

Controlled backfill in oncology dose-finding trials

Hakim-Moulay Dehbi^{a,*}, John O'Quigley^b and Alexia Iasonos^c

^a*Comprehensive Clinical Trials Unit at UCL, London, UK*

^b*Department of statistical science, University College London, London, UK*

^c*Memorial Sloan Kettering Cancer Center, New York, USA*

ARTICLE INFO

Keywords:

backfill
dose-finding
early phase
continual reassessment method
efficacy
response
plateau
oncology

ABSTRACT

The use of backfill in early phase dose-finding trials is a relatively recent practice. It consists of assigning patients to dose levels below the level at which the trial is at. The main reason for backfilling is to collect additional information on efficacy, in order to assess whether a plateau may exist on the dose-efficacy curve, which is possible with molecularly targeted agents or immunotherapy in oncology. Recommending a dose level lower than the maximum tolerated dose at the end of study could be supported in such situations. How to best backfill patients is not yet established. In this paper we propose to randomise backfill patients between the doses below the dose where the dose-escalation experiment is at. A refinement of this would consist of discontinuing dose levels that show insufficient efficacy compared to higher dose levels, starting at dose level 1 and repeating this process sequentially. At the end of the study, data from all patients (both the backfill patients and the dose-finding patients) are used to estimate the dose-efficacy curve. The fit from a change point model is compared to the fit of a monotonic model to identify a potential plateau. Using simulations, we show that this approach can identify the plateau on dose-efficacy curves when such a plateau exists, allowing the recommendation of a dose level lower than the maximum tolerated dose for future studies. This contribution provides a methodological framework for backfilling, from the perspective of both design and analysis.

1. Introduction

As a key component of the drug development process in oncology, Phase I dose-finding trials have traditionally focused on assessing safety and any drug related toxicity due to cytotoxic agents. The importance of the dose-finding component to the whole process has been the subject of recent work by Conaway and Petroni [2]. In recent years, phase I trials have increasingly attempted simultaneous evaluation of efficacy endpoints in various patient populations via dose-expansion cohorts, with the aim of identifying which patient population(s) to be the focus of future studies [6]. With chemotherapy the dose-response curve was traditionally assumed to be monotonic and increasing with dose thus the working assumption was that a higher dose would be more efficacious, as long as it remains well tolerated. In other words toxicity was used as a surrogate for efficacy. In such settings, the aim was to identify the maximum tolerated dose (MTD). More recently with molecularly-targeted agents such as monoclonal antibodies the working assumption that the dose-efficacy and dose-toxicity curves are increasing with dose has been challenged [3]. Targeted therapies have shown that they can have a dose level, above which the dose-response curve is no longer increasing. This level will frequently be below the MTD [15, 16]. For this reason the concepts of biologically effective dose (BED), or optimal biological dose (OBD), have been increasingly used in dose-finding studies of monoclonal antibodies. Similarly, some immune therapy agents have shown signs of a dose-response curve that is rather flat or reaches a plateau in the sense that efficacy is similar at nearby dose levels [1]. If the dose-efficacy curve displays a plateau, a dose level lower than the MTD may be

recommended as the recommended phase II dose (RP2D), as it reduces the risk of toxicity for patients, incurs no notable penalty efficacy-wise and can potentially be more cost-effective.

While using toxicity data to guide the dose-finding exercise, there are two main ways to assess efficacy in settings where we hypothesize that a plateau on the dose-efficacy curve may exist. Firstly, we could identify the MTD in a dose-escalation study and at end of the dose-escalation part, include one or several expansion cohorts near the MTD [6, 5, 8, 7, 10]. Secondly, during the dose-escalation study, we could simultaneously backfill dose levels below the dose level where the experimentation is at, as these doses would have been declared “safe”. The objective behind the inclusion of backfill patients would be to collect additional information on efficacy on the lower part of the curve below the MTD. This second approach has recently been implemented in a few recent studies [12, 11, 18]. For example, in the dose-escalation phase of a recent dose-finding study of carfilzomib and panobinostat for patients with relapsed/refractory multiple myeloma, there were four possible dose levels. Dose levels 1 and 2 were each backfilled with three patients. Level 1 was backfilled when the experimentation was at level 3, and level 2 when the experimentation was at level 4. The timing of backfill was pre-specified in the study protocol.

Currently, there appears to be no statistical framework in the literature to provide structure to backfill [9]. Yet spreading or experimenting far away from the MTD still must respect statistical, coherency and ethical principles. Here, we consider approaches that allocate backfill patients in a way that maximizes their potential benefit, while still allowing the evaluation of the shape of the dose-efficacy curve. One question of particular interest is whether randomization can be used to allocate backfill patients to doses, and which doses

*Corresponding author
ORCID(s):

should constitute the backfill set. If randomization is to be used, it would be important to discontinue backfilling to the first dose level (and subsequently to the second dose level et cetera) as soon as there is enough evidence that its efficacy is lower than the efficacy of other dose levels.

2. Statistical features of the design

In this Section we describe the main elements of our statistical approach. In Section 2.1 we gather together the main notation that we use. Following this we describe in Section 2.2 how to implement an actual trial and, finally, Section 2.3 indicates how the recommended dose is obtained.

2.1. Notation and observations

The study concerns m dose levels and a maximum of n patients. The goal is to locate the recommended phase II dose (RP2D) in an efficient and ethical way. The j th patient, $j = 1, \dots, n$ is allocated to one of the m dose levels. This level is denoted $X_j \in \{d_1, \dots, d_m\}$. We define Y_j to be a binary variable (0, 1) where 1 denotes that patient j suffered a dose-limiting toxicity (DLT) and is zero otherwise. The probability of DLT at dose level $d_i, i = 1, \dots, m$, is denoted $R(d_i)$ so that, $R(d_i) = \Pr(Y_j = 1 | X_j = d_i)$ and where $R(d_i) < R(d_k)$ when $i < k$. For the j th patient we will also record two additional binary variables, V_j and W_j . The first of these indicates tumour response. The true probability of a positive tumour response at $X_j = d_i$ is denoted by $Q(d_i) = \Pr(V_j = 1 | X_j = d_i), i = 1, \dots, m$. The second variable, W_j , indicates which set of patients they are part of: 1 for dose-finding patient and 0 for backfill patient. The data can be represented sequentially by $\Omega_j = \{(x_\ell, y_\ell, v_\ell, w_\ell), \ell = 1, \dots, j\}$. We are interested in the running tally of successes at dose levels $d_i, i = 1, \dots, m$, and consider the following two quantities:

$$\bar{V}_j(d_i) = n_j^{-1}(d_i) \sum_{\ell=1}^j V_\ell I(d_i = x_\ell),$$

$$\bar{V}_j(d_i^*) = n_j^{-1}(d_i^*) \sum_{\ell=1}^j V_\ell I(d_i > x_\ell),$$

where I is the indicator function, $n_j(d_i) = \sum_{\ell=1}^j I(d_i = x_\ell)$, $n_j(d_i^*) = \sum_{\ell=1}^j I(d_i > x_\ell)$ and where we define $0/0 = 0$.

As far as the rates of DLT are concerned we lean on a one-parameter model-based design, $\psi(x_j, a)$ where a is the model parameter, which will help us identify the MTD. In terms of efficacy data, our model $\phi(x_j, b)$ where b is a vector of parameters, can take two different shapes, one with a plateau and one without a plateau, each one resulting in a specific approach to dose recommendation.

During the study, we need to dynamically determine the set of acceptable doses for backfill based on accumulating efficacy data. The highest dose level in this set is always one dose level lower than the dose where the dose-finding experiment is at. The lowest dose level is level 1 as soon as the

second dose level is opened for dose-finding patients. Using our running tallies $\bar{V}_j(d_i)$ and $\bar{V}_j(d_i^*)$, we define the hypothesis H_0 to help us decide whether the first dose level, and subsequently the second dose level et cetera, can be removed from the set. In statistical terms we write $H_0 : E[\bar{V}_j(d_i)] = E[\bar{V}_j(d_i^*)], i = 1$. This process is iterative. As long as we remain under H_0 then backfill to the lowest dose level can be justified as there is not enough evidence that signals higher efficacy rates at higher levels. When H_0 is rejected for the first time, we discard the first dose level from the set, and redefine H_0 , now as $E[\bar{V}_j(d_i)] = E[\bar{V}_j(d_i^*)], i = 2$. This results in the sequential elimination of doses that show comparatively insufficient efficacy. We present in Section 2.2 how we intend to test H_0 in practice.

At the end of the study, we contrast the following two hypotheses to make a final dose recommendation that combines efficacy and toxicity considerations:

- $H_1 : Q(d_i) < Q(d_k)$ when $i < k$.
- $H_2 : \text{There exists } h, 1 < h < m, \text{ such that } Q(d_i) < Q(d_k) \text{ when } i < k < h \text{ and } Q(d_i) = Q(d_k) \text{ when } i \geq h, k \geq h.$

Under H_2 the probability of tumour response is assumed to increase with dose up to a certain dose level, and to then be reasonably well approximated by a plateau. There are a number of different ways in which we can decide between H_1 and H_2 which of the two is best supported by the efficacy data at the end of the study. We provide ways to test these hypotheses in Section 2.3.

In our framework, we have 3 different expressions of the likelihood to consider. The first relates to the dose-DLT model $\psi(x_j, a)$. The second and third expressions relate to the dose-efficacy curve and correspond to the hypotheses outlined by H_1 and H_2 . From a frequentist perspective our running estimate of the parameter a comes from the estimating equation, $U_n^{(1)}(\hat{a}) = 0$ where:

$$U_n^{(1)}(a) = \sum_{j=1}^n w_j \left[y_j \frac{\psi'}{\psi} \{x_j, a\} + (1 - y_j) \frac{-\psi'}{1 - \psi} \{x_j, a\} \right]. \quad (1)$$

Under H_1 , our running estimate of the vector b comes from the zeros of $U_n^{(2)}(\hat{b})$ where:

$$U_n^{(2)}(b) = \sum_{j=1}^n \left[v_j \frac{\phi'}{\phi} \{x_j, b\} + (1 - v_j) \frac{-\phi'}{1 - \phi} \{x_j, b\} \right]. \quad (2)$$

Under H_2 , and, for given $h, 1 < h < m$, our running estimate of the vector b comes from the zeros of $U_n^{(3)}(\hat{b})$ where:

$$U_n^{(3)}(b) = \sum_{j=1}^n \mathcal{I}_{jh} \left[v_j \frac{\phi'}{\phi} \{x_j, b\} + (1 - v_j) \frac{-\phi'}{1 - \phi} \{x_j, b\} \right] \\ + \sum_{j=1}^n (1 - \mathcal{I}_{jh}) \left[v_j \frac{\phi'}{\phi} \{d_h, b\} + (1 - v_j) \frac{-\phi'}{1 - \phi} \{d_h, b\} \right], \quad (3)$$

in which we use \mathcal{I}_{jh} to abbreviate $I(x_j < d_h)$. $U_n^{(1)}(a)$ is orthogonal to both $U_n^{(2)}(b)$ and $U_n^{(3)}(b)$ but, of course, $U_n^{(2)}(b)$ and $U_n^{(3)}(b)$ are not orthogonal to one another.

From an operational standpoint we have two sequential dose allocation schemes taking place iteratively during the study. First, we make use of the continual reassessment method (CRM) [13] to identify the MTD, using DLT data from dose-finding patients, and, secondly, we carry out sequential testing of H_0 . These schemes impact the allocation of dose-finding and backfill cohorts, which we consider in the next section.

2.2. Allocation of cohorts of patients during the study

We need keep track of the two distinct cohorts of patients:

1. So called *dose-finding patients* are those whose observed DLT data contribute to the dose-escalation decisions during the course of the study. If patient j belongs to this cohort then $W_j = 1$.
2. The *backfill patients* are the patients that are randomised to dose levels below the level that is explored with dose-finding patients. The observed DLT data for these patients do not influence the dose-escalation decisions during the study. If patient j belongs to this cohort then $W_j = 0$.

In practice dose-finding patients are recruited first when a new dose level is opened following dose escalation. Thereafter, backfill patients are individually randomised to the dose levels that are part of the backfill set, in the way that described below in this Section.

Dose-finding patients

We make use of the Bayesian CRM as model-based design [13, 17] to allocate dose-finding patients to the available dose levels. The dose-finding exercise starts at the first dose level, and dose-skipping is not allowed. Within these constraints, the allocation of the dose level for the next cohort of patients is to the dose level with an estimated DLT risk closest to the target toxicity level (TTL) denoted α . To model the dose-DLT relationship, we assume a working model of the form $\psi(d_i, a) = \beta_i^a$, where β_i are the standardised doses (skeleton) representing the discrete dose levels d_i and a follows an exponential distribution with λ parameter equal to 1. Given that the prior mean of the model parameter is equal to 1, the standardised dose is equal to the prior estimate of DLT probability at each dose level. We employ the posterior plug-in mean estimate of DLT probability for dose-escalation decisions [13].

Backfill patients

We use the dose-finding patients to guide us to our best current estimate of the MTD. When this estimate is at level 2 or higher then we can address the question of whether and how to *backfill patients* to those levels lower than the current estimate of the MTD. We will associate a probability of a backfill patient being allocated to any of these lower levels, including a probability of zero when the evidence is such as to suggest insufficient efficacy at that level. Specifically, we proceed as follows:

1. For as long as H_0 has not been rejected, we randomise, according to $G(d_i)$ to all dose levels below the level where the dose-finding experimentation is currently at. We use a discrete uniform distribution for $G(d_i)$, resulting in equal probability of randomisation at each dose level of the set of acceptable doses for backfill;
2. As soon as H_0 has been rejected for the first dose level, we randomise, according to $G(d_i)$, $i = 2, \dots, m$ to all dose levels excluding the lowest level. The same idea then applies as we continue the study.

Any acceptable test can be used to sequentially monitor H_0 . In our case we chose to carry out a test of H_0 on the basis of a Bayesian Beta-Binomial model, assuming a uniform prior on the (0,1)-interval for the true (but unknown) probabilities of response that govern both $\bar{V}_j(d_i)$ and $\bar{V}_j(d_i^*)$. We specify the test of H_0 by picking some difference between the rates at the lowest level available and the higher levels, and an associated probability of it. Specifically we chose a difference of zero and a posterior probability of a difference of 80% as threshold to reject H_0 .

2.3. Dose recommendation at study completion

The sequential allocation to the current estimate of the MTD results in a final estimate of the MTD at study completion, this estimate being the level the next patient would have received, had one further patient been included in the study. No additional model fitting is carried out. The data involved in this estimate come from information from the dose-finding patients. In particular the information from *backfill patients* is not used. However, the determination of RP2D will make use of all of the observations, those tumour response observations from the *backfill patients* as well as those from the *dose-finding patients*. These observations are used to calculate the observed efficacy rates. They are also used in the more flexible models for $\phi(x_j, b)$. Not only do we use all of the data to estimate the unknown components of the vector b but also we will use the full data set to decide which of two forms for $\phi(x_j, b)$ has the strongest support. Choosing between competing model forms come under the heading of model choice. The first model form is one where the success rate is monotonically increasing over the doses and any choice for $\phi(x_j, b)$ that satisfies the strict monotonicity constraint could be used. Here, and in our illustrations below, we appeal to the well-known logistic transformation in which $b = (\beta_0, \beta_1)$, where β_0 is the intercept parameter and $\beta_1 > 0$, the slope parameter.

The second form of interest for modelling $\phi(x_j, b)$ is slightly more complex as it involves a plateau occurring for doses higher than h , a conceptual continuous value for a dose. We treat h as an unknown parameter so that one possible parameterization for $\phi(x_j, b)$ would specify b as of dimension 3, with a logistic specification for doses below h , followed by a plateau for those doses above h . When $d_1 < h < d_m$, the model is well-specified and all 3 parameters, β_0 , β_1 and h can be estimated. We might push h to the limit of the above parameterization, i.e., $h = d_m$. In this case we have a strictly monotonic dose-efficacy curve. We specify the second form

in the following way:

$$\log \frac{Q(d_i)}{1 - Q(d_i)} = \beta_0 + \beta_1 \{I(x_j \leq h)x_j + I(x_j > h)h\}, \quad (4)$$

$$d_1 < h < d_m, \beta_1 > 0.$$

In some formal way then, we can express the problem of model choice as one between two hypotheses: the first in which $h = d_m$ (for which we use the first model specification, i.e. the logistic model without plateau) and the second in which $d_1 < h < d_m$. Our estimated RP2D will depend upon which one of these two hypotheses is best supported by the data. We can include any prior information on the parameters of the two models. In our illustrations the prior distribution for h is taken to be uniform on the range of explored dose levels. For β_0 and β_1 we may use minimally informative prior distributions, using normal distributions centered at zero with large variances. The same prior distributions for β_0 and β_1 were used in the models with a change point and the models without a change point in our illustrations below. How to make best use of the posterior distributions is a topic we have not studied and, for our examples, we simply used the posterior mean. From the computational angle, note that a model without a plateau is a conventional Bayesian logistic regression that can be fit with standard software.

Our running estimate of the MTD will also be our RP2D if we conclude that a model without a plateau provides the best explanation of the data. If though we conclude in favor of a model with a plateau then it makes sense to choose a dose as low as possible but on the plateau. This will have the advantage of lowering the rate of adverse effects but ought only to make a very small difference to the rate of efficacy, or none at all in the case of a pure plateau. In this case we revise our final estimate of the MTD from the running estimate based on the *dose-finding patients* alone to a lower dose. This is then the RP2D. Although unlikely, it is theoretically possible for the estimated plateau to take place beyond the currently estimated MTD. We say unlikely because we have little in the way of observations beyond the MTD. If this should occur then we would take the estimated MTD as the RP2D and make no use of the plateau. Model choice is a very broad topic and the many approaches within this topic, in particular as they relate to change point models, are relevant to our work here. We have not carried out any study into the many options available and chose one well understood and popular method based on leave-one-out cross-validation [4, 20]. This is a technique that selects the model with greatest predictive ability on an empirical basis.

3. Simulations

3.1. Illustration of a single hypothetical dose-finding trial with backfill

We present here as illustration a dose-finding trial that was simulated according to the dose-efficacy and dose-DLT curves shown in Figure 1A. We set the target toxicity level α at 25%. From a toxicity perspective, dose level 5 corresponds to the MTD. The dose-efficacy curve presents a

plateau from dose level 3 at 25% for tumour response rate. Recommendation-wise at the end of study, the RP2D should be dose level 3, since it provides the same efficacy as dose level 5 but with a lower DLT rate.

In this simulated trial a one-parameter Bayesian CRM model was employed to model the dose-DLT curve, as described in Section 2.2. The CRM's skeleton was equal to the true probabilities of toxicity. The study started at dose level 1 and dose-skipping was not allowed.

The evolution of the trial is shown in Figure 2. Ten dose-finding cohorts of three patients and nine backfill cohorts of three patients were allocated during the study, for a total of 57 patients. The first DLT was observed at dose level 4 in a patient from the fourth dose-finding cohort. Until then no positive response had been seen in the dose-finding patients. At this point backfill patients had been randomised to dose levels 1 to 3. No DLTs and no positive responses had been seen in the backfill patients. The study continued to dose level 5, with three additional dose-finding patients on dose level 5, and randomisation of 3 backfill patients to dose levels 1, 3 and 4. This backfill patient at dose level 4 experienced the first positive tumour response in the study. The trial continued in this fashion until 57 patients in total had been dosed, as shown on Figure 2. Randomisation to dose level 1 was discontinued for backfill patients after the 39th patient. Five responses had been seen by then at dose levels 2 and higher among 30 patients, and none among the 9 patients at dose level 1. The posterior probability of a difference in probability between dose 1 on the one hand, and dose levels 2 and higher combined on the other hand, was 84.8%, greater than the threshold of 80% proposed in Section 2.2.

At the end of the study, dose level 5 was the MTD with respect to the TTL of 25%. The posterior plug-in estimate of DLT risk was 22.5% at dose level 5, and 32.7% at dose level 6. Efficacy-wise, a model with a plateau was retained as better fit to the data, in comparison to a model without plateau. Indeed the comparison of predictive accuracy, using leave-one-out-cross-validation, between a change point model and a conventional logistic regression for efficacy (assuming monotonicity between probability of response and dose level) favoured the change point model. The mean of the posterior distribution of the change point was 3.72. The resulting modeled dose-efficacy curve is shown in red with its confidence band on Figure 1B, in addition to the true and observed dose-efficacy curves.

Combining efficacy and toxicity considerations, the RP2D was dose level 4. Its modeled efficacy rate was equal to that of dose level 5. From a toxicity perspective, the posterior plug-in estimates of DLT risk were 14.6% and 22.5% at dose levels 4 and 5 respectively.

3.2. Simulations of six different scenarios

To assess the performance of our proposed backfill strategy, we considered six scenarios, each with seven dose levels. These are presented in Figure 3 and described quantitatively in Table 1. In all scenarios the target toxicity level was set at 25%.

In the first two scenarios (scenarios A and B), the RP2D should be dose level 3, given that it corresponds to the start of the plateau efficacy-wise at a lower dose level than the MTD at dose level 5. Indeed, recommending dose level 3 would provide an equal efficacy rate compared to level 5, with reduced toxicity. In scenarios C and D, the RP2D should be dose level 7, corresponding to the MTD, given that there is no plateau on the dose-efficacy curves. In the last two scenarios (E and F), given that the plateau starts at the MTD (scenario E) or a dose level higher than the MTD (scenario F), the MTD itself, which is dose level 4, should be the RP2D.

As in Section 3.1, a Bayesian one-parameter power model was employed to model the dose-DLT curve. A Gamma(1,1) prior was used for the model parameter. The CRM's skeleton was equal to the true dose-DLT curve in the first two scenarios, but was different to the true dose-DLT curve in the remaining four scenarios.

The simulated trials started at dose level 1. Dose-skipping was not allowed. Ten dose-finding cohorts of three patients were used, and randomised backfill cohorts of three patients were used as long as the experimentation was at dose level 2 or higher. In order to discontinue dose levels, a threshold of 80% was used for the probability of a difference in efficacy rates between the first dose level (and subsequently the second dose level etc) and the remaining dose levels combined. 1000 simulations were performed per scenario. The results are reported in Tables 1, 2 and 3.

In scenarios A and B, dose levels 3 and 4 together were recommended 75% and 78% of the time respectively. The MTD toxicity-wise, dose level 5, was recommended in 16% and 15% of the simulations in scenarios A and B respectively. A plateau on the dose-efficacy was identified 70% and 72% of the time in scenario A and B respectively, leading to a RP2D lower than the estimated MTD in 66% and 69% of the simulations. In terms of discontinuation of levels for the backfill patients, only 14% of the simulated trials in scenario A did not discontinue the first dose level. This percentage was 39% in scenario B, due to the less steep slope of the dose-efficacy curve prior to the plateau. Discontinuation percentages for the first, second and the remaining dose levels are provided in Table 3.

In scenarios C and D where there was no plateau on the dose-efficacy curve, the true MTD, dose level 7, was recommended in 51% and 60% of the simulations respectively. In scenario C, dose level 4 was recommended 30% of the time, due to the shallow nature of the dose-efficacy curve. The first dose level was not discontinued for the backfill patients in 27% and 11% of the simulations for scenarios C and D respectively, and was discontinued in 33% and 26% respectively. In scenario D, in 76% of the simulations, either dose level 1, or dose levels 1 and 2, or dose levels 1, 2 and 3, were discontinued for the randomisation of backfill patients.

In both scenarios E and F, dose level 3 was recommended 45% of the time. Dose level 4, the true MTD, was recommended 36% and 34% respectively. A plateau was identified in 56% of the simulations for both scenarios, and led to a

RP2D at a level lower than the estimated MTD in 42% and 39% of the simulations in scenario E and F respectively. The first dose level was not discontinued in 33% and 36% of the simulations in scenario E and F respectively, and was discontinued 31% of the time in both cases.

These results support the fact that a systematic and ethical approach to backfilling lower dose levels can be put in place, in order to accurately find the RP2D. In most scenarios studied, there were no losses in terms of accurately finding the MTD or allocating backfill patients at the right levels as determined by efficacy, while the sample size remained feasible and small as consistent with clinical protocols of early phase trials.

4. Discussion

The introduction of backfill into early phase dose-finding trials is a relatively recent development [12, 11, 18, 9]. It is important to note that the goals of backfill are very different to the ones of approaches such as EffTox [19] that use both efficacy and toxicity data to guide the dose-finding exercise. With backfilling, the dose-escalation process is guided by toxicity, even though in practice pharmacokinetics and/or pharmacodynamics data can be formally or informally incorporated into dose-(de)escalation decisions. Efficacy data in the form of partial or complete responses is used to refine the backfill set during the escalation process, and at the end of the study to make a dose recommendation. Comparing operating characteristics between a backfill design and EffTox via simulations would be challenging on two aspects. Firstly the backfill design aims at spreading backfill patients on dose levels below the MTD. An EffTox approach aims at concentrating the patients on the optimal dose by incorporating both efficacy and toxicity information in each dose-escalation decision. For this reason it is not easy to compare both approaches with respect to the proportion of patients allocated to dose levels. Secondly, it is possible to construct situations where the optimal dose depends on the method. In Figure 1, the optimal dose level is dose level 3 according to the backfill approach. However, with an EffTox approach, depending on chosen utilities, dose level 2 may be the optimal dose. Consequently, these approaches cannot be compared objectively with respect to the proportion of simulated studies choosing the optimal dose, as the definition of the optimal dose varies.

Part of our purpose in this paper is to argue that a development such as backfilling needs to take place within the framework of some formal statistical structure. At the most elementary level, we would like to be sure that the motivation for backfill is well grounded. Early in the study, before much dose escalation has taken place there may be little to concern us by allocating patients to levels lower than the current one: these levels are likely to be not far removed from the dose that is currently under scrutiny. Given the limited amount of data available at this point it is not necessarily feasible to discard any dose level from the backfill set due to insufficient efficacy. However, as we move away from the lowest levels, and as we gather more information, we can

find ourselves in a situation in which, not only is the need for an answer more compelling but our ability to provide an answer improves with each new inclusion. In this work we initially assume a simple structure whereby all the levels up to the current estimate of the MTD enjoy a comparable rate of efficacy. A sequential test of the hypothesis governing the rate of efficacy at the lowest level when contrasted against all the other levels allows the continuation of backfill to all of the lower levels. Once the hypothesis is rejected then the level is no longer a part of the dose-finding algorithm, or at least no longer a part of the backfill component of the dose-finding algorithm.

We can pursue the same strategy to the backfill patients at levels 2, 3 and above. Starting with the lowest, these levels can be eliminated from the backfill allocation set - one by one - as we obtain enough evidence to reject a hypothesis that the current lowest level enjoys an efficacy rate comparable to the higher levels. Here, we chose to work with a Bayesian Beta-Binomial model in order to base our decision on the posterior distribution of the difference in tumour response rates between dose level 1 and the remaining dose levels. This is not the only way to achieve our objective. Other models could be used or, indeed, we could make a simple comparison based on the empirical rates. Different approaches to this comparison could be a subject worthy of further study. Our approach appeared to work quite well here. For those simulated scenarios where the dose-efficacy curve was increasing over at least the first two dose levels, this approach discontinued either the first or the first two dose levels about 50% of the time. Another possible, potentially more realistic, model might postulate a steady dose-efficacy function and one that includes a change point in the slope at some point on the dose-efficacy curve. Doses below the change point would be considered unacceptable and removed from the backfill set. A very extensive literature on change point detection, in particular when it relates to rates, is available and could provide some initial input into a further investigation of the sequential identification of the backfill set. There is some room then for further study of the best way in which to structure the backfill component of the study.

A final ingredient to the backfill component of the study remains and that is the distribution, $G(d_i)$ to the backfill doses. Here, for simplicity, we have taken this to be discrete uniform over the backfill set. However there is no added difficulty in working with some other choice for $G(d_i)$, possibly one that puts more weight on those doses closer to the MTD than those further removed from it. More research on $G(\cdot)$ is needed. For example, a tendency to keep most of the backfill set close to the MTD, while reducing the risk of allocation to ineffective doses will also have the result of reducing the amount of information we need to discard ineffective doses. So, the impact of an apparently conservative approach to backfill is not immediately clear. Greater complexity would follow if we allow $G(d_i)$ to also depend on the observed rates of efficacy. It would be quite a challenge to find a way to decide among the various options which ones are the most promising. In some instances, including the

lowest dose level in the backfill set may not be appealing to investigators and sponsor. Indeed there might be prior information (from pharmacokinetic data for example) that the efficacy rate at the starting dose level may not be sufficiently high to justify backfilling this level. In our framework, this would be achieved by allocating a fixed weight of zero to dose level 1, i.e. $G(d_1) = 0$.

The main goal of backfill, as far as our final recommendation is concerned, is to be able to have enough information to decide whether or not a dose lower than the estimated MTD may be a more suitable dose to take forward as our RP2D. Again, there are many potential approaches to providing an answer to that question. We have taken a simple approach in which we contrast hypotheses H_1 and H_2 . The approach is promising as illustrated by scenarios A and B where the plateau was identified some 70% of the time. Under H_1 , we assumed the dose-efficacy function could be approximated by a logistic model. This could be given more study. Under H_2 , we used a similar function for the first part of the curve, followed by a change point at which no further increase in the probability of efficacy is apparent. As for all modelling, we do not suppose that H_1 or H_2 describe precisely the mechanisms we are studying but that they provide working approximations to them. Our interest is focused on the difference between H_1 and H_2 rather than the exactitude of either so that, if we feel H_2 provides a better description of the observations then we use this information to locate the approximate beginning of the plateau. We then argued that this point provides a more effective RP2D than the MTD itself. Our results showed this quite convincingly in some cases.

Clearly we cannot hope to come out on top in all situations and the most challenging situation here would be one in which the dose-efficacy curve is not particularly steep. Such a case is shown in Scenario C and does not involve any plateau. The chances here of choosing incorrectly an RP2D lower than the MTD are higher since a simple change point model with a plateau can prove itself to better reflect the observations than the more accurate monotonic logistic model. Examples in the dose-finding literature where an incorrect model provides a better fit than the correct model are quite common and stem from the paucity of information consequent upon small sample sizes. The error here is not great since a shallow dose efficacy curve tells us that the lower recommendation is not too far below the potentially more effective one at the MTD. All of these issues need to be brought out at the design stage in order to tie down a particular design with features that, overall and on average, result in the kind of operating characteristics that the study investigators deem suitable.

It was necessary to make a number of choices in order to find a design that can be implemented in practice. Some of these may seem quite arbitrary such as the class of dose-finding models we chose to work with. Other choices relate to design calibration and, under this heading, we might wish to investigate various possibilities for the selection of a suitable $G(d_i)$. A final set of choices were made in order to gain

an angle on statistical properties. Under this heading there is our choice to employ a subset of patients to identify the MTD and the whole cohort to identify which doses can be considered acceptable for the backfill set. A significant body of theoretical study is available to us to confirm that this approach avoids any potential bias in estimating the MTD [14]. We could potentially make fuller use of all of the data and there may be ways in which information gleaned from the backfill set could be employed, at least in part, to help improve our sequential location and final estimation of the MTD. This would be worthy of further study.

It is important to note that backfilling may delay the completion of the study. However, backfilling is done in practice nowadays because investigators believe that this may be a price worth paying in order to collect additional information during the study, including assessing whether there is a plateau on the dose-efficacy curve. Nonetheless, as illustrated in Figure 2, backfilling takes place simultaneously as the dose-finding cohorts. This saves time compared to backfilling after completing the dose-escalation phase.

In this paper we have provided a comprehensive framework to backfill patients in a controlled way in oncology dose-finding trials that takes into account the ethical constraints that we should not backfill to dose levels as soon as there is evidence of insufficient efficacy.

4.1. Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

4.2. Declaration of conflicting interests

The Authors declare that there is no conflict of interest.

References

- [1] Agrawal, S., Feng, Y., Roy, A., Kollia, G., Lestini, B., 2016. Nivolumab dose selection: challenges, opportunities, and lessons learned for cancer immunotherapy. *Journal for immunotherapy of cancer* 4, 72.
- [2] Conaway, M.R., Petroni, G.R., 2019. The impact of early-phase trial design in the drug development process. *Clinical Cancer Research* 25, 819–827.
- [3] Corbaux, P., El-Madani, M., Tod, M., Péron, J., Maillet, D., Lopez, J., Freyer, G., You, B., 2019. Clinical efficacy of the optimal biological dose in early-phase trials of anti-cancer targeted therapies. *European Journal of Cancer* 120, 40–46.
- [4] Gelman, A., Hwang, J., Vehtari, A., 2014. Understanding predictive information criteria for bayesian models. *Statistics and computing* 24, 997–1016.
- [5] Iasonos, A., O'Quigley, J., 2013. Design considerations for dose-expansion cohorts in phase I trials. *Journal of Clinical Oncology* 31, 4014.
- [6] Iasonos, A., O'quigley, J., 2015. Early phase clinical trials—are dose expansion cohorts needed? *Nature reviews Clinical oncology* 12, 626–628.
- [7] Iasonos, A., O'Quigley, J., 2016. Integrating the escalation and dose expansion studies into a unified phase I clinical trial. *Contemporary clinical trials* 50, 124–134.
- [8] Iasonos, A., O'Quigley, J., 2017. Sequential monitoring of phase I dose expansion cohorts. *Statistics in medicine* 36, 204–214.
- [9] Iasonos, A., O'Quigley, J., 2021. Randomised phase I clinical trials in oncology. *British Journal of Cancer* .
- [10] Iasonos, A., O'Quigley, J., 2016. Dose expansion cohorts in phase I trials. *Statistics in biopharmaceutical research* 8, 161–170.
- [11] Jung, K.H., LoRusso, P., Burris, H., Gordon, M., Bang, Y.J., Hellmann, M.D., Cervantes, A., de Olza, M.O., Marabelle, A., Hodi, F.S., et al., 2019. Phase I study of the indoleamine 2, 3-dioxygenase 1 (ido1) inhibitor navoximod (gdc-0919) administered with pd-11 inhibitor (atezolizumab) in advanced solid tumors. *Clinical Cancer Research* 25, 3220–3228.
- [12] Manasanch, E.E., Shah, J.J., Lee, H.C., Weber, D.M., Thomas, S.K., Amini, B., Olsem, J., Crumpton, B., Morphe, A., Berkova, Z., et al., 2020. Phase I/Ib study of carfilzomib and panobinostat with or without dexamethasone in patients with relapsed/refractory multiple myeloma. *haematologica* 105, e242.
- [13] O'Quigley, J., Pepe, M., Fisher, L., 1990. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* , 33–48.
- [14] O'Quigley, J., 2006. Theoretical study of the continual reassessment method. *Journal of Statistical Planning and Inference* 136, 1765–1780.
- [15] Patil, V.M., Noronha, V., Joshi, A., Abhyankar, A., Menon, N., Banavali, S., Gupta, S., Prabhaskar, K., et al., 2019. Low doses in immunotherapy: Are they effective? *Cancer Research, Statistics, and Treatment* 2, 54.
- [16] Sachs, J.R., Mayawala, K., Gadamsetty, S., Kang, S.P., de Alwis, D.P., 2016. Optimal dosing for targeted therapies in oncology: drug development cases leading by example. *Clinical Cancer Research* 22, 1318–1324.
- [17] Sweeting, M., Mander, A., Sabin, T., 2013. Bcrn: Bayesian continual reassessment method designs for phase I dose-finding trials. *Journal of Statistical Software* 54.
- [18] Tao, J.J., Cangemi, N.A., Makker, V., Cadoo, K.A., Liu, J.F., Rasco, D.W., Navarro, W.H., Haqq, C.M., Hyman, D.M., 2019. First-in-human phase I study of the activin a inhibitor, stm 434, in patients with granulosa cell ovarian cancer and other advanced solid tumors. *Clinical Cancer Research* 25, 5458–5465.
- [19] Thall, P.F., Cook, J.D., 2004. Dose-finding based on efficacy–toxicity trade-offs. *Biometrics* 60, 684–693.
- [20] Vehtari, A., Gelman, A., Gabry, J., 2017. Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and computing* 27, 1413–1432.

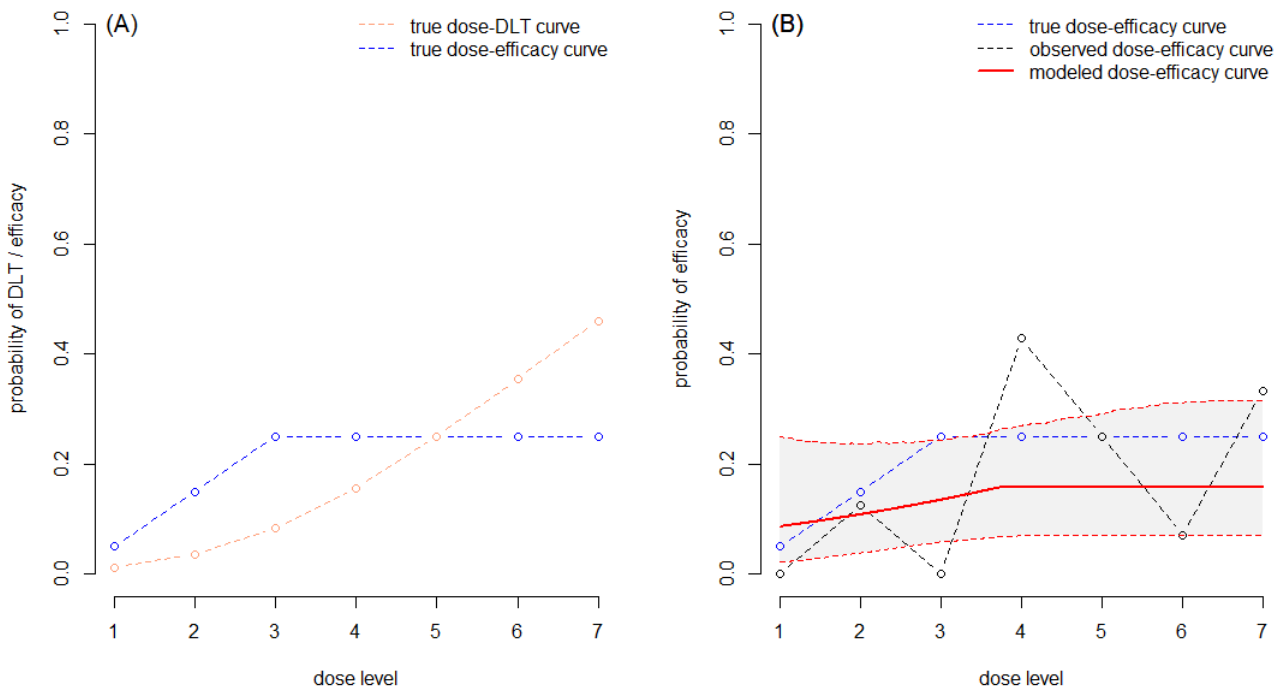


Figure 1: (A) True dose-DLT and dose-efficacy curves from which the hypothetical trial data were simulated, (B) true, observed and modeled dose-efficacy curves in the hypothetical trial after 10 cohorts and 57 patients in total. The shaded grey area corresponds to the limits of the 2.5%-97.5% quantiles of the posterior distribution of the modeled dose-efficacy curve

Controlled backfill

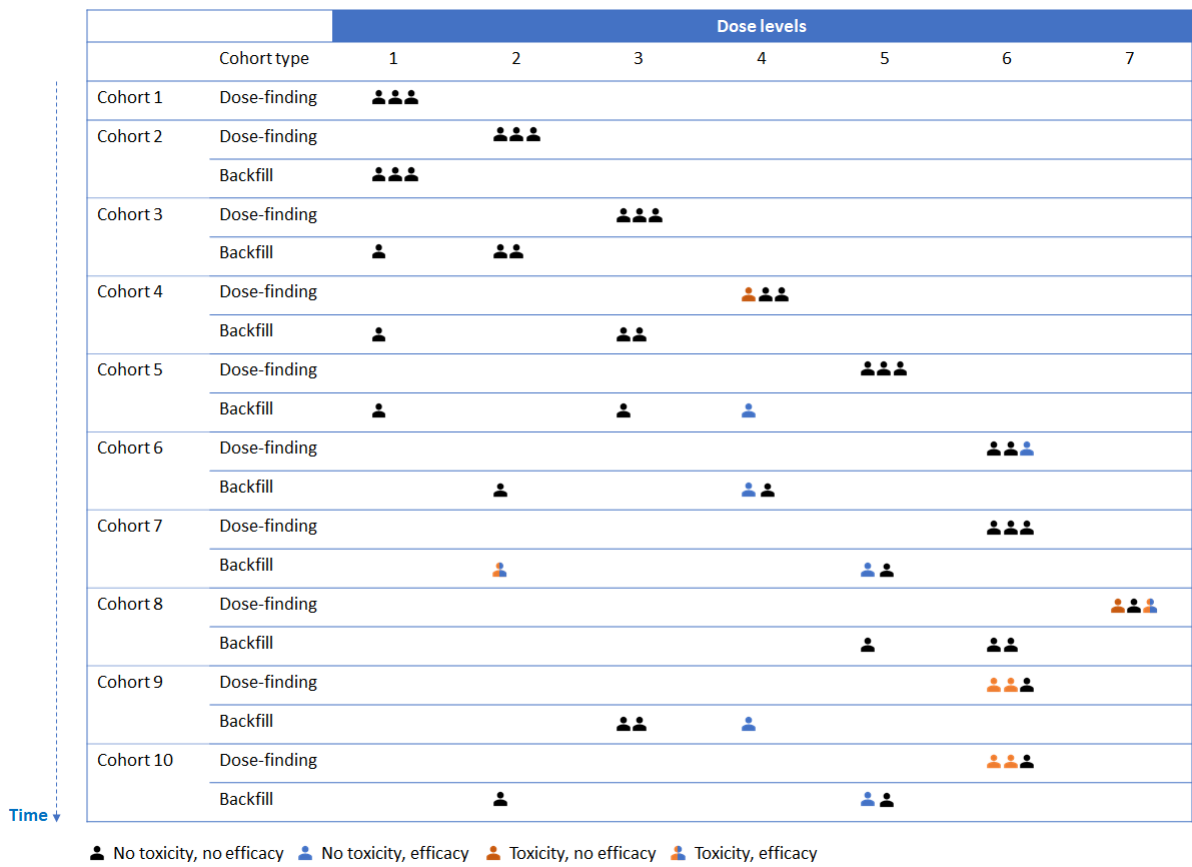


Figure 2: Evolution of the hypothetical dose-finding trial

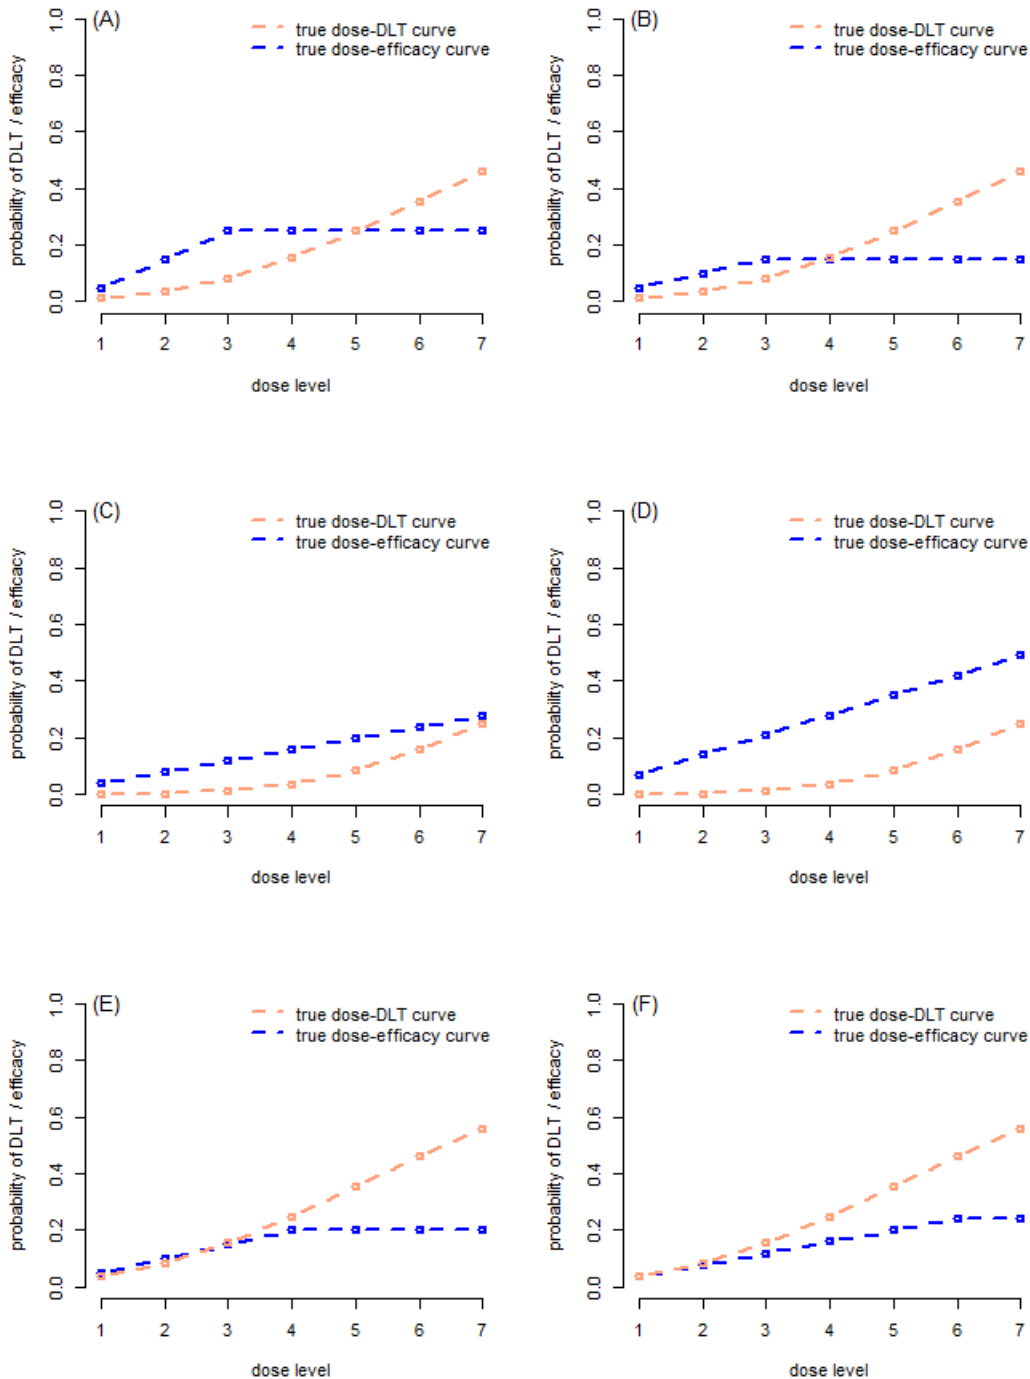


Figure 3: Simulation scenarios: (A) dose-efficacy curve displays a plateau at 25% tumour response probability from dose level 3 and MTD is dose level 5 (TTL is 25% for all scenarios), (B) dose-efficacy curve displays a plateau at 15% tumour response probability from dose level 3 and MTD is dose level 5, (C) monotonic dose-efficacy curve reaching 28% at dose level 7, which is also the MTD, (D) monotonic dose-efficacy curve reaching 49% at dose level 7, which is also the MTD, (E) dose-efficacy curve displays a plateau at dose level 4 at 20% tumour response probability and MTD is dose level 4, (F) dose-efficacy curve displays a plateau at dose level 6 at 24% tumour response probability and MTD is dose level 4

Table 1
Dose-efficacy curve, dose-DLT curve, skeleton, recommendations and in-trial allocations of the six simulated scenarios

	Dose level						
	1	2	3	4	5	6	7
Scenario A							
dose-efficacy curve	5%	15%	25%	25%	25%	25%	25%
dose-DLT curve	1%	4%	8%	16%	25%	35%	46%
skeleton	1%	4%	8%	16%	25%	35%	46%
% of recommendation for dose	0%	0%	29%	46%	16%	8%	1%
% of patients receiving dose	16%	18%	20%	19%	17%	9%	2%
Scenario B							
dose-efficacy curve	5%	10%	15%	15%	15%	15%	15%
dose-DLT curve	1%	4%	8%	16%	25%	35%	46%
skeleton	1%	4%	8%	16%	25%	35%	46%
% of recommendation for dose	0%	0%	32%	46%	15%	6%	1%
% of patients receiving dose	18%	19%	18%	18%	17%	8%	2%
Scenario C							
dose-efficacy curve	4%	8%	12%	16%	20%	24%	28%
dose-DLT curve	0%	0%	1%	4%	8%	16%	25%
skeleton	5%	10%	15%	20%	25%	30%	35%
% of recommendation for dose	0%	0%	0%	30%	10%	9%	51%
% of patients receiving dose	17%	16%	14%	13%	11%	10%	19%
Scenario D							
dose-efficacy curve	7%	14%	21%	28%	35%	42%	49%
dose-DLT curve	0%	0%	1%	4%	8%	16%	25%
skeleton	5%	10%	15%	20%	25%	30%	35%
% of recommendation for dose	0%	0%	0%	14%	13%	13%	60%
% of patients receiving dose	15%	15%	14%	13%	12%	11%	19%
Scenario E							
dose-efficacy curve	5%	10%	15%	20%	20%	20%	20%
dose-DLT curve	4%	8%	16%	25%	35%	46%	56%
skeleton	3%	10%	17%	25%	32%	40%	48%
% of recommendation for dose	0%	5%	45%	36%	11%	2%	0%
% of patients receiving dose	21%	22%	24%	19%	10%	3%	1%
Scenario F							
dose-efficacy curve	4%	8%	12%	16%	20%	24%	24%
dose-DLT curve	4%	8%	16%	25%	35%	46%	56%
skeleton	3%	10%	17%	25%	32%	40%	48%
% of recommendation for dose	0%	6%	45%	34%	13%	3%	0%
% of patients receiving dose	21%	22%	23%	19%	11%	3%	1%

Table 2

Plateau identification in the six scenarios and recommendation of RP2D at a dose level lower than MTD

[% of simulations]	plateau identification	RP2D < MTD due to existence of plateau
Scenario A	70%	66%
Scenario B	72%	69%
Scenario C	41%	41%
Scenario D	31%	30%
Scenario E	56%	42%
Scenario F	56%	39%

Table 3
Backfill dose discontinuation percentages in the six scenarios

[discontinuation %]	Dose level(s) discontinued					
	none	1	1+2	1+2+3	1+2+3+4	1+2+3+4+5
Scenario A	14%	37%	39%	9%	1%	0%
Scenario B	39%	36%	20%	3%	1%	0%
Scenario C	27%	33%	25%	11%	3%	0%
Scenario D	11%	26%	30%	20%	11%	2%
Scenario E	33%	31%	21%	12%	2%	0%
Scenario F	36%	31%	19%	10%	3%	0%