose for the second patient is determined by the dose optimisation criterion in (4). On obtaining the response from the second patient, the model parameters are estimated again, and a dose is chosen for the third patient. A trial continues until it meets the stopping rules in section 3.2. The MTD is then determined so that the absolute difference between the estimated probability of toxicity and the target toxicity rate is minimum, as shown in (5).

The proposed stopping rule is investigated for the variable width $w\hat{\vartheta}_{2k}$ determined at stage $k = 15$ for $w = 2/3$. It is also investigated for the fixed widths 0.50, 0.40, 0.30, and 0.20. The maximum number of patients n is taken as 60. Each of the six scenarios is investigated through 1000 simulated trials using a self-written *R* code.

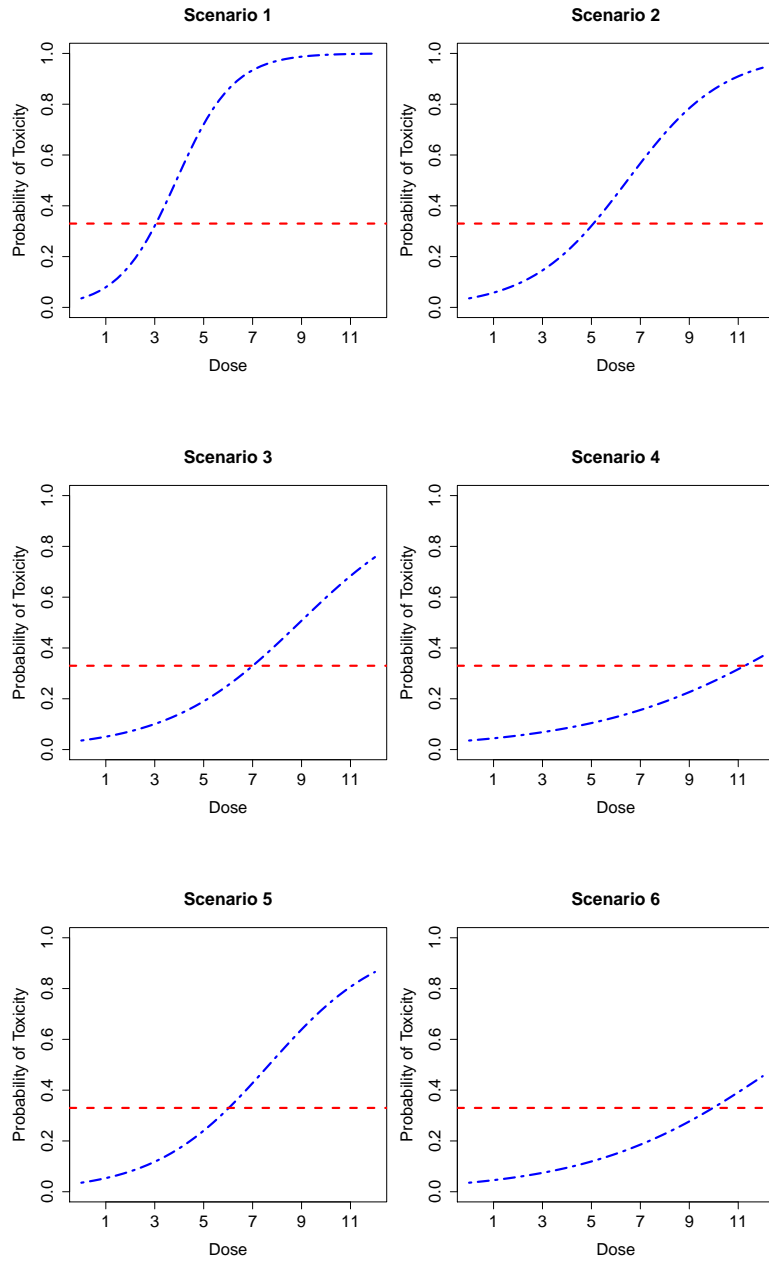


Figure 1. Assumed dose-response scenarios for simulation study; the parameter values in the respective scenarios are $\vartheta = (-3.3, 0.85)$, $\vartheta = (-3.3, 0.51)$, $\vartheta = (-3.3, 0.37)$, $\vartheta = (-3.3, 0.23)$, $\vartheta = (-3.3, 0.43)$ and $\vartheta = (-3.3, 0.26)$; the target toxicity rate γ is indicated by the dotted horizontal line

6. Simulation findings

The results obtained for the fixed stopping width are presented in Table 1. In the table, the bold percentage for each scenario is the rate of identification of the true MTD. A common finding which is consistent across scenarios is that as the value of the fixed width decreases, the accuracy of identification of the MTD increases. For example, Scenario 1, which has dose 3 as the true MTD, achieves 98.3% accuracy of MTD identification when the stopping width is set to 0.50. However, when the value is set to 0.20, the accuracy improves to 99.9%. Similar patterns can be found in other scenarios. Another finding is that when the true MTD lies at the end of the dose region, it is identified with lower accuracy. For instance, in Scenario 4, where the true MTD is dose 11, the accuracy percentages range between 73.2% and 80.8%, compared with the much higher range of 82-95.8% for Scenario 2, which has 5 as the true MTD. In Scenario 5, the true MTD is 6. Since this dose is not available in the dose vector \mathcal{X} , the design mostly recommends the neighbouring doses 5 and 7 as the MTD. Similarly, the true MTD in Scenario 6 lies between doses 9 and 11, and the design recommends these doses as the MTD most often. As in the previous scenarios, identification improves with a decrease in width.

Although the average numbers of patients used at widths 0.50 and 0.20 are similar across the scenarios, this is not the case for the other two widths. To achieve the stopping widths 0.40 and 0.30, higher numbers of patients are found to be engaged when the MTD lies towards the beginning of the dose region than in scenarios where the MTD lies towards the end. This is because the true parameter ϑ_2 is relatively large in the early-MTD scenarios, and therefore, to achieve the desired width, it takes a large number of patients. On the other hand, in the scenarios where ϑ_2 has small values, a relatively small number of patients is required to reach the desired width.

While Table 1 suggests that identification of MTD improves with lower values of the stopping width, it also shows that the accuracy comes at the cost of involving many patients, which makes a trial unnecessarily long. For instance, Scenario 1 requires 59.9 patients, almost the entire sample on average, to achieve the impressive 99.9% accuracy in MTD identification when the width is set to 0.20. Since in this scenario the true MTD lies towards the beginning of the dose region, ideally it should not take that many patients to recommend a MTD for further development in the next phase. Also, unless we have some idea about the dose-response curve, it is quite impossible to suggest a width for stopping. That is to say, in some

scenarios a width of 0.20 may be a good choice, whereas in other scenarios it may be either too narrow or too wide to be achieved. Hence, setting a fixed width in advance of a trial has some serious limitations, and this motivated us to develop the dynamic stopping rule.

Table 1. Average number of patients used, distribution of the MTD, and dose allocation (in parentheses) obtained at various widths for the scenarios; the bold numbers indicate the percentage of correct MTD selection for the respective scenarios

Scenario	Width	Patients	Dose					
			1	3	5	7	9	11
1	0.50	19.6	0.1(22.1)	98.3 (25.8)	1.6(19.9)	0.0(28.1)	0.0(3.6)	0.0(0.4)
	0.40	42.4	0.8(28.5)	99.1 (20.9)	0.0(27.7)	0.1(21.6)	0.0(1.2)	0.0(0.1)
	0.30	57.8	0.3(29.2)	99.5 (20.2)	0.2(28.3)	0.0(20.9)	0.0(1.2)	0.0(0.2)
	0.20	59.9	0.0(24.4)	99.9 (24.6)	0.1(22.6)	0.0(26.6)	0.0(1.6)	0.0(0.2)
2	0.50	17.3	0.0(4.7)	3.7(34.4)	82.0 (4.6)	14.2(31.3)	0.1(14.1)	0.1(10.9)
	0.40	28.8	0.0(7.4)	6.4(30.6)	82.7 (7.3)	9.9(28.7)	0.8(12.5)	0.2(13.4)
	0.30	55.6	0.0(2.7)	2.3(35.6)	92.8 (2.6)	4.9(30.8)	0.0(12.5)	0.0(15.6)
	0.20	59.1	0.0(2.5)	2.0(35.5)	95.8 (2.6)	2.2(30.3)	0.0(11.9)	0.0(17.1)
3	0.50	15.2	0.0(10.4)	0.0(20.1)	30.9(9.8)	43.7 (18.2)	17.8(12.8)	7.6(28.6)
	0.40	16.9	0.0(9.5)	0.0(21.1)	22.9(8.8)	55.0 (19.3)	15.0(12.3)	7.1(28.9)
	0.30	32.7	0.0(5.7)	0.0(24.0)	12.2(5.5)	71.8 (21.6)	14.7(7.4)	1.3(35.8)
	0.20	59.5	0.0(4.2)	0.0(25.5)	9.6(4.6)	83.3 (22.8)	7.0(5.5)	0.1(37.4)
4	0.50	15.0	0.0(22.4)	0.0(4.8)	0.6(19.7)	7.8(4.7)	18.4(18.5)	73.2 (29.8)
	0.40	15.0	0.0(22.4)	0.0(4.8)	0.6(19.7)	7.6(4.7)	17.3(18.5)	74.5 (29.8)
	0.30	29.8	0.0(21.2)	0.0(4.2)	0.0(19.1)	2.9(4.2)	22.6(19.4)	74.5 (31.8)
	0.20	59.8	0.0(21.3)	0.0(4.7)	0.0(18.5)	0.5(4.5)	18.7(18.4)	80.8 (32.7)
5	0.50	15.7	0.0(7.9)	1.2(25.9)	61.9 (7.8)	32.1(23.4)	4.2(12.4)	0.6(22.6)
	0.40	19.6	0.0(6.3)	0.0(26.9)	51.1 (6.2)	43.5(24.6)	4.6(12.3)	0.8(23.6)
	0.30	43.1	0.0(3.1)	0.1(29.9)	53.3 (3.2)	45.2(26.1)	1.3(7.9)	0.1(29.7)
	0.20	59.0	0.1(2.5)	0.2(29.7)	56.7 (3.3)	42.9(25.6)	0.1(6.2)	0.0(32.8)
6	0.50	15.0	0.0(19.8)	0.0(7.6)	2.8(17.4)	16.0(7.3)	25.4 (16.8)	55.8(31.1)
	0.40	15.1	0.0(19.8)	0.0(7.7)	0.7(17.3)	16.8(7.3)	26.3 (16.5)	56.2(31.5)
	0.30	29.3	0.0(18.5)	0.1(7.2)	0.1(16.8)	8.4(7.1)	37.3 (17.2)	54.1(33.3)
	0.20	59.7	0.0(17.3)	0.2(9.1)	0.1(15.1)	4.2(8.6)	43.7 (15.1)	51.8(34.8)

The results for the proposed stopping rule are presented in Table 2. As mentioned earlier, here we use parameter estimates after a few stages to set the dynamic width. This is a more natural way of setting width, since in an unknown situation we will be completely ignorant of its value. Also, to achieve the same level of accuracy in identifying the MTD, the width should be different across the scenarios. It is seen from the table that in Scenario 1 the true MTD is identified 99.5% of times utilising on average 20.7 patients. The same percentage can be achieved, at the cost of a large number of patients, when fixed width is used. In Scenario 2, the

new stopping rule identifies the true MTD in 88.5% of the trials with an average of 40.3 patients used. The rate of correct identification of the MTD in Scenario 3 is 75.6%, at the cost of 46.8 patients per trial. As the MTD lies towards the end of the dose region in Scenario 4, the average number of patients utilised is 59.2, a relatively large number compared with the previous scenarios. Dose 5 is identified as the MTD in 57.1% of the trials in Scenario 5, while dose 7, which is also close to the true MTD, is identified in 42.2% of the trials. These two doses are identified in a total of 99.3% of the trials, with an average of 41.1 patients. In Scenario 6, doses 9 and 11 are identified as the MTD in 94.8% of the trials, the average number of patients used being 58.3.

Table 2. Average number of patients used, distribution of the MTD, and dose allocation (in parentheses) obtained for the scenarios assuming the width as $(2/3)\hat{\vartheta}_{2k}$; the bold numbers indicate the percentage of correct MTD selection for the respective scenarios

Scenario	Patients	Dose					
		1	3	5	7	9	11
1	20.7	0.1(22.1)	99.5 (26.0)	0.4(19.6)	0.0(28.7)	0.0(3.1)	0.0(0.4)
2	40.3	0.1(3.4)	6.7(35.3)	88.5 (3.4)	4.7(31.0)	0.0(11.2)	0.0(15.8)
3	46.8	0.0(5.0)	0.1(25.0)	17.2(4.9)	75.6 (22.7)	7.1(6.2)	0.0(36.2)
4	59.2	0.0(21.6)	0.0(4.4)	0.0(18.7)	1.3(4.3)	17.3(18.6)	81.4 (32.6)
5	41.1	0.1(3.3)	0.4(29.7)	57.1 (3.6)	42.2(26.1)	0.2(7.00)	0.0(30.4)
6	58.3	0.0(18.2)	0.0(8.1)	0.2(15.8)	5.0(7.7)	39.8 (15.9)	55.0(34.4)

7. Conclusion

This study has been undertaken to investigate the proposed dynamic stopping rule. The essence of the rule is that it does not require the width to be specified at the beginning of a trial. Instead, it can be set based on parameter estimates obtained after accumulating information from the first few patients. The motivation comes from the fact that it is hard to suggest a pre-specified width which will work in different dose-response situations. Our simulation results indicate that the proposed stopping rule can reduce the challenge in setting a width. Also, it can decrease the number of patients to be employed in a trial. Therefore, the proposed stopping rule can be regarded as a simple alternative to the fixed-width case.

Instead of using $w = 2/3$, one may utilise a small weight so that the width becomes smaller. However, this will increase the average number of patients to be used, although MTD identification will improve further. Similarly, we can choose the weight in such a way that the width becomes larger.

This will use fewer patients to recommend an MTD. Determining the optimal value for the weight is an open research problem requiring further exploration. The presented approach of setting the width is a more general one. It is expected to work well with other model-based adaptive designs for phase I trials. Other dose-response models may be used as well. Also, a similar approach can be investigated in the settings of seamless phase I/II clinical trials.

REFERENCES

- Alam M.I. (2016): A comparison between the continual reassessment method and D-optimum design for dose finding in phase I clinical trials. *Biometrical Letters* 53(2): 69–82.
- Atkinson A.C., Fedorov V.V., Herzberg A.M., Zhang R. (2014): Elemental information matrices and optimal experimental design for generalized regression models. *Journal of Statistical Planning and Inference* 144(1): 81–91.
- Brunier H.C., Whitehead J. (1994): Sample sizes for phase II clinical trials derived from Bayesian decision theory. *Statistics in Medicine* 13(23-24): 2493–2502.
- Chernoff H. (1953): Locally optimal designs for estimating parameters. *The Annals of Mathematical Statistics* 24(4): 586–602.
- Heyd J.M., Carlin B.P. (1999): Adaptive design improvements in the continual reassessment method for phase I studies. *Statistics in Medicine* 18(11): 1307–1321.
- Korn E.L., Midthune D., Chen T.T, Rubinstein L.V., Christian M.C., Simon R.M. (1994): A comparison of two phase I trial designs. *Statistics in Medicine* 13(18): 1799–1806.
- O’Quigley J. (2002): Continual reassessment designs with early termination. *Biostatistics* 3(1): 87–99.
- O’Quigley J., Reiner E. (1998): A stopping rule for the continual reassessment method. *Biometrika* 85(3): 741–748.
- Stallard N. (1998): Sample size determination for phase II clinical trials based on Bayesian decision theory. *Biometrics* 54(1): 279–294.
- Thall P.F., Russell K.E. (1998): A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* 54(1): 251–264.
- Tighiouart M., Rogatko A. (2012): Number of patients per cohort and sample size considerations using dose escalation with overdose control. *Journal of Probability and Statistics* 2012: Article ID 692725, pp. 16, doi:10.1155/2012/692725.
- Zohar S., Chevret S. (2001): The continual reassessment method: comparison of Bayesian stopping rules for dose-ranging studies. *Statistics in Medicine* 20(19): 2827–2843.

Zohar S., Latouche A., Taconnet M., Chevret S. (2003): Software to compute and conduct sequential bayesian phase I or II dose-ranging clinical trials with stopping rules. *Computer Methods and Programs in Biomedicine* 72(2): 117–125.