
Backfilling Cohorts in Phase I Dose-Escalation Studies

Journal Title
XX(X):2–18
© The Author(s) 2016
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

Helen Barnett^{1,2}, Oliver Boix³, Dimitris Kontos⁴ and
Thomas Jaki^{1,5}

Abstract

The use of ‘backfilling’, assigning additional patients to doses deemed safe, in phase I dose-escalation studies has been used in practice to collect additional information on the safety profile, pharmacokinetics and activity of a drug. These additional patients help ensure that the Maximum Tolerated Dose (MTD) is reliably estimated and give additional information in order to determine the recommended phase II dose. In this paper, we study the effect of employing backfilling in a phase I trial on the estimation of the MTD and the duration of the study. We consider the situation where only one cycle of follow-up is used for escalation as well as the case where there may be delayed onset toxicities. We find that, over a range of scenarios, there is an increase in the proportion of correct selections and a notable reduction in the trial duration at the cost of more patients required in the study.

Keywords

Dose-Finding; Dose-Escalation; Backfilling; Phase I Trials; Model-Based; Late-onset Toxicity.

Introduction

In Phase I dose-finding studies, the main objective is often to find the Maximum Tolerated Dose (MTD) or the recommended Phase II dose, the dose recommended for further testing in Phase II. The MTD is defined as highest dose that has an acceptable level of toxicity¹, most often corresponding to a certain probability of occurrence of a Dose Limiting Toxicity event (DLT). In oncology, a DLT is frequently defined as a grade 3 or higher toxicity by the grading scale of the National Cancer Institute². In the following, we use the terms DLT and toxicity interchangeably.

In a Phase I dose-escalation study, a set of doses is investigated, patients are recruited in cohorts and an escalation procedure is used to carefully escalate from low doses that are expected to be very safe to dose levels that have an acceptable level of toxicity and at the same time induce some desirable activity in a patient. The escalation procedure can be rule-based¹, model-based³ or model-assisted⁴ and cohorts are most commonly small⁵, with a size of 3 often used in such trials.

However, with such small sample sizes, comes greater uncertainty in the estimate of interest. It is, for example, desirable to establish the MTD quickly and accurately but with no more patients than necessary. As with any clinical trial, a balance must be taken between the accuracy, the duration of the study and the trial size. Typically, the

¹MRC Biostatistics Unit, University of Cambridge

²Lancaster University

³Bayer AG

⁴ClinBAY

⁵University of Regensburg

Corresponding author:

Helen Barnett

Email: helenyvetteb@gmail.com

larger the trial, the higher the accuracy, but also the longer the duration of the study.

An approach that has gained popularity in the recent years^{6,7}, is the use of ‘backfilling’ of cohorts on lower doses⁸. The principle is that, once a dose is deemed safe enough to escalate to a higher dose-level, additional patients may be allocated to lower doses to increase the understanding of the safety, tolerability and activity of these doses. The decision to backfill a dose may be taken solely on the criteria that the dose is deemed ‘safe’, or it may require the additional condition that an activity signal must be seen.

Whilst it is clear that backfilling will result in a, potentially large, increase in patient numbers, additional insight around the safety, tolerability and potential activity of the treatment is gained. Moreover, there is the potential that the trial duration could be substantially reduced due to the improved understanding of the dose-toxicity relationship. The exact nature of this relationship between trial duration and sample size is, however, unclear - in particular in the setting where late-onset toxicities are of concern.

In this work we investigate the impact of backfilling on relationship between probability of correctly selecting the MTD, the duration and the sample size in phase I dose-escalation trials, with a motivation of reducing the total duration.

Motivating Trial Examples

The first in-human Phase I study⁹ investigated the safety and activity of the activin A inhibitor, STM 434, in advanced solid tumours. The study initially considered five doses levels, 0.25, 0.5, 1, 2, and 4 mg/kg administered

every 4 weeks. The treatment scheduled subsequently changed to bi-weekly after the half-life was estimated to be lower than anticipated and an additional dose of 8 mg/kg was added later due to lower than anticipated predicted exposures. A 3+3 design¹ was used to guide dose-escalation. Dose-limiting toxicities (DLT) were defined as any grade ≥ 3 non-haematologic toxicity, any grade ≥ 4 haematologic toxicity lasting 7 days, febrile neutropenia, or grade 3 thrombocytopenia with active bleeding and the DLT assessment period was 28 days. A minimum of 3 evaluable patients were required prior to dose-escalation and "Backfill" slots were permitted at doses that had been declared safe.

A total of 32 patients participated in the trial of which three experienced a DLT and of 28 patients that were evaluable according to RECIST¹⁰, 16 achieved stable disease while the remaining had progressive disease (Table 1). Since two of the DLTs occurred in the highest dose, the MTD was declared to be 4 mg/kg every 2 weeks but no dose-expansion was undertaken following the safety review committees recommendation on the basis of the overall safety profile observed.

[Table 1 about here.]

From the results of the study we can see that backfilling slots were indeed used as several of the cohorts have more than three (or six) patients required for the 3+3 design.

Note, that the 3+3 design does, however, not utilize these additional data when recommending subsequent doses. In the remainder of this work we will therefore explore the impact of backfilling on the operating characteristics of a trial that utilizes a model-based escalation approach. This exploration is motivated by a recent study of Thorium-227 in combination with an antibody (NCT03507452)¹¹, a potential treatment for Advanced Recurrent Epithelioid

Mesothelioma, Serous Ovarian Cancer and Metastatic or Locally Advanced Pancreatic Ductal Adenocarcinoma. As this study utilized doses of Thorium-227 of 1.5 MBq in steps of 1.0 or 1.5 MBq, with antibody doses of 10 mg, we will consider the same dose-levels in our subsequent evaluations. Moreover, due to the radioactive nature of Thorium-227, late-onset toxicities – toxicities that occur after the first treatment cycle – are of potential concern. We therefore also explore the impact of backfilling under a traditional CRM-type model as well as a time-to-event model.

Methods

In this section we outline the methods used in this work. Firstly, we illustrate and discuss the concept of backfilling in more detail. We then describe the statistical models used for dose-escalation, and finally outline the rules used in the design implementation.

Single Simulation Examples

To illustrate the concept of backfilling, and the difference between trials that use backfilling and those that do not, we present an example of a single trial simulation. Following the motivating study (NCT03507452) the following six doses are investigated: 1.5MBq , 2.5MBq , 3.5MBq , 4.5MBq , 6.0MBq , 7.0MBq. **Figure 1 illustrates a trial conducted without backfilling (left panel) and with full backfilling (right panel). Every other aspect of the trials was comparable.**

Cohorts of size three enter the trial, starting at the lowest dose. If no backfilling is implemented, then cohorts are assigned to doses according to the escalation procedure, until some stopping rule is triggered. If full backfilling is implemented, when a dose is considered safe enough that

the escalation continues above it, the next two cohorts are 'backfilling cohorts' and are recruited to that same dose. Then the next cohort is assigned to the escalated dose. Of course one may also choose to backfill more cautiously, but in this illustration we use two cohorts as standard.. Patients are followed up for three cycles of treatment, with new cohorts assigned after each cycle, with a cycle lasting 6 weeks.

In this case, the trial without backfilling determined the MTD to be the 4.5MBq dose while the trial using backfilling recommended 3.5MBq. Figure 1 shows that the escalation in both settings is the same - they both reach the 4.5MBq dose, but in the trial without backfilling two further cohorts are allocated to this dose before being recommended as the MTD. When backfilled cohorts were used, more information was available at the lower doses allowing the precision stopping rule to be reached quickly and hence the lower dose is recommended. In addition to the difference in recommended dose one can also observe that the duration of the study is notably shorter when backfilling is used.

[Figure 1 about here.]

BLRM and TITE-BLRM

In order to investigate the impact of backfilling on the estimation of the MTD and the duration of the trial, we consider two settings. The first setting assumes that patients are followed up for one cycle of treatment only (i.e. DLT period 1 cycle), and the next cohort of patients is assigned once the previous cohort's follow-up period has been fully observed. The second setting assumes that there may be late onset toxicities. Patients are therefore followed up for three cycles of treatment. A new cohort of patients is admitted every cycle, so that only partial information is available for the previous two cohorts, as

their full follow-up period has not yet been observed. In the first setting we use a Bayesian Logistic regression model (BLRM)¹², and in the second setting a Time-To-Event version of the BLRM (TITE-CRM)¹³.

In each setting, a set of J doses labelled d_j for $j = 1, \dots, J$ are investigated, with patients labelled $i = 1, \dots, n$.

The BLRM is conducted in the following way. Starting at the lowest dose d_1 , cohorts of patients enter the trial. After each cohort has been fully observed, the dose assignment of the next cohort is decided. A two-parameter logistic model is used to describe the dose-toxicity relationship:

$$F(d, \boldsymbol{\beta}) = \frac{\exp(\beta_0 + \beta_1 d)}{1 + \exp(\beta_0 + \beta_1 d)},$$

where $F(d, \boldsymbol{\beta})$ is the probability of DLT at dose d , and $\boldsymbol{\beta} = (\beta_0, \beta_1)$ is parameter vector with prior:

$$\begin{pmatrix} \beta_0 \\ \log(\beta_1) \end{pmatrix} \sim N_2 \left(\begin{pmatrix} c_1 \\ c_2 \end{pmatrix}, \begin{pmatrix} v_1 & 0 \\ 0 & v_2 \end{pmatrix} \right).$$

The posterior distribution of $\boldsymbol{\beta}$ is updated after each cycle using the likelihood

$$\mathcal{L}(\boldsymbol{\beta}) = \prod_{i=1}^n F(d, \boldsymbol{\beta})^{y_i} \{1 - F(d, \boldsymbol{\beta})\}^{1-y_i},$$

where y_i is an indicator taking the value 1 if patient i had observed a DLT and 0 otherwise. The updated posterior for $\boldsymbol{\beta}$ is then used to estimate the probability of DLT at each dose. The dose assignment of the next cohort is then dose d_j that minimizes $|F(d_j, \hat{\boldsymbol{\beta}}) - \tau_1|$, where τ_1 is the target DLT rate for one cycle of follow-up, subject to certain rules - see the next section. The final dose recommendation is then the dose d_j that minimizes $|F(d_j, \boldsymbol{\beta}) - \tau_1|$ once a stopping rule has been triggered.

In the second setting, the TITE-CRM¹³ is used. Here, patients are followed up for three cycles of treatment. However, if it was required to wait until the entire follow-up for the previous cohort had been completely observed to assign the dose for the next patient, the trial length would be very undesirably long. Therefore, each new cohort is assigned their dose once the previous cohort has been observed for one cycle of treatment. There is therefore only partial information available for the previous two cohorts. The TITE-CRM takes this into account by weighting the observations in the following way.

The dose **toxicity** model $F(d, \boldsymbol{\beta})$ is identical to in the non-time-to-event version, and is weighted to form $G(d, w, \boldsymbol{\beta})$:

$$G(d, w, \boldsymbol{\beta}) = wF(d, \boldsymbol{\beta}),$$

where the weights $0 \leq w \leq 1$ are a function of time-to-event of a patient **toxicity**. The posterior distribution of $\boldsymbol{\beta}$ is updated after each cycle using likelihood

$$\mathcal{L}(\boldsymbol{\beta}) = \prod_{i=1}^n G(d_{[i]}, w_{i,n}, \boldsymbol{\beta})^{y_{i,n}} \{1 - G(d_{[i]}, w_{i,n}, \boldsymbol{\beta})\}^{1-y_{i,n}},$$

where n is the number of patients that have been treated so far, $d_{[i]}$ is the dose assigned to patient i and $y_{i,n}$ is an indicator which takes the value 1 if patient i has observed a DLT after the n patients treated so far have been observed for at least one cycle, and 0 otherwise.

The updated posterior for $\boldsymbol{\beta}$ is used to estimate the probability of DLT at each dose, as in the one cycle setting. The dose assignment of the next cohort is then the dose d_j that minimizes $|F(d_j, \boldsymbol{\beta}) - \tau_3|$, where τ_3 is the target DLT rate for three cycles of follow-up, again subject to certain rules - see the next section.

In this implementation, we use the simple specification of weights suggested by Cheung and Chappell¹³:

$w_{i,n} = u_{i,n}/S$, where $u_{i,n}$ is the current number of cycles for which patient i has been observed and S is the total number of cycles in the follow-up period. If patient i observes a DLT, then $w_{i,n} = 1$. The final dose recommendation is the dose level that minimises $|F(d_j, \hat{\beta}) - \tau_3|$ once a stopping rule has been implemented and the follow-up for all enrolled patients has been completed.

Due to the added complexity of the method, the TITE-CRM uses a start-up period such that the dose assignment is escalated one level at a time until a DLT is observed. Once a DLT is observed, then the TITE-CRM model is used. Although in the original description of the methodology, this start-up period requires each patient to be followed up for their entire follow-up time before the next patient's dose is assigned, in our implementation, only one cycle is required for follow-up before the next is assigned, which is in line with the rest of the trial.

Simulations Set-up

To evaluate the performance in the setting of a Phase I dose-escalation trial, we use enforcement and stopping rules that could be used in such a trial. We define p_{s,d_j} as the $P(DLT)$ in cycles up to and including cycle s for any given dose d_j . These rules are described in the Online Supplementary Materials.

In order to investigate the effect of backfilling on the operating characteristics of the dose-finding designs, we conduct a simulation study. The same set of six doses as used in the motivating example are used: 1.5MBq , 2.5MBq , 3.5MBq , 4.5MBq , 6.0MBq , 7.0MBq. Seventeen scenarios are considered, to cover a wide range of potential dose responses. Table 2 gives the probability of a DLT in the first six week cycle for each of the considered

scenarios, with the MTD associated with the target, $\tau_1 = 0.3$, highlighted in boldface. Scenarios 1 – 6 represent cases where each level in turn is the MTD, with higher and lower doses equidistant in terms of probability of DLT. Scenarios 7 and 8 are non-linear around the MTD. In scenarios 9 – 12, no dose is exactly on target, with all doses unsafe in scenario 9. Scenarios 13 – 17 are a set of varying scenarios, typically used to test performance of a dose-finding algorithm.

[Table 2 about here.]

For the setting where the follow-up period is three cycles, the conditional probability of DLT in subsequent cycles is multiplied by a factor of $1/3$. So that

$$p_{2,d_j} = p_{1,d_j} + (1 - p_{1,d_j})\frac{p_{1,d_j}}{3},$$

and

$$p_{3,d_j} = p_{1,d_j} + (1 - p_{1,d_j})\frac{p_{1,d_j}}{3} + (1 - p_{1,d_j})(1 - \frac{p_{1,d_j}}{3})\frac{p_{1,d_j}}{9}.$$

For example if $p_{1,d_j} = 0.3$, then $p_{2,d_j} = 0.37$ and $p_{3,d_j} = 0.391$, hence we use $\tau_3 = 0.391$ as the target toxicity for three cycles of follow-up. **To generate patient toxicity, a Uniform(0,1) random variable, U , is generated for each patient, and a DLT is observed in cycle 1 on dose d_j if $U \leq p_{1,d_j}$, in cycle 2 if $p_{1,d_j} < U \leq p_{2,d_j}$, in cycle 3 if $p_{2,d_j} < U \leq p_{3,d_j}$ and no DLT is observed if $U > p_{3,d_j}$.**

The maximum sample size is chosen to be $n_{\max} = 54$ which is relatively large for a Phase I trial, but has been chosen to allow escalation to dose d_6 with backfilling implemented.

We conduct 5,000 simulations for each scenario, and compare the performance to the non-parametric

benchmark¹⁴. This benchmark gives an indication of the ‘difficulty’ of a scenario, so that we can quantify any differences between the approaches accordingly. Note that this uses the maximum sample size in every simulation, and is not subject to any stopping rules.

In these implementations, as well as comparing backfilling all doses considered safe (Fully Backfilled), and not (Not Backfilled), we also consider the setting where a dose level is only backfilled once an activity signal has been observed at that dose or any dose below (Partially Backfilled). For example, if an activity signal is first seen in the third dose level and not the first or second, then the decision to escalate to the fourth dose level would mean the additional backfilling cohorts are assigned to dose three but not doses one and two. **Since we employ backfilling once an activity signal has been seen, we use the probability of at least one complete response in a cohort. In this partially backfilling setting, the underlying probability of observing an activity signal is 0.00, 0.15, 0.30, 0.45, 0.60, 0.75 at each of the six doses respectively, corresponding to an individual’s activity probability of (0, 0.05, 0.11, 0.18, 0.26, 0.37) at each of the six doses.**

We consider two options for the values of the hyper-parameters of the prior distribution; non-calibrated and calibrated. In the first option, the hyper-parameters for the prior used are chosen so that the prior is relatively vague, with a mean effective prior sample size of 1.3 patients per dose level. These choices are in line with those used by Neuenschwander et al.¹². The second option is to calibrate the values of these hyper-parameters over a small range of scenarios, to choose the values yielding the best performance across a wide range of settings (see for example Mozgunov et al.¹⁵). Further details, including the values of the hyper-parameters are available in the online supplementary materials.

Results

The most often considered metrics of performance in dose-finding trials are the proportion of correct selections and the proportion of acceptable selections. A correct selection is defined as selecting the MTD as defined in Table 2, or as a safety stopping rule being correctly triggered. An acceptable selection is defined as a dose whose true probability of toxicity during the first cycle is between 0.18 and 0.33, based on the overview of Phase I trials by¹⁶.

Here we discuss the results for the non-calibrated prior, with the calibrated prior discussed in the online supplementary materials.

Non-Calibrated Prior

[Figure 2 about here.]

[Figure 3 about here.]

Figures 2 and 3 show the **proportion of correct and acceptable selections** in the setting with one cycle and three cycles respectively. **The Monte Carlo simulation errors are at most $\pm 1.4\%$** . From Figure 2, it can be seen that, when one cycle of follow-up is considered, employing backfilling increases the **proportion of correct and acceptable selections** in almost all scenarios, as expected. The only scenario which did not see an increase in **the proportion of correct selections** is scenario 7, where there is a marginal decrease in **the proportion of correct selections**, but a noticeable increase in **the proportion of acceptable selections**. This is a particularly challenging scenario, where the dose below the MTD has true probability of DLT only 5% below target. The largest increase is seen in scenario 13, where full backfilling gives a 9% increase in correct selections. In nearly all scenarios, partial backfilling gives a performance between no backfilling and fully backfilling as expected. Interestingly in scenario 17,

the performance of partial backfilling is worse than both no and full backfilling. With such an ‘easy’ scenario, the **proportion of correct selections** is already very high, and the decrease is due to a shift to recommending the unsafe fourth dose. The average increase in correct selections from no backfilling to fully backfilling is 4%, and is 5% for acceptable selections. The approach to partially backfilling we have taken means the probability of backfilling any given level dose increases with dose. Therefore when the lower doses are more toxic, fewer doses will be backfilled and partial backfilling gives more similar results to no backfilling. Although in all cases it is noticeable that the performance is well below that of the benchmark, it is worth noting that the benchmark has the advantage of a much larger sample size on average.

To compare the measures of trial sizes of the two approaches, Figure 4 shows the relationship between the mean total sample size and the mean trial duration for each scenario **and most importantly how this relationship changes for differing levels of backfilling**. The diamonds represent the setting with one cycle of follow-up, with blue indicating that full backfilling was used, purple indicating partial backfilling and black indicating no backfilling. It is clear to see that the use of full backfilling substantially increases the total sample size in scenarios where the MTD is at least the second dose level. **This increase in sample size owed to the use of backfilling decreases the average trial duration notably, seen by the upward left drift of the blue diamonds compared to their black counterparts.** On average across scenarios, full backfilling increases sample size by 12 patients on average but reduces the trial duration by 6 weeks (**one treatment cycle**). It appears that on average each additional patient reduces the duration by half a week. For partial backfilling, although the overall increase in sample size is smaller, as is the decrease in trial duration, the trade-off is still on average each additional

patient reduces the duration by half a week. **Although such a trade-off may seem small, it also brings with it the increase in precision of MTD estimation.**

The other important metrics used in dose-finding trials concern the safety of the patients within the trial. Figures in the online supplementary materials show the distribution of mean percentage of DLTs and the mean number of patients assigned to overly toxic doses. Although backfilling increases the number of DLTs observed in every scenario, the percentage of patients observing a DLT decreases slightly when full backfilling is used, and is largely unchanged for partial backfilling. For the most part, the increase in numbers of patients exposed to overly toxic doses are in scenarios where all doses are unsafe. Hence one mistaken escalation increases this exposure by three cohorts when backfilling instead of one when not. The benefit in reducing the number of patients exposed to overly toxic doses comes in scenarios where the MTD is in the middle of the dose range, and the backfilling expansion provides more information on the lower doses, resulting in more cautious escalation.

[Figure 4 about here.]

Interestingly, the patterns observed in the setting with one cycle of follow-up are not all duplicated in the setting with three cycles of follow-up. It is not the case here that the **proportion of correct and acceptable selections** increase across all scenarios when backfilling is employed. In fact in some scenarios, there is a noticeable decrease in **proportion of correct selections** when backfilling is employed. For example, in scenarios 5 – 7, where the **proportion of correct selections** is low in both settings, it is lower when backfilling is used. This reflects further the point raised earlier that the backfilling leads to a more cautious escalation. In most cases, the **proportion**

of acceptable selections is higher when backfilling is implemented.

In terms of mean sample size and trial duration, the circles on Figure 4 display this relationship. This relationship is similar to that observed with one cycle. For every additional patient, the trial duration decreases by 0.3 weeks, for both full and partial backfilling. Again, the magnitude of this overall increase or decrease is larger for fully backfilling. Likewise, the comparisons of the safety aspects of the trials are similar. With three cycles of observation, fully backfilling increases the total number of DLTs in all scenarios, decreases the percentage of DLTs overall, and decreases the number of patients exposed to overly toxic doses in some scenarios, and increases in others. On average across scenarios, there is a slight decrease in the number of patients exposed to overly toxic doses.

Conclusions

In this work, we have investigated the effect of backfilling on the operating characteristics of dose-escalation studies. The main reasons to utilize backfilling include to gain better understanding of the safety, tolerability and activity of the treatment under investigation and to aid determining the recommended Phase II dose. In this work we have focused on the implications of backfilling on the estimation of the MTD and the duration of the study. We found that backfilling increases the chance of identifying the MTD while reducing the duration of the study, on average by 14%. This comes at the cost of an increased number of patients required in studies that use backfilling. The impact of backfilling on the accuracy is larger in the setting with a DLT assessment period of one cycle than when three cycles of follow-up are used. The added value of backfilling depends on the escalation scheme used, with

most value potentially added for a less conservative scheme.

Here we have used a target DLT rate of 0.391 for the setting with three cycles of follow-up, as this corresponds to the target of 0.3 for one cycle of follow up. However, similar patterns of results have been observed when the target is 0.3 over three cycles.

In our evaluations, patients for backfilling are available for recruitment immediately which clearly is an optimistic assumption. Moreover the use of backfilling does increase the number of patients in the study. Nevertheless we did see consistent benefits of backfilling. Specifically, one additional patient did yield a reduction of the study duration of approximately half a week.

Our investigations explore two different settings with respect to the prior distributions used. In the first setting the prior is chosen with a study in mind that does not plan to use backfilling, while the second considers backfilling as an option at the outset. Encouragingly we find that, irrespective of the setting, the benefits of backfilling on estimation of the MTD and duration of the study are fairly consistent.

Data Availability Statement

All data is simulated according to the specifications described.

Acknowledgements

This report is supported by the NIHR Cambridge Biomedical Research Centre (BRC-1215-20014). The views expressed in this publication are those of the authors and not necessarily those of the NHS, the National Institute for Health Research or the Department of Health and Social Care (DHSC). T Jaki and H Barnett received funding from UK Medical Research Council (MC_UU_00002/14).

References

1. Storer BE. Phase I trials. In Redmond C and Colton T (eds.) *Biostatistics in Clinical Trials*. Chichester: John Wiley & Sons, Ltd, 2001. pp. 337–342.
2. National Cancer Institute. Common Terminology Criteria for Adverse Events (CTCAE) v5.0, 2017. URL http://ctep.cancer.gov/protocolDevelopment/electronic_applications/ctc.htm.
3. Wheeler GM, Mander AP, Bedding A et al. How to design a dose-finding study using the continual reassessment method. *BMC Medical Research Methodology* 2019; 19(1): 1–15.
4. Liu S and Yuan Y. Bayesian optimal interval designs for phase I clinical trials. *Journal of the Royal Statistical Society: Series C: Applied Statistics* 2015; : 507–523.
5. Zhou Y and Whitehead J. Practical implementation of bayesian dose-escalation procedures. *Drug information journal: DIJ/Drug Information Association* 2003; 37(1): 45–59.
6. Hamilton EP, Dees EC, Wang JSZ et al. Phase I dose escalation of H3B-6545, a first-in-class highly Selective ER α Covalent Antagonist (SERCA), in women with ER-positive, HER2-negative breast cancer (HR+ BC). *Journal of Clinical Oncology* 2019; 37(15_suppl): 1059–1059. DOI:10.1200/JCO.2019.37.15_suppl.1059. URL https://doi.org/10.1200/JCO.2019.37.15_suppl.1059.
7. Pauff JM, Papadopoulos KP, Janku F et al. A phase I study of LY3410738, a first-in-class covalent inhibitor of mutant IDH1 in cholangiocarcinoma and other advanced solid tumors. *Journal of Clinical Oncology* 2021; 39(3_suppl): TPS350–TPS350. DOI:10.1200/JCO.2021.39.3_suppl.TPS350. URL https://doi.org/10.1200/JCO.2021.39.3_suppl.TPS350.
8. Dehbi HM, OQuigley J and Iasonos A. Controlled backfill in oncology dose-finding trials. *Contemporary clinical trials* 2021; : 106605.
9. Tao JJ, Cangemi NA, Makker V et al. 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* 2019; 25(18): 5458–5465.
10. Eisenhauer EA, Therasse P, Bogaerts J et al. New response evaluation criteria in solid tumours: revised recist guideline (version 1.1). *European journal of cancer* 2009; 45(2): 228–247.
11. U.S. National Library of Medicine An Open-label, First-in-human, Multi-center Study to Evaluate the Safety, Tolerability, Pharmacokinetics and Anti-tumor Activity of a Thorium-227 Labeled Antibody-chelator Conjugate, BAY2287411 Injection, in Patients With Solid Tumors Known to Express Mesothelin URL <https://clinicaltrials.gov/ct2/show/NCT03507452>

12. Neuenschwander B, Branson M and Gsponer T. Critical aspects of the Bayesian approach to phase I cancer trials. *Statistics in Medicine* 2008; 27: 2420–2439. DOI:10.1002/sim.
13. Cheung YK and Chappell R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* 2000; 56(4): 1177–1182. DOI:10.1111/j.0006-341x.2000.01177.x.
14. O’Quigley J, Paoletti X and Maccario J. Non-parametric optimal design in dose finding studies. *Biostatistics* 2002; 3(1): 51–56. DOI: 10.1093/biostatistics/3.1.51.
15. Mozgunov P, Knight R, Barnett H et al. Using an interaction parameter in model-based phase I trials for combination treatments? A simulation study. *International Journal of Environmental Research and Public Health* 2021; 18(1): 1–19. DOI:10.3390/ijerph18010345.
16. Iasonos A and O’Quigley J. Adaptive dose-finding studies: A review of model-guided phase I clinical trials. *Journal of Clinical Oncology* 2014; 32(23): 2505–2511. DOI:10.1200/JCO.2013.54.6051. URL <https://doi.org/10.1200/JCO.2013.54.6051>. PMID: 24982451, <https://doi.org/10.1200/JCO.2013.54.6051>.




Figure 1. A single trial simulation with and without backfilling. Note that the backfilled cohorts in grey all enter the trial at the same time as the escalation cohort above in black, the horizontal offset is for ease of reading.




Figure 2. Proportion of correct and acceptable selections across scenarios for one cycle of follow-up




Figure 3. Proportion of correct and acceptable selections across scenarios for three cycles of follow-up




Figure 4. Relationship between mean total sample size and mean trial duration across scenarios for differing levels of backfilling. Darker colours indicate scenarios where a higher number of doses are unsafe.

Dose (mg/kg)	DLT	Stable Disease	Total subjects
0.25 (4 weekly)	0	1	4
0.50 (4 weekly)	1	3	6
0.50 (2 weekly)	0	3	4
0.75 (2 weekly)	0	1	3
1.00 (2 weekly)	0	2	4
2.00 (2 weekly)	0	3	4
4.00 (2 weekly)	0	2	4
8.00 (2 weekly)	2	1	3

Table 1. Results of the the first in-human Phase I study reported in⁹ by number of subjects and DLT per dose-level.

Scenario	1.5MBq	2.5MBq	3.5MBq	4.5MBq	6.0MBq	7.0MBq
1	0.30	0.40	0.50	0.60	0.70	0.80
2	0.20	0.30	0.40	0.50	0.60	0.70
3	0.10	0.20	0.30	0.40	0.50	0.60
4	0.05	0.10	0.20	0.30	0.40	0.50
5	0.05	0.10	0.15	0.20	0.30	0.40
6	0.02	0.05	0.10	0.15	0.20	0.30
7	0.15	0.20	0.25	0.30	0.45	0.60
8	0.05	0.15	0.30	0.35	0.40	0.45
9	0.40	0.45	0.50	0.55	0.60	0.65
10	0.05	0.15	0.25	0.35	0.45	0.55
11	0.15	0.20	0.35	0.40	0.45	0.50
12	0.05	0.10	0.15	0.20	0.25	0.40
13	0.06	0.07	0.08	0.09	0.11	0.12
14	0.10	0.14	0.21	0.30	0.46	0.58
15	0.16	0.30	0.50	0.70	0.89	0.95
16	0.55	0.91	0.99	1.00	1.00	1.00
17	0.05	0.05	0.05	0.80	0.80	0.80

Table 2. $p_{1,d_j} = P(DLT)$ in cycle 1 for each dose d_j across 17 defined simulation scenarios. MTD in **bold**.