



Published in final edited form as:

Stat Methods Med Res. 2022 February ; 31(2): 334–347. doi:10.1177/09622802211064996.

Stopping rules for phase I clinical trials with dose expansion cohorts

Sean M. Devlin¹, Alexia Iasonos¹, John O’Quigley²

¹Memorial Sloan Kettering Cancer Center, New York, USA

²Department of Statistical Science, University College London, U.K.

Abstract

Many clinical trials incorporate stopping rules to terminate early if the clinical question under study can be answered with a high degree of confidence. While common in later-stage trials, these rules are rarely implemented in dose escalation studies, due in part to the relatively smaller sample size of these designs. However, even with a small sample size, this paper shows that easily implementable stopping rules can terminate dose escalation early with minimal loss to the accuracy of MTD estimation. These stopping rules are developed when the goal is to identify one or two dose levels, as the MTD and co-MTD. In oncology, this latter goal is frequently considered when the study includes dose expansion cohorts (DECs), which are used to further estimate and compare the safety and efficacy of one or two dose levels. As study protocols do not typically halt accrual between escalation and expansion, early termination is of clinical importance as it either allows for additional patients to be treated as part of the DEC to obtain more precise estimates of the study endpoints or allows for an overall reduction in the total sample size.

Keywords

Phase I clinical trials; Oncology; Early termination; Expansion cohorts; Sample size; Dynamic stopping rules

1 Introduction

The landscape of early phase clinical trials in oncology has been evolving steadily. Dose expansion cohorts (DEC) have emerged as a nexus between phase I and phase II clinical trials. Distinct from the goal of identifying the maximum tolerated dose (MTD) location in the phase I study and the hypothesis test for efficacy in the phase II study, the scientific goal of the DEC falls under various estimation objectives: a more precise estimate of toxicity, a preliminary estimate of efficacy, or an estimate of the plasma concentration of the drug over time. These estimation objectives are carried out at the MTD, and frequently along with a co-MTD, to support decisions about the recommended phase II dose. In order

Data Accessibility

The R code for the various stopping rules along with the code to generate all simulation scenarios can be obtained from www.github.com/sedevlin.

to achieve these objectives, the expansion cohorts typically represent more homogeneous patient populations to align with what could be used in the next study phase.

Given the resources required to open multiple DEC's, each of which typically includes 20 or more patients, it is imperative to conclude the dose escalation and commence the DEC as soon as there is evidence the escalation design has settled on a dose level or any additional patients are unlikely to change the recommended dose. Early termination of dose escalation can reduce the overall sample size, or expand the number of patients treated in the DEC to obtain more precise estimates of the study endpoints.

In this paper, we consider some existing and some new dynamic stopping rules for early termination of dose escalation within the framework of model based designs, such as the continual reassessment method (O'Quigley et al., 1990; Storer, 1989). The main feature of these designs is that the current estimate of the MTD is sequentially updated using all patients treated up until that point using a single parameter dose-toxicity working model. New patients, either alone or in a cohort, are treated at the dose estimated to be closest to the MTD based on a target toxicity rate (α). A fixed final sample size n is decided upon for the study and the estimated MTD is the dose that would be allocated to patient $n + 1$ were they to be included in the trial. One of our goals is to evaluate whether stopping before patient n has an impact on the accuracy if various stopping criteria are achieved.

Various rules have been proposed for the early termination of CRM trials (O'Quigley et al., 1990; O'Quigley and Reiner, 1998; Heyd and Carlin, 1999). The rule suggested by O'Quigley et al. (1990), later given a theoretical justification by Shen and O'Quigley (1996) and further studied by Heyd and Carlin (1999), decides whether enough precision has been obtained, on the basis of a confidence interval, to bring the study to an early close. There are two limitations to this approach. Firstly, most commonly accepted levels of precision would typically require many more patients than available in phase I studies so that, in practice, the trial is not likely to halt before n patients have been included. Secondly, it is not clear that obtaining some fixed level of precision for the probability of toxicity at the recommended dose is of itself a major objective before the DEC. The maximum tolerated dose will correspond to some point percentile but, in practice, there is some room for flexibility around this point.

Alternatively, O'Quigley and Reiner (1998) proposed early stopping rules that consider the potential toxicity information of all remaining patients that are required to reach the final sample n . With these different potential paths to study completion, the authors estimate the probability that the current recommended dose level will turn out to be the final recommended level and potentially that, in addition, all remaining patients will be treated at the current level. The principle behind the method leans to some degree on the property of CRM to converge to some dose level and then recommend treating all remaining patients at this same level. If able to predict with high probability what the final recommendation will be, then CRM can be terminated early.

This property of convergence suggests a yet simpler rule. One can keep track of the number of times each dose is recommended during the trial and take as a stopping rule the following

condition: stop when the dose recommended for a subsequent patient is the same dose recommended for the previous k consecutive patients, where k is some number fixed at the beginning of the trial. This has the advantage of great simplicity and, as we show below, works out reasonably well in practice.

This paper gives further consideration to these rules and takes a close look at their operational characteristics in practical settings. As the goal is to evaluate the stopping rules in the context of phase I trials with dose expansion cohorts, we extend the stopping criteria to incorporate the selection of a co-MTD in association with the MTD itself. Dealing with two levels instead of a single dose to recommend adds a layer of overall complexity. In this paper, we extend and build upon existing concepts in order to enable the estimation of the various probabilities that can provide guidance for stopping early in this setting.

2 Model structure and estimands

The overall purpose of the trial is to identify the MTD. For this we first need to precisely define what we mean by the MTD. We have available for study m ordered dose levels, d_1, d_2, \dots, d_m . The probability, β_i , of encountering a dose limiting toxicity (DLT) at dose level d_i is unknown although we do know that $\beta_1 < \beta_2 < \dots < \beta_m$, i.e., we have a monotonic increasing relationship between the dose and the probability of a DLT at that dose. These ideas can be generalized, and some assumptions relaxed. For now, we can work within this framework. For the purpose of notational completeness and harmony, we add two doses to our set-up: doses d_0 and d_{m+1} , where $\beta_0 \approx 0.0$ and $\beta_{m+1} \approx 1.0$. These two doses are of no practical concern.

We define the MTD by the parameter $\theta \in (d_1, \dots, d_m)$, as $\theta = \operatorname{argmin}_i |\beta_i - a|, i = 1, \dots, m$, which is defined with respect to some “acceptable” target DLT rate a . We make the theoretical working assumption that, for no i , does $\beta_i = a$. Under our assumptions, the parameter θ is uniquely defined. Very closely related to θ is the parameter-pair, (θ^-, θ^+) . If we look at pairs of doses (d_i, d_{i+1}) then $(\beta_i - a) \times (\beta_{i+1} - a) < 0$, defines uniquely the pair $(d_i, d_{i+1}) = (\theta^-, \theta^+)$. Clearly, $\theta \in \{\theta^-, \theta^+\}$. A common goal would be to identify the pair (θ^-, θ^+) and to take it forward to the dose expansion cohort.

Progress is made by assuming some working model. Suppose for the dose toxicity curve, we have $\beta_i = \psi(d_i, a)$, $i = 1, \dots, m$, to model and reproduce the probabilities of interest. We focus on the much stricter parameterization $a = a_1 \dots = a_m$, and that, for some a , $\beta_i = \psi(d_i, a)$, $i = 1, \dots, m$, which greatly reduces the complexity of the problem. In the place of m parameters on which to focus our estimating power, we now only have one. When the aim is to make inference over the whole range of β_i , this will generally not work well, and the biases can be great. However, when the goal is to treat patients close to the MTD, the one-parameter model is adequate. Indeed, increasing the dimension from one parameter to two parameters, even when the latter model correctly generates the observations, will lead to poorer performance (Iasonos et al., 2016). For the purposes here, we assume a one-parameter continual reassessment model. As well as leading to improved operational performance, the use of a single parameter greatly simplifies the theoretical development of the following sections.

The overall purpose is constrained by other considerations. We would like to estimate as efficiently as possible while respecting ethical constraints such as the minimization of the number of patients treated at doses either too far above, or too far below, the MTD. As the trial progresses, pairs of data (X_j, Y_j) are observed, where the binary variable Y_j is the drug-related toxicity for individual j , equal to 1 if a DLT is observed after the individual is treated at $X_j = x_j$, where $x_j \in \{d_1, d_2, \dots, d_m\}$. We let $\Omega_j = \{(x_\ell, y_\ell), \ell = 1, \dots, j\}$ denote the dose-toxicity information after the first j patients are treated. This information is then included in the likelihood expression which can be written as

$$L(a; \Omega_j) = \prod_{\ell=1}^j \psi^{y_\ell}(x_\ell, a) \{1 - \psi(x_\ell, a)\}^{(1 - y_\ell)}, \tag{1}$$

where $\psi(x, a)$ is the assumed dose-toxicity working model. Since we anticipate our model to be misspecified, the standard likelihood and Bayesian theory cannot be applied automatically, and more is required. Following Shen and O'Quigley (1996), this model is selected with the following properties: the parameter a belong to a finite interval $[A, B]$; for a fixed a , $\psi(x, a)$ is continuous and strictly increasing in x ; for a fixed x , $\psi(x, a)$ is continuous and strictly decreasing in a . These properties are not restrictive, enabling us to avoid problems to do with singularities, as well as guaranteeing the existence and uniqueness of the dose and model parameter given the other.

Under a Bayesian framework with a prior density $g(a)$, the posterior distribution for the parameter a is

$$f(a | \Omega_j) = H^{-1}(\Omega_j) \times g(a)L(a; \Omega_j),$$

where $H(\Omega_j) = \int_{a \in [A, B]} g(u)L(u; \Omega_j)du$. Estimates for a are sequentially updated, either via maximum likelihood or by working with the posterior distribution whereby $\hat{a}_j = \int_{a \in [A, B]} a f(a | \Omega_j) da$. In the same way we can immediately obtain Bayesian estimates,

$$E\{\psi(d_i, a)\} = \int_{a \in [A, B]} \psi(d_i, a) f(a | \Omega_j) da,$$

or make use of plug-in estimates, $\psi(d_i, \hat{a})$, for the probabilities of DLT at dose levels $d_i, i = 1, \dots, m$. In turn, these result in estimates for the MTD θ , and the MTD/co-MTD pair (θ^-, θ^+) , denoted by $\hat{\theta}(\Omega_j)$ and $\{\hat{\theta}^-(\Omega_j), \hat{\theta}^+(\Omega_j)\}$, where we make explicit the dependence of these estimators on the current set of observations, Ω_j . Once we have this information we can then decide on the next level to which the incoming patients may be allocated.

3 Dynamic Stopping Rules

Following observations on the currently treated patient, or cohort of patients, the objective is to determine whether the dose escalation study should continue accruing the full n patients initially planned or whether we have enough evidence to bring the study to an early close. Broadly, this decision can be made if it is likely that additional patients will not change the

recommended MTD, that the method has stopped oscillating between dose levels and began to plateau in the last few patients, or there is high probability that the current recommended dose level after j patients is the true MTD.

3.1 Posterior Probability

In order to make a decision based on the posterior probability, we can calculate an odds and, for sufficiently large or small values of the odds, we can bring the study to a close. Denoting the odds by \mathbb{R} , three possible odds have particular appeal. These are: (1) $\mathbb{R}_1(d_i)$, the odds that d_i , as opposed to a different dose, corresponds to the true MTD, (2) $\mathbb{R}_2(d_{i-1}, d_i)$, the odds that the target rate of DLT lies in the interval (d_{i-1}, d_i) as opposed to any other interval, and (3) $\mathbb{R}_3(d_{i-1}, d_i)$, the odds that one of d_{i-1} or d_i corresponds to the MTD as opposed to neither in the pair being the MTD. Which particular odds to work with would be a design parameter chosen during protocol development. The first of these odds would be suitable when we only want a single estimated MTD to take forward to the DEC. The second and third of these odds can be appealing when the goal is to identify a pair of doses to be taken forward to the DEC. The following subsection indicates how to structure these choices precisely.

Under the assumptions of the working dose-toxicity model discussed in Section 2, we have m constants $a_1, \dots, a_m \in [A, B]$ such that for $1 \leq i \leq m$, $\psi(d_i, a_i) = \beta_i$, $\psi(d_i, B) < \alpha < \psi(d_i, A)$, and for a unique $a_M \in (a_1, \dots, a_m)$, $\psi(d_M, a_M) = \beta_M$, where d_M is the MTD. We anticipate β_M to be closest to α although, in the practical setting, we do not expect them to exactly coincide. From these assumptions we have the following results:

Lemma 1—For each $1 \leq i \leq m$, there exists a unique constant τ_i such that $\psi(d_i, \tau_i) = \alpha$.

Lemma 2—For each $1 \leq i \leq m$, there exists a unique constant κ_i such that

$$\alpha - \psi(d_i, \kappa_i) = \psi(d_{i+1}, \kappa_i) - \alpha > 0.$$

These two lemmas allow us to partition the interval $[A, B]$ into a union of non overlapping intervals so that:

$$[A, B] = \bigcup_{i=1}^m S_i = \bigcup_{i=1}^m T_i,$$

where

$$S_1 = [A, \kappa_1], S_2 = [\kappa_1, \kappa_2], \dots, S_m = (\kappa_{m-1}, B],$$

and

$$T_1 = [A, \tau_1], T_2 = [\tau_1, \tau_2], \dots, T_m = (\tau_{m-1}, B].$$

Figure 1 illustrates these sub-intervals for five dose levels using a simulation set-up discussed later in Section 4.

The importance of these partitions of the parameter space are reflected in the following corollaries and theorem.

Corollary 1—*Suppose that \hat{a}_j is the estimate of the parameter a after the inclusion of j patients. If $\hat{a}_j \in S_i$ then d_i is the level recommended to patient $j + 1$.*

Corollary 2—*Suppose that \hat{a}_j is the estimate of the parameter a after the inclusion of j patients. If $\hat{a}_j \in T_i$ then the estimated MTD, the level recommended to patient $j + 1$, is either d_{i-1} or d_i where $\psi(d_{i-1}, \hat{a}_j) < \alpha < \psi(d_i, \hat{a}_j)$.*

Instead of focusing our development on point estimates, $\hat{a}_j, j = 1, \dots, n$, it can be advantageous to think in a more Bayesian way and to view the intervals, S_j and T_j , as random variables with respect to a probability measure, \mathbb{P} , based on the sets Ω_j . Specifically,

Definition 1—*Given the density $g(a)$ and the set Ω_j , we define a probability measure, \mathbb{P} , mapping the sub-interval, $V \in [A, B]$ to $(0,1)$ by:*

$$\mathbb{P}(V) = H^{-1}(\Omega_j) \int_{u \in [A, B]} \mathcal{I}_{u \in V} L(u; \Omega_j) g(u) du \quad (2)$$

where $H(\Omega_j) = \int_{u \in [A, B]} L(u; \Omega_j) g(u) du$ is a normalizing measure over $[A, B]$ and where $\mathcal{I}_C = 1$ when C holds and is zero otherwise.

Rather than base allocation decisions on \hat{a}_j we might then base these decisions upon $\max_i \mathbb{P}(S_i), i = 1, \dots, m$ (Iasonos and O'Quigley, 2016). Both approaches are common and are unlikely to disagree apart from very early in the trial when the prior weights may have an impact.

We further make use of these intervals when we wish to stop on the basis of two levels, an estimated MTD together with a corresponding co-MTD estimate, and the way to do this can be seen via the following lemma. This lemma extends Corollary 1 to the situation of two dose levels.

Lemma 3—*Suppose that \hat{a}_j is the estimate of the parameter a after the inclusion of j patients. The pair, $\{\hat{\theta}^-(\Omega_j), \hat{\theta}^+(\Omega_j)\} = (d_i, d_{i+1})$, if and only if $\hat{a}_j \in T_{i+1} \cap \{S_i \cup S_{i+1}\}$.*

The lemma is readily demonstrated and allows us (below) to work out the above probabilities. Note that we are not identifying which one of the pair (d_i, d_{i+1}) is the currently estimated MTD and, furthermore, as j increases to n , while the pair can remain the same the estimated MTD itself can change, possibly more than once. This lemma can guide decision making when we treat the intervals, S_j and T_j , as fixed. When we view them as random with

associated probability measures, $\mathbb{P}(S_i)$ and $\mathbb{P}(T_i)$, then we have a direct analogy to the above lemma:

Lemma 4—*The pair $\{\tilde{\theta}^-(\Omega_j), \tilde{\theta}^+(\Omega_j)\} = (d_\ell, d_{\ell+1})$ if and only if $\operatorname{argmax}_i \mathbb{P}[T_{i+1} \cap \{S_i \cup S_{i+1}\}] = \ell$.*

These lemmas are no more than working tools that can help with decision making.

Corollary 3—*Under strict monotonicity, the above lemmas simplify since $T_{i+1} \cap \{S_i \cup S_{i+1}\} = T_{i+1}$.*

The corollary is readily seen since, in our set-up, T_{i+1} is strictly contained within the set $S_i \cup S_{i+1}$. Corollary 3 defines the various intervals on the parameter space which correspond to different odds ratios as defined in Theorem 1 below. As such, we are able to estimate whether the target toxicity rate falls within an interval between two dose levels using the single sub-interval interval T_i .

Theorem 1—*All 3 odds described above can now be defined in terms of \mathbb{P} as;*

$$\begin{aligned} \mathbb{R}_1(d_i) &= \frac{\mathbb{P}(S_i)}{1 - \mathbb{P}(S_i)}, \quad \mathbb{R}_2(d_{i-1}, d_i) = \frac{\mathbb{P}(T_i)}{1 - \mathbb{P}(T_i)}, \quad \mathbb{R}_3(d_{i-1}, d_i) \\ &= \frac{\mathbb{P}(S_{i-1} \cup S_i)}{1 - \mathbb{P}(S_{i-1} \cup S_i)} \end{aligned} \quad (3)$$

We can base decisions on when to stop on any or all of $\mathbb{R}_1(d_i)$, $\mathbb{R}_2(d_{i-1}, d_i)$ and $\mathbb{R}_3(d_{i-1}, d_i)$. As a suggestion, an odds of 3 or more would be enough evidence to bring this part of the trial to a halt and to proceed to the next phase. The decision of whether to bring one or two doses forward may depend on which of the odds we make use of. Such decisions are easily incorporated into the protocol itself.

As an example, estimates of these three odds for one example trial is provided in Figure 2. From these figures, if the goal is to select two dose levels for the DEC, the odds are highest that d_2 or d_3 correspond to the MTD as opposed to neither; however, if using an odds threshold of 3, there is not sufficient evidence to halt the study at this point. Depending on the goals of the study, these intervals can also be combined or modified to estimate other quantities of interest, which can be similarly estimated using Eq. (2). As shown in Figure 3, one such example is that we can estimate the odds each dose level is either the MTD or co-MTD as opposed to neither using the intervals T_i . From this example, the odds are high that d_3 is the MTD or co-MTD.

Basing decision on these odds and probabilities is less computationally burdensome than the idea behind the tree based rules of O'Quigley and Reiner (1998). The tree-based rules, recalled and modified in the following section, can quickly run into issues of combinatorial complexity.

3.2 Tree-based Rules

As noted above, if given enough patients, any CRM trial will go from a period of oscillation between dose levels to a more or less a plateau, where the dose recommended at the end of the plateau is the MTD. O'Quigley and Reiner (1998) considered to what extent we could anticipate at some earlier point in the study that we are indeed on a plateau. In some cases this can be stated with certainty, i.e. whatever toxicity is observed in the remaining patients will not change the recommended level. More generally, we can estimate the probability of staying and remaining at some level for the remainder of the study.

As an initial example, which is discussed again later, consider Figure 4. After 17 patients have been treated, the current recommended dose level is 3. We can consider the exhaustive ordering of all possible DLT events for the remaining patients and estimate the associated probabilities of each potential final recommended dose level using the current estimate \hat{a}_{17} . It is under this general framework the tree-based rules were further developed.

There are two probabilities of potential interest: the probability that the dose recommended after the inclusion of all n patients will be the same as that recommended after the inclusion of the first j patients, and the probability of the more restrictive event in which, not only are the recommendations following j and n inclusions the same, but all inclusions between j and n have the same level. We denote these probabilities by:

$$\mathcal{P}_1(j) = P[\hat{\theta}(\Omega_n) = \hat{\theta}(\Omega_j)]; \quad \mathcal{P}_2(j) = P[\hat{\theta}(\Omega_n) = \hat{\theta}(\Omega_{n-1}) = \dots = \hat{\theta}(\Omega_j)]. \quad (4)$$

In the same way, we are interested in:

$$\begin{aligned} \mathcal{P}_1^-(j) &= P[\hat{\theta}^-(\Omega_n) = \hat{\theta}^-(\Omega_j)]; & \mathcal{P}_2^-(j) &= P \\ &[\hat{\theta}^-(\Omega_n) = \hat{\theta}^-(\Omega_{n-1}) = \dots = \hat{\theta}^-(\Omega_j)], \end{aligned} \quad (5)$$

as well as analogous definitions for $\mathcal{P}_1^+(j)$ and $\mathcal{P}_2^+(j)$. We make use of these when we wish to stop on the basis of two levels, an estimated MTD together with a corresponding co-MTD estimate. To align with Section 3.1, these quantities can be converted to an odds to use as an early stopping criterion.

In order to evaluate the probabilities of different paths we first construct a binary tree of all possible outcomes of the trial between patient j and the dose that would be recommended for patient $n+1$. We then write:

$$\hat{\mathcal{P}}_2(j) = \sum_{c \in \mathcal{C}(j)} \{1 - \psi(x_{j+1}, \hat{a}_j)\}^{n-j-|c|} \{\psi(x_{j+1}, \hat{a}_j)\}^{|c|},$$

where $\mathcal{C}(j)$ denotes the set of paths in the tree for which the level recommended to subject $n+1$ is the same as the level that has been recommended to subject j , and for which there is no change in level for the last $n-j$ subjects. The number of toxicities encoded by a path $c \in \mathcal{C}(j)$ is written $|c|$. We next note:

Lemma 5—*The probability that we remain at the same level, x_{j+1} , following the inclusion of a further $n - j$ patients, is estimated consistently by $\widehat{\mathcal{P}}_2(j)$.*

The lemma is shown in O’Quigley and Reiner (1998). In that paper the authors also considered $\mathcal{P}_1(j)$. To our knowledge, this probability has never been deeply studied, and it would likely be difficult to do so. This is due to the fact that if we remain at a single level, then the probability of toxicity at that level is well estimated. Once we expand our investigation to levels removed from the current estimate of the MTD, we are faced with a very difficult estimation problem, even if just recommending a single level for the DEC. This will be even more complex when studying two dose levels.

3.2.1 Evaluation of tree based rules—The difficulties in projecting forward are mostly of a combinatorial nature. The number of possibilities grows exponentially such that, in practice, we are not able to look too far ahead in the number of patients. In relatively simple cases we can work out the relevant probabilities. We can also take on board the error in estimates and check to what extent the final recommendations depend on these errors. This is further developed in Supplemental Section S1. While such calculations are possible, it remains unclear whether the potential gains in further understanding the precision offset the additional complexity in implementation. This warrants further investigation.

In way of illustration, suppose that the goal is to identify one dose level for the DEC. As shown in Figure 4, dose level 3 is recommended to patient 18 using \hat{a}_{17} , the estimate after 17 patients have accrued. Dose level 2 or level 3 will be recommended for patient 19 if patient 18 does or does not have a DLT, respectively. Dose recommendations are recursively identified until the final $n+1$ recommendation. As shown at the bottom of this figure, $\widehat{\mathcal{P}}_2(17)$, the estimated probability that level 3 is the recommended level for all remaining patients is $0.14+0.45=0.59$, or an odds of 1.42.

As shown in Figure 5, this approach can be similarly applied when the goal is to identify two levels for the DEC. After patient 17, the recommended two levels are 2 and 3. $\widehat{\mathcal{P}}_2^-(17)$, or equivalently $\widehat{\mathcal{P}}_2^+(17)$, is $0.17+0.15+0.14=0.47$, or an odds of 0.87.

Whether or not to bring the escalation part of the study to a halt and initiate the DEC on the basis of $\widehat{\mathcal{P}}_2(17)$ or $\widehat{\mathcal{P}}_2^-(17)$ will be specified in the study protocol using a predefined odds threshold.

3.3 Allocation Limits

A very simple rule follows from the basic idea behind $\mathcal{P}(j)$, which is the number of consecutive times any level has been recommended. The more time spent at a dose level, the more likely it is, under sequential model updating, that this level will turn out to be the MTD. This idea is behind the simple stopping rule proposed by Goodman et al. (1995) in which, after some fixed number of patients have been consecutively treated at some level – and that same level would be recommended to the subsequent patient if enrolled – the study is brought to a close. Intuitively, the idea is to set a maximum number, k , of times any level

should be consecutively allocated. A trial is halted as soon as this maximum is reached or exceeded. Formally, let k_j denote the number of consecutive patients treated prior to patient j who received the current estimate of the MTD $\hat{\theta}(\Omega_j)$. Termination is called as soon as $k_j = k$, for some $j < n$. For large k , theoretical justification could follow the same reasoning underlying the rule for binary outcome trees.

The allocation rule naturally extends when recommending two dose levels for the DEC. Analogously, the trial will stop if the same two dose levels have been recommended for k consecutive patients.

One clear advantage to a rule based on allocation limits is its great simplicity. Once we have fixed the experimental design parameter k , no further calculation is needed. This contrasts sharply with the binary tree rule which necessitates combinatorial calculations and the inevitable use of computer algorithms. In practice k will be small, and so it is more useful to consider behavior via simulations when comparing these methods.

4 Simulations

The performance of the three CRM stopping rules were evaluated under various scenarios when the study goal was either to identify one dose level or two dose levels for further evaluation in the DEC. These scenarios were designed to investigate the stopping rules under a variety of different toxicity rates for a total of five dose levels. These scenarios are described in Table 1. The target toxicity rate α was set to 0.20, and the skeleton was selected based on Lee and Cheung (2009), which is also provided in Table 1.

The maximum sample size for all scenarios was 20 patients. When identifying a single dose level, the tree-based method used the odds corresponding to $\mathcal{P}_2(j)$, and the posterior probability approach used \mathbb{R}_1 . When identifying two levels, $\mathcal{P}_2^-(j)$ and \mathbb{R}_3 were used as stopping criterion. Using these, the simulated trial would stop early if the estimated odds exceeded 3. The allocation rule stopped if 6 consecutive patients were treated at the same level, and the recommend level if a subsequent patient accrued would remain the same (deemed a 6+1 allocation rule). All methods were first evaluated after 15 patients accrued. The accuracy of these stopping rules were benchmarked against the standard CRM with a fixed sample size of 20 patients.

4.1 Example Trial

To illustrate the stopping rules and how information accumulates as patients accrue, we selected a single trial from Scenario 2 when the goal is either to select one dose level or two dose levels for the DEC. As shown in the first plot in Figure 6, CRM correctly identified dose level 3 in the simulated trial using a fixed sample size of 20 patients (no early stopping rule) as the level with a toxicity rate closest to our target α of 0.2. As shown in the second plot, the dose levels 2 and 3 were recommended after 20 patients accrued when the goal is to identify two levels.

Implementing early stopping when recommending a single dose level, the binary tree approach stopped the study early after 18 patients accrued. Neither the posterior probability approach nor the allocation rule of 6+1 stopped early.

When selecting two dose levels, the posterior probability approach stopped early after 15 patients, the tree-based rule stopped after 16 patients, and the allocation rule stopped after 17 patients. All approaches selected dose levels 2 and 3 as the recommended levels for the DEC.

4.2 Operating Characteristics

For each of the five toxicity scenarios, the three different stopping rules were evaluated based on the proportion of the 10,000 simulated trials that selected each dose level as one of the levels for the DEC along with the average sample size. Figure 7 shows the simulation results when the goal was to identify one dose level as the MTD, and Figure 8 shows the results when the goal was to identify two dose levels, one as the MTD and the other as the co-MTD.

For Scenario 1 in the top row of Figure 7, CRM with a fixed sample size selected dose level 2, which true toxicity rate aligned with α , in 55% of simulated trials. Importantly, there was minimal-to-no decrease in performance when one of three early stopping rules was implemented: the binary tree approached selected the correct dose level 54% of the simulated trials, the 6+1 rule 54%, and the posterior probability rule 55% of the time. As shown in the subsequent figure, the binary tree and 6+1 stopping rules saved, on average, 2 patients, corresponding to a 10% reduction in sample size from the maximum of 20 patients. The posterior probability infrequently stopped early with an average sample size of 19.3.

Overall, similar results were observed for the other four different toxicity scenarios. There was a negligible decrease in accuracy for the binary tree and 6+1 stopping rules compared to CRM with a fixed sample size, and the overall decrease in the sample size for these methods ranged from 1.4 to 2.9 patients. The posterior probability approach rarely stopped, though did so most frequently in Scenario 5, saving, on average, 1 patient.

Figure 8 provides the same results when the goal is to identify two levels for the expansion cohort. Across the scenarios, a high percentage of trials selected the correct dose level as one of the two levels. Note the percentages associated D1-D5 are not mutually exclusive. All three methods were able to stop the trial earlier, on average, and with higher accuracy compared to when the goal is to select only one dose level for the DEC.

5 Discussion

Although well studied for phase II and III clinical trials, early stopping rules for phase I dose escalation trials are not often used outside of the 2-out-of-6 rule of the standard 3+3 algorithm-based design. This lack of implementation may be due to different underlying reasons. For one, stopping rules based on confidence intervals (Heyd and Carlin, 1999) are limited due to the large variance that accompanies binomial sampling. Also, a confidence interval-based approach somewhat indirectly answers the relevant question: how likely

is the trial's conclusions to change if accrual continues and the full sample size is reached. O'Quigley and Reiner (1998) addressed this question directly and worked out a combinatorial solution to the problem. While performance appears promising, the burden of calculation can be challenging, particularly when full accrual includes more than 5 or so additional patients. A further consideration, and a contributing factor to the reluctance to implement these methods, may be that most phase I trials already anticipate quite small sample sizes. Therefore, any gains, consequent upon the implementation of an efficient stopping rule, may not have a large practical impact.

However, the situation has changed over the last ten years (Iasonos and O'Quigley, 2015). This is driven by three main reasons: (1) early phase trials are no longer carried out on very few patients and will typically involve relatively DEC's (Manji et al., 2013; Dahlberg et al., 2014); (2) the start of the DEC is now likely to involve not just a single estimated MTD but, rather, two dose levels, an MTD and a co-MTD; and (3) while there was traditionally a significant delay between the different study phases in drug development – diminishing the benefit to saving 2–3 patients in the escalation design – phase I studies with DEC's have no such delay. Therefore, initiating the DEC early, if only a few patients ahead of the planned sample size, can be of clinical importance. In consequence, the relevance of an efficient stopping rule has very greatly increased, as there is a need to not treat additional patients in the initial escalation phase of the study once we know the dose or doses likely to advance to the DEC. Once provided with this information, we would like to either decrease the overall study sample size or include these additional patients in the expansion to aid in the estimation objectives of the DEC.

Our purpose in this work has two conflicting aims: (1) to make the practical evaluation simpler thereby enabling an easier implementation of already available dose stopping rules and, (2) to make the evaluation more comprehensive so that we now take on board the need to follow the paths of two dose levels, rather than a single level, to take forward to the DEC. We present analytic tools and a theoretical framework to address both aims. Stopping rules based on the posterior probability addresses the first aim. The different partitions of the parameter space in the posterior probability approach make it straightforward to calculate the various odds that can guide our decision making, which addresses the second aim.

Corollary 3 makes practical implementation of the stopping rules quite easy in that it provides a single sub-interval T_j over which to estimate whether the target toxicity rate falls within an interval between two dose levels. It is likely that this degree of simplification may not be attainable for more complex situations, for example when dealing with heterogeneous groups (Horton et al., 2019), or when dealing with the case of partial ordering that arises in combination studies (Wages et al., 2011). We have not studied this in any depth but it would be something worthy of investigation. Lemma 3 and lemma 4 will be of value in this respect. In this current work we have only considered the relatively more elementary situation of a simple dose finding set-up. The gains that can be made via the use of an efficient stopping rule are significant, dependent, at least in part, on the accrual rate, the associated costs of treatment, and when the DEC is initiated.

To this end, among the methods discussed here, our preference is to reason in terms of the probabilities of different levels being the MTD and/or the co-MTD, which has more of a Bayesian flavor. The approach of O'Quigley and Reiner (1998) is different and, rather than ask the question about how great are the chances that any level will, ultimately, turn out to be the MTD, they asked a more limited question: if we continue enrolling patients, what are the chances that we will settle at the current level. While this method performs well when identifying a single level as the MTD, the underlying question of this approach has less relevance in the context of taking more than a single dose forward to the DEC.

Lastly, with good overall performance, the 6+1 rule provides a simple approach to implementing an early stopping rule in a phase I study. However, the approach lacks a summary measure that an analyst can use to evaluate the evidence that the current dose level(s) are the MTD/co-MTD. We feel that such a measure will be useful when implementing trials in many clinical settings, including those with DEC's or dose optimization cohorts (Ratain et al., 2021). For example, if presented with a figure similar to Figure 3 while running a trial, an analyst may decide to take d_2 and d_3 to the DEC. However, depending on the severity of toxicities observed, another decision may be to take three levels, d_2 , d_3 , and d_4 , to the DEC given that the odds are similar for d_2 and d_4 . Or, alternatively, the analyst may decide to continue the phase I study to full accrual to better estimate the odds associated with each dose level. If using the 6+1 rule, there is no parallel way to review and summarize the trial data-to-date to make a more informed decision about which dose levels to take to the DEC. Such measures are readily available with posterior probability stopping rules, which is a distinctive strength of the approach.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This work was supported by the NIH Grant P30CA008748.

References

- Dahlberg SE, Shapiro GI, Clark JW, and Johnson BE. Evaluation of statistical designs in phase I expansion cohorts: the Dana-Farber/Harvard Cancer Center experience. *J Natl Cancer Inst*, 106(7), Jul 2014.
- Goodman SN, Zahurak ML, and Piantadosi S. Some practical improvements in the continual reassessment method for phase I studies. *Stat Med*, 14(11):1149–1161, Jun 1995. [PubMed: 7667557]
- Heyd JM and Carlin BP. Adaptive design improvements in the continual reassessment method for phase I studies. *Stat Med*, 18(11):1307–1321, Jun 1999. [PubMed: 10399198]
- Horton BJ, Wages NA, and Conaway MR. Shift models for dose-finding in partially ordered groups. *Clin Trials*, 16(1):32–40, 02 2019. [PubMed: 30309262]
- Iasonos A and O'Quigley J. Clinical trials: Early phase clinical trials-are dose expansion cohorts needed? *Nat Rev Clin Oncol*, 12(11):626–628, Nov 2015. [PubMed: 26441082]
- Iasonos A and O'Quigley J. Integrating the escalation and dose expansion studies into a unified Phase I clinical trial. *Contemp Clin Trials*, 50:124–134, 09 2016. [PubMed: 27393122]

- Iasonos A, Wages NA, Conaway MR, Cheung K, Yuan Y, and O'Quigley J. Dimension of model parameter space and operating characteristics in adaptive dose-finding studies. *Stat Med*, 35(21):3760–3775, 09 2016. [PubMed: 27090197]
- Lee SM and Cheung YK. Model calibration in the continual reassessment method. *Clinical Trials*, 6:227–238, 2009. [PubMed: 19528132]
- Manji A, Brana I, Amir E, Tomlinson G, Tannock IF, Bedard PL, Oza A, Siu LL, and Razak AR. Evolution of clinical trial design in early drug development: systematic review of expansion cohort use in single-agent phase I cancer trials. *J Clin Oncol*, 31(33):4260–4267, Nov 2013. [PubMed: 24127441]
- O'Quigley J. Continual reassessment designs with early termination. *Biostatistics*, 3(1):87–99, Mar 2002. [PubMed: 12933626]
- O'Quigley J and Reiner E. A stopping rule for the continual reassessment method. *Biometrika*, 85:741–48, 1998.
- O'Quigley J, Pepe M, and Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics*, 46(1):33–48, Mar 1990. [PubMed: 2350571]
- Ratain MJ, Tannock IF, and Lichter AS. Dose Optimization of Sotorasib: Is the US Food and Drug Administration Sending a Message? *J Clin Oncol*, page JCO2101371, Sep 2021.
- Shen LZ and O'Quigley J. Consistency of Continual Reassessment Method Under Model Misspecification. *Biometrika*, pages 395–405, Jun 1996.
- Storer BE. Design and analysis of phase I clinical trials. *Biometrics*, 45(3):925–937, Sep 1989. [PubMed: 2790129]
- Wages NA, Conaway MR, and O'Quigley J. Continual reassessment method for partial ordering. *Biometrics*, 67(4):1555–1563, Dec 2011. [PubMed: 21361888]




Figure 1:
Partitioning of $[A, B]$ used to define the sub-intervals S_j and T_j for five dose levels.

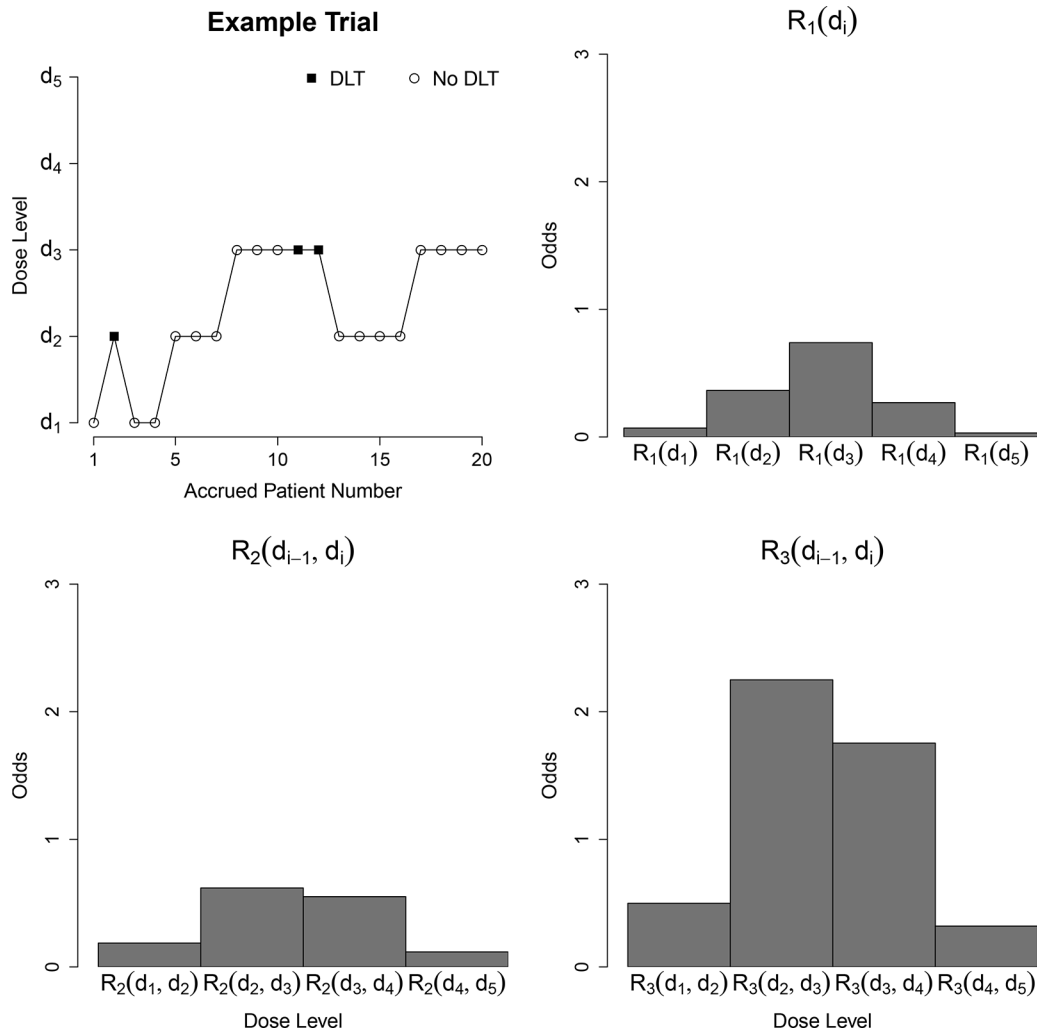


Figure 2: One example trial with a total of five dose levels. The first plot shows the dose-limiting toxicities (DLT) observed as patients are treated at each dose level. The remaining plots provide the three odds parameters R_1 , R_2 , and R_3 , as defined in Section 3.1. The odds are highest that d_3 is the MTD (R_1) and that d_2 or d_3 correspond to the MTD as opposed to neither (R_3). Lastly, the odds (R_2) are similar that the target MTD rate lies between d_2 and d_3 or d_3 and d_4 . If the goal were to select two dose levels for the dose expansion cohort, d_2 and d_3 would be the likely candidates based on these estimates.

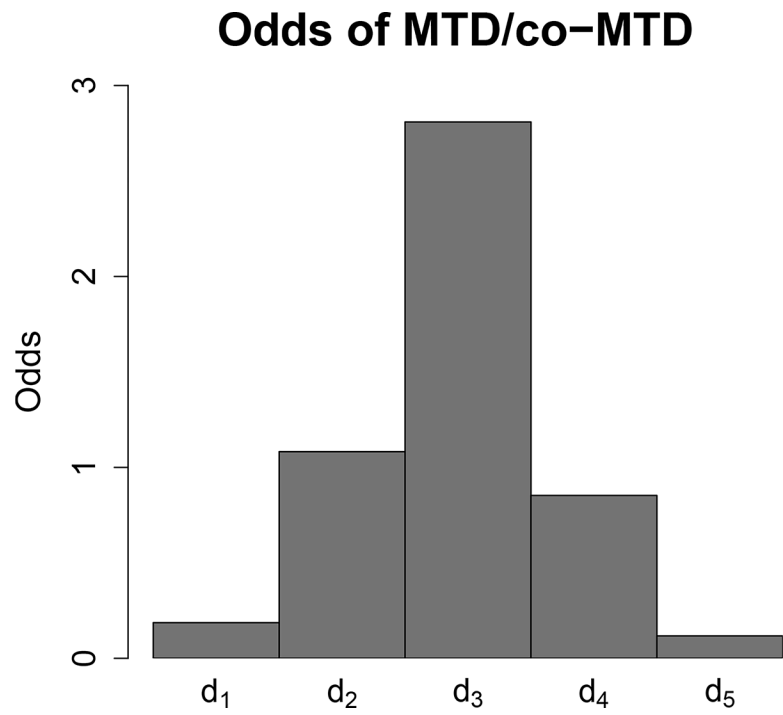


Figure 3:
The odds that each dose level is either the MTD or the co-MTD using the same example trial as Figure 2.

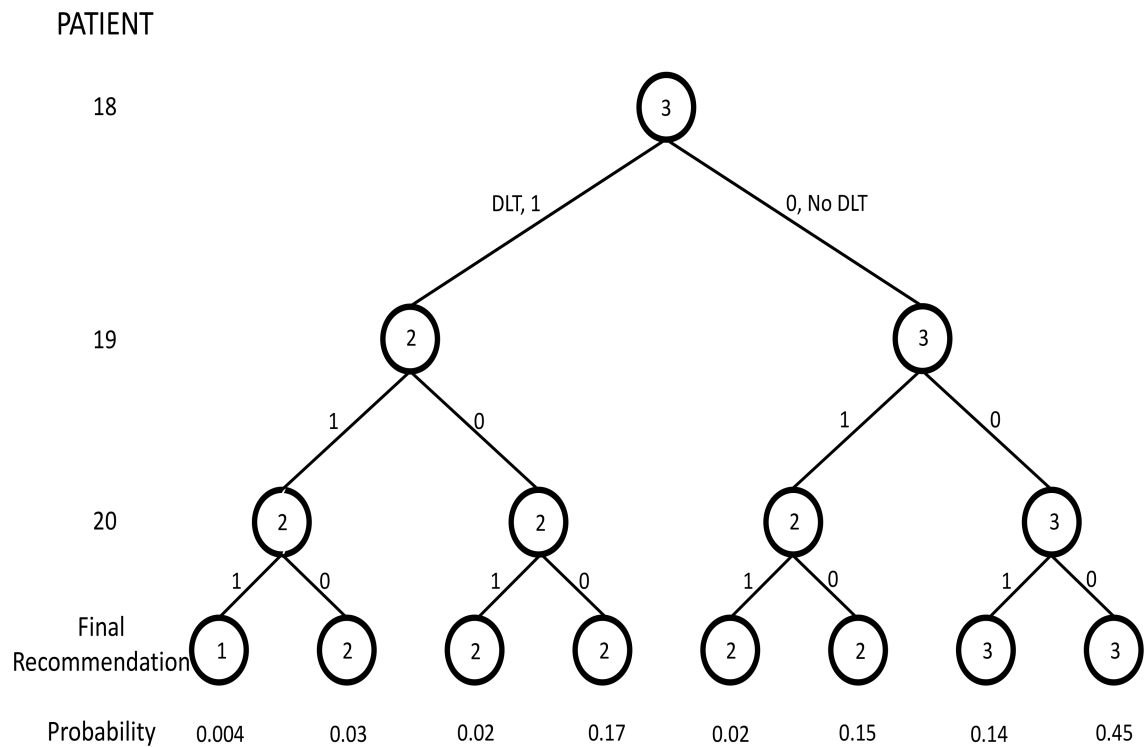


Figure 4:

Construction of the binary stopping rule after 17 patients have been treated when the goal is to identify one dose level for the DEC. Using \hat{a}_{17} , the recommended dose level for patient 18 is 3. From the tree, the probability that all remaining patients receive level 3 in addition to this level being the final recommended dose is 0.59 or an odds of 1.42. If the threshold for stopping is an odds of 3, this trial would continue accruing patients.

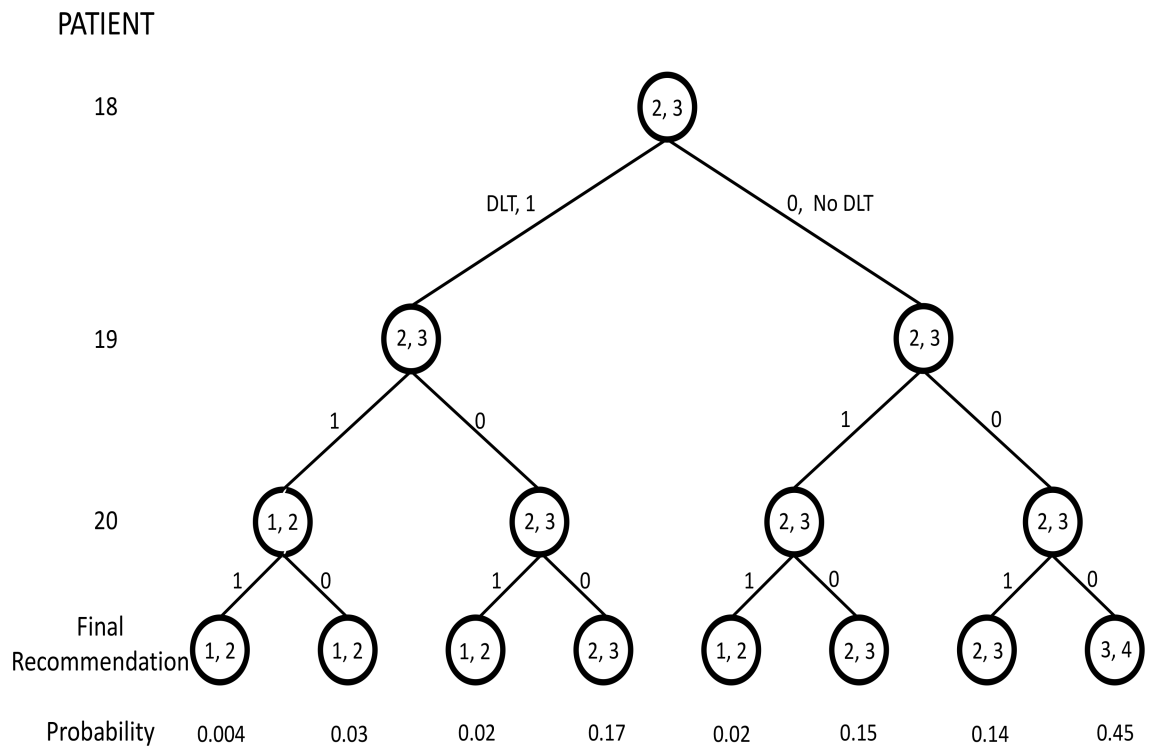


Figure 5: Construction of the binary stopping rule when the goal is to identify two dose levels for the DEC. The recommended two dose levels after 17 patients have been treated are 2 and 3. The probability that these two dose levels remain the two recommended levels throughout the remainder of the study is 0.47 or an odds of 0.87.

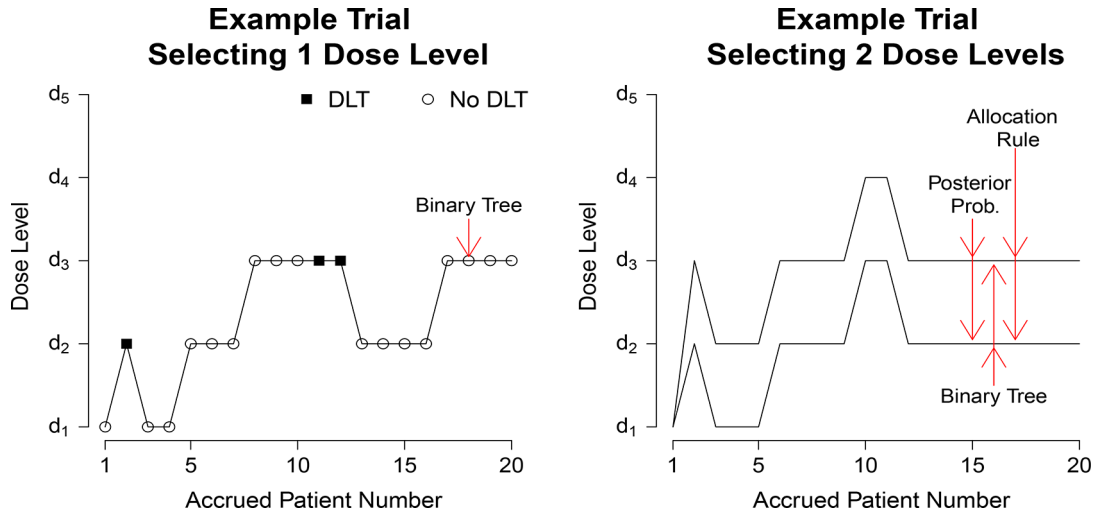


Figure 6: The first plot shows the observed dose-limiting toxicities (DLTs) based on the assigned dose level when the goal is to identify one dose level for the dose expansion cohort. The rule based on the binary tree approach (\mathcal{P}_2) stopped the study after the 18th patient selecting level d_3 . Neither the allocation rule nor the posterior probability approach stopped the study before it reached full accrual (N=20). The second plot shows the same trial but when the goal is to identify two dose levels for the DEC. The stopping rule based on the posterior probability stopped the study after patient 15, selecting d_2 and d_3 for the DEC. The binary tree approach and the allocation rule stopped the study after patient 16 and 17, respectively, both selecting d_2 and d_3 .

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Figure 7: Percentage of simulated trials that selected each of the five dose levels (D1-D5) as the maximum tolerated dose (MTD) along with average total sample size when selecting a single dose level for the dose expansion cohort (DEC). The four methods under evaluation include: the model-based design without early stopping (fixed 20); early stopping based on the binary tree approach (\mathcal{P}_2); early stopping based on the 6+1 allocation rule; early stopping rules based on the posterior probability.

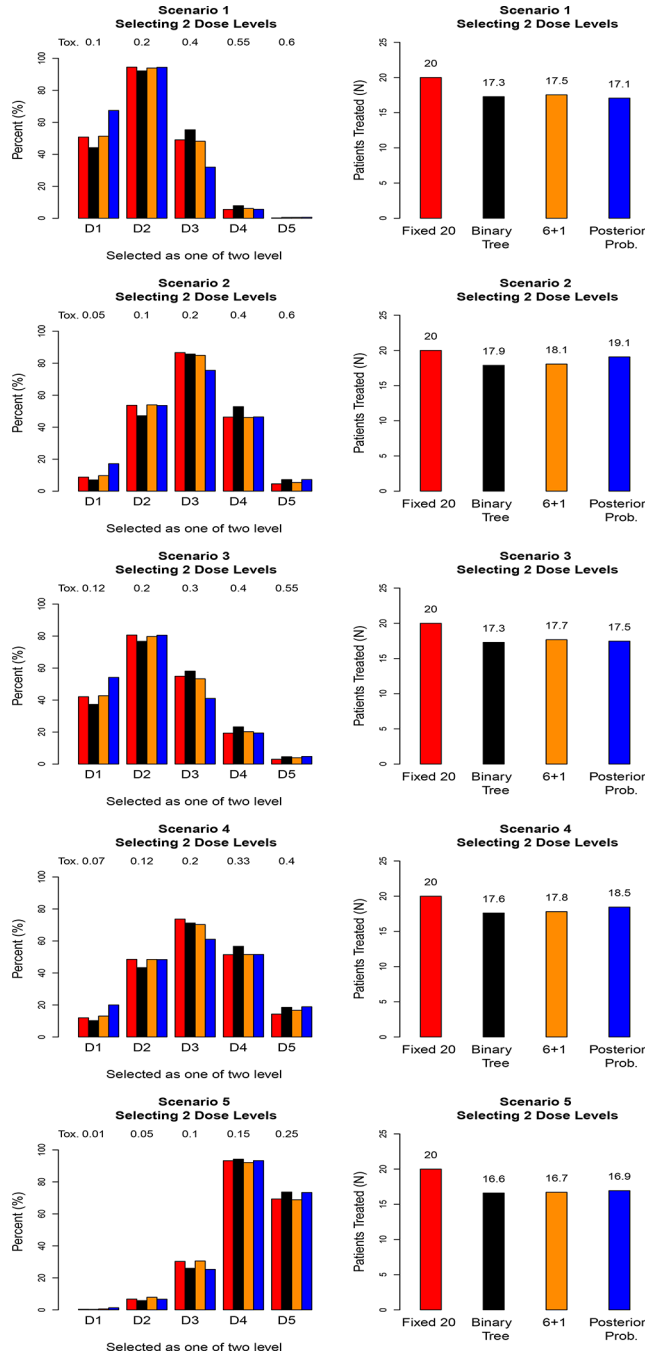


Figure 8: Percentage of simulated trials that selected each of the five dose levels (D1-D5) as the maximum tolerated dose (MTD) along with average total sample size when selecting two dose levels for the dose expansion cohort (DEC). The four methods under evaluation include: the model-based design without early stopping (fixed 20); early stopping based on the binary tree approach (\mathcal{P}_2); early stopping based on the 6+1 allocation rule; early stopping based on the posterior probability.

Table 1:

The parameters used for the simulation study including the skeleton values and true toxicities rates for the five dose levels.

	Level 1	Level 2	Level 3	Level 4	Level 5
Skeleton (α)	0.049	0.111	0.200	0.308	0.423
Toxicity Rates					
Scenario 1	0.10	0.20	0.40	0.55	0.60
Scenario 2	0.05	0.10	0.20	0.40	0.60
Scenario 3	0.12	0.20	0.30	0.40	0.55
Scenario 4	0.07	0.12	0.20	0.33	0.40
Scenario 5	0.01	0.05	0.10	0.15	0.25