Designing phase I oncology dose escalation using dose–exposure–toxicity models as a complementary approach to model-based dose–toxicity models

Kristyn Pantoja[1,2] | Shankar Lanke[2] | Alain Munafo[3] | Anja Victor[4] |
Christina Habermehl[4] | Armin Schueler[4] | Karthik Venkatakrishnan[2] |
Pascal Girard[3] | Kosalaram Goteti[2]

[1]Department of Statistics, Texas A&M University, College Station, Texas, USA

[2]EMD Serono Research Institute, Billerica, Massachusetts, USA

[3]Merck Institute for Pharmacometrics, Lausanne, Switzerland

[4]Merck KGaA, Darmstadt, Germany

Correspondence
Kosalaram Goteti, 45E Middlesex Turnpike, Billerica, MA 01821, USA.
Email: kosalaram.goteti@emdserono.com

Abstract

One of the objectives of oncology phase I dose-escalation studies has been to determine the maximum tolerated dose (MTD). Although MTD is no longer set as the dose for further development in contemporary oncology drug development, MTD determination is still important for informing the therapeutic index. Bayesian adaptive model-based designs are becoming mainstream in oncology first-in-human trials. Herein, we illustrate via simulations the use of systemic exposure in Bayesian adaptive dose–toxicity models to estimate MTD. We extend traditional dose–toxicity models to incorporate pharmacokinetic exposure, which provides information on exposure–toxicity relationships. We pursue dose escalation until the maximum tolerated exposure (corresponding to the MTD) is reached. By leveraging pharmacokinetics, dose escalation considers exposure and interindividual variability on a continuous rather than discrete domain, offering additional information for dose-escalation decisions. To demonstrate this, we generated 1000 simulations (starting dose of 1/25th the reference dose and six dose levels) for several different scenarios. Both rule-based and model-based designs were compared using metrics of potential safety, accuracy, and reliability. The mean results over simulations and different toxicity scenarios showed that model-based designs were better than rule-based methods and that exposure–toxicity model-based methods have the potential to valuably complement dose–toxicity model-based methods. Exposure–toxicity model-based methods had decreased underdose risk accompanied by a relatively smaller increase in overdose risk, resulting in improved net reliability. MTD estimation accuracy was compromised when exposure variability was large, emphasizing the importance of appropriate control of pharmacokinetic variability in phase I dose-escalation studies.

Study Highlights

WHAT IS THE CURRENT KNOWLEDGE ON THE TOPIC?

There are several experimental designs for phase I clinical trials in oncology. Bayesian methods offer more flexibility in the dose–toxicity model used in model-based designs such as continual reassessment method and Bayesian logistic regression model. Our method explores the performance of an exposure–toxicity model for dose escalation by relying on the relationship between the dose and the exposure metric.

WHAT QUESTION DID THIS STUDY ADDRESS?

We address one method for incorporating exposure metrics into experiment designs for dose escalation in phase I oncology clinical trials and compare performance to existing designs. We also evaluate the impact of pharmacokinetic (PK) variability on performance characteristics.

WHAT DOES THIS STUDY ADD TO OUR KNOWLEDGE?

We demonstrate that exposure can contribute to the efficiency of dose escalation and the accuracy of maximum tolerated dose (MTD) estimation and show that higher interpatient variability in systemic exposure can compromise MTD estimation accuracy.

HOW MIGHT THIS CHANGE DRUG DISCOVERY, DEVELOPMENT, AND/OR THERAPEUTICS?

Our results encourage evaluation of exposure–toxicity models to supplement dose–toxicity models in Bayesian dose-escalation designs in oncology and reinforce the importance of appropriate control of PK variability.

INTRODUCTION

One of the objectives of oncology phase I first-in-human (FIH) dose-escalation studies has been to determine the maximum tolerated dose (MTD) of a new drug. The characteristics of such studies have been described in literature.[1] Taken together with safety and tolerability, pharmacokinetics (PK), and pharmacodynamics (PD) for characterization of the bioactive dose/exposure range, the recommended dose range is defined for subsequent expansions.[2] Although it is important to note that the MTD is not of primary relevance for dose-selection decisions in the development of contemporary anticancer therapies,[3] its determination nevertheless informs the broader characterization of therapeutic index/window and evaluation of the consequence of supratherapeutic exposure settings associated with intrinsic or extrinsic factor effects. Safety and tolerability are usually highly correlated with dose or systemic exposures, forming the basis for dose–toxicity modeling for defining the MTD. In the in silico investigation described in the present study, we evaluate the performance characteristics of a design method based on a modeling framework that leverages the available information on PK exposure for a more informed dose escalation.

The MTD is typically defined as the highest dose with an acceptable toxicity. For cytotoxic drugs, there is a positive association between the toxicity and therapeutic effect of the drug: although higher doses are potentially more effective, they are also more toxic to the patient. Because the relationship between the dose and its toxicity is generally assumed to be monotonic, the dose is escalated in phase I clinical trials until the MTD is reached. Hence an acceptable level of toxicity must be specified. With the advent of new agents with different modes of action in oncology, the first objective of a dose-escalation study remains defining the toxicity profile of the drug including either the definition of an MTD (e.g., for cytotoxic agents) or a dose range that can be considered safe and pharmacologically active based on PK/PD considerations (e.g., for molecularly targeted agents). A broader objective is the definition of recommended dose(s) for expansion.[4]

Usually a dose-limiting toxicity (DLT) is defined as any adverse event (AE) that is of potential clinical significance such that further dosing of the patient or dose escalation would expose patients to unacceptable risk. A monotonic nature of the dose–toxicity relationship is typically assumed such that, as the dose increases, the probability of a DLT (toxicity level) also increases. Although the event of a DLT is unfavorable and should be avoided, it is acceptable that a small proportion of patients experience them at the MTD. In other words, the goal is to find the dose that corresponds to a level of toxicity (target toxicity level [TTL]) that is deemed acceptable. Usual target levels for MTD are 0.3 or 0.25.[5] Dose escalation occurs until a reliable estimate of the MTD is obtained. Although the MTD itself is not to be equated with a dose level that is suitable

for further expansion or evaluation in phase II studies, it represents the upper end of short-term tolerability, providing an upper boundary for further evaluations of PK, PD, safety, and long-term (multicycle) tolerability for selecting the recommended dose range for expansion and evaluation of the therapeutic index.

Dose-escalation designs for phase I clinical trials in oncology begin by assigning the first patient (or cohort of patients) to a starting dose and then continuing by subsequently assigning patients to adaptively chosen doses until some stopping rule is fulfilled or the MTD is determined. Doses are typically selected from a set of prespecified dose levels.

The historically most widely used dose-escalation design is the 3 + 3 design.[6,7] Dose escalation is based on a set of rules that determine when to escalate to the next dose, in some cases when to de-escalate, when to stop dose escalation, and how to determine the MTD. Such rule-based designs were favored for their simplicity and ease of use. However, the 3 + 3 design is known to be slow in dose escalation and less accurate compared with other methods because of its memoryless property by which it only considered the data from the last observed cohort.[8–10]

A more efficient and more accurate rule-based design is the Bayesian optimal interval (BOIN) method.[11] BOIN is considered a "model-assisted" rule-based design because it assumes an underlying dose–toxicity relationship where the MTD is defined in terms of the TTL, and dose escalation/de-escalation is determined by a set of rules that depend on certain prespecified thresholds for the sample proportion observed. Unlike the 3 + 3 design, BOIN offers more flexibility in sample size of cohorts.[12]

Other more efficient and accurate alternatives to the 3 + 3 design include model-based designs that have become increasingly applied in contemporary oncology drug development programs.[10] They define a parametric model for the dose–toxicity relationship to use all available toxicity information. In the Bayesian framework, the posterior probabilities of DLT at each dose are sequentially updated and summarized using all available data, which results in faster dose escalation and more accurate MTD estimates than the memoryless 3 + 3 design.[5,13,14] The MTD is often chosen as the dose with the posterior mean/median toxicity probability closest to the TTL or as the dose maximizing the probability of being in a target interval around the TTL. In addition, the model-based designs offer the potential to include additional data (e.g., PD data) besides DLT only.

The continual reassessment method (CRM)[15] is the first and most widely used model-based design. The original proposed method used a one-parameter power model and summarized the posterior probabilities of toxicity at dose d using the posterior mean. The dose with a posterior mean toxicity probability closest to the TTL is selected as

the next dose to gather data on. This procedure continues until the stopping rule is fulfilled. The original CRM has been criticized for its lack of safety and small-sample properties.[16,17] Neuenschwander et al.[16] claim that these shortcomings are due to a lack of expressiveness in the chosen model and an inadequate summarization of the posterior toxicity probabilities. They justify the use of a two-parameter logistic regression model and a posterior summary based on a loss function over the entire posterior distribution. The resulting method is called the Bayesian logistic regression model (BLRM).

The aforementioned designs in their most commonly used formulations do not account for interindividual variability in PK (i.e., systemic exposures) that should in principle explain in part the variability in the drug's safety and tolerability. For example, it is well recognized that failure to appropriately control important sources of PK variability (e.g., pharmacogenetic variation or clinically relevant drug–drug interactions with coadministered drugs that alter systemic exposures of the investigational drug) in dose-escalation trials can bias determination of the MTD.[18,19] In this article, we extend the dose–toxicity model to incorporate PK exposure, which provides information on the benefit–risk relationship of the drug, and we pursue dose escalation until the maximum tolerated exposure (MTE; corresponding to the MTD) is reached. By leveraging PK into the model, dose escalation occurs based on predictor data collected on a continuous rather than discrete domain, which allows for more informed dose escalation. It should be noted that methods for incorporating PK data into dose-escalation decisions have been suggested previously,[20–22] although opportunities remain to enable pragmatic integration of these frameworks in drug development settings.

In this article, we pragmatically describe how PK exposure can be incorporated with a BLRM-based, dose-escalation design. We compare different design methods and present our design comparisons based on metrics discussed in this manuscript. Furthermore, we perform scenario analyses of the impact of population variability in systemic exposures in the dose–exposure relationship on performance characteristics of dose-escalation designs incorporating exposure inputs. We trust that the results of our in silico evaluations will be useful for the oncology clinical research community to inform application of the exposure–toxicity model-based approach discussed herein or its variations as a complementary method alongside dose–toxicity models that are used in dose-escalation designs.

Synthetic (virtual) data generation for the analyses presented here was model free[23,24] and independent of each of the designs like 3 + 3, BOIN, CRM, BLRM, and other methods. This ensures that we have a fair comparison between model-based and rule-based designs because there is no model assumed in the data-generating process. We

show that our dose–exposure–toxicity model-based methods result in more efficient dose escalation compared with their dose–toxicity model-based method counterparts while maintaining performance in other areas. All designs are compared using measures of safety (overdose rate), accuracy (percent correct selection [PCS]), reliability (overdose risk, underdose risk), therapeutic efficacy (DLT) rate, and underdose rate.

METHODS

In this section, we introduce some notations and review two rule-based and two variants of a model-based design for dose escalation that we consider in our in silico experiments. Then, we describe a new dose-escalation design that incorporates PK exposure. Afterward, we discuss the setup of the in silico experiment and approach to the assessment of performance characteristics.

Let $\Omega_d = \{d_1, \ldots d_K\}$ denote the set of available doses, $d_1 < d_2 < \cdots < d_K$, where the starting dose d_1 has been determined from preclinical/translational investigations. This has traditionally been using animal species[25,26] and is increasingly leveraging quantitative pharmacologic integration and mechanism-based human translation, especially for novel mechanisms (e.g., immunotherapy) where considerations of the minimum anticipated biological effect level may be crucial.[25,27,28] The MTD is often defined as the dose $d_i \in \Omega_d$ with a toxicity level closest to the TTL, denoted by θ. Let Y denote the occurrence of a DLT, so that $Y = 1$ when a DLT has occurred and $Y = 0$ when a DLT does not occur. The toxicity level, or probability of a DLT, $P(Y = 1d)$, where d is a given dose, is estimated by a Bayesian model in model-based dose-escalation designs.

Dose-escalation designs

Review of rule-based designs: 3 + 3 and BOIN

3 + 3 design

The 3 + 3 design[7] is carried out by assigning a cohort of three patients to a dose and following a set of rules to determine whether to escalate to the next dose level for the next cohort. After assigning the first cohort to the first dose level d_1, the dose assignments for the next cohorts are determined by the number of DLTs in the preceding cohort. The MTD is declared after observing two or more DLTs at a given dose level as the highest tested dose where the observed DLT incidence is no greater than one in six evaluable patients. There are several variations of the 3 + 3 design implementation. We use the implementation described in Hansen et al.[6]

BOIN design

Given a prespecified TTL θ, BOIN[11] chooses corresponding escalation/de-escalation boundaries for observed DLT incidence at a given dose (λ_e and λ_d, respectively), which are used to determine whether to escalate/de-escalate after observing a patient/cohort. These boundaries are functions of ϕ_1, the highest DLT rate that is deemed to be underdosing, and ϕ_2, the lowest DLT rate that is considered overdosing. After the trial has completed, DLT rates for each of the doses are estimated using isotonic estimates, and the dose that has the isotonic estimate closest to θ is the selected MTD. For our simulation study, we use the local BOIN design with "standard" boundaries for $\phi_1 = 0.6\theta$ and $\phi_2 = 1.4\theta$, as given in Liu and Yuan.[11]

Model-based designs: BLRM variants

All model-based dose-escalation designs rely on a model to capture the dose–toxicity relationship, such as the two-parameter logistic regression model used by the BLRM design.[16]

$$\operatorname{logit}\{\pi_\theta(d)\} = \log(\beta_0) + \beta_1 \times \log\left(\frac{d}{d^*}\right) \quad (1)$$

where d is a dose in the set of prespecified doses Ω_d, $\pi_\theta(d) = P(\mathrm{DLT}|d, \beta_0, \beta_1)$ is the probability of a DLT, and β_0 and β_1 are parameters. A reference dose d^* is allowing for interpretation of β_0 as the odds of a DLT at d^*. We use normal priors for $\log(\beta_0)$ and $\log(\beta_1)$ with mean zero, standard deviation one, and no correlation.

BLRM using posterior mean
A posterior distribution for each potential dose can be derived, and dose escalation in a Bayesian model-based design will depend on the measure that is used to summarize this distribution. Our experiments include a variant of BLRM that uses the posterior mean as a summary. The next dose level is chosen as that for which the posterior mean estimate is closest to and lower than this posterior mean estimate for the MTD. This BLRM variant is similar to CRM[15] except it uses the two-parameter logistic regression model instead of the power model with one parameter. We call this variant BLRM1 in our experiment results where MTD was chosen as the dose with the posterior mean closest to TTL.

BLRM using a loss
Rather than taking a point summary of each dose's posterior distribution, the original BLRM[16] summarizes the whole distribution using a loss function that is defined over a partition of the distribution, such as:

$$L(\theta, d) = \begin{cases} l_1 = 1, & \text{if } \pi_\theta(x) \in (0, 0.2] \\ l_2 = 0, & \text{if } \pi_\theta(x) \in (0.2, 0.35] \\ l_3 = 2, & \text{if } \pi_\theta(x) \in (0.35, 0.6] \\ l_4 = 3, & \text{if } \pi_\theta(x) \in (0.6, 1] \end{cases} \quad (2)$$

where 1-0-2-3 is the loss function we used and {0.2, 0.35, 0.6} are the toxicity cutoff points for the partition. The loss function can be adjusted, and in fact we chose this more conservative loss function, compared with the original loss 1-0-1-2, for increased penalization for posterior probabilities that exceeded the target interval. Then dose escalation proceeds by recommending the dose that minimizes the Bayes risk: $l_1 P(\pi_\theta(x) \in (0, 0.2]) + l_2 P(\pi_\theta(x) \in (0.2, 0.35]) + l_3 P(\pi_\theta(x) \in (0.35, 0.6]) + l_4 P(\pi_\theta(x) \in (0.6, 1])$. In our experiment results, we call this variant BLRM2. Similar to BLRM1, the MTD is chosen as the dose with posterior mean closest to θ at the end of the trial.

Incorporating exposure into model-based designs

To incorporate PK exposure, we instead consider the exposure–toxicity relationship using the same two-parameter model in Equation (1):

$$\text{logit}\{\pi_\theta(x)\} = \log(\beta_0) + \beta_1 \times \log\left(\frac{x}{x^*}\right) \quad (3)$$

where x is the PK exposure, and x^* is a reference value for exposure. See Figure 1 for an illustration of the exposure–toxicity model.
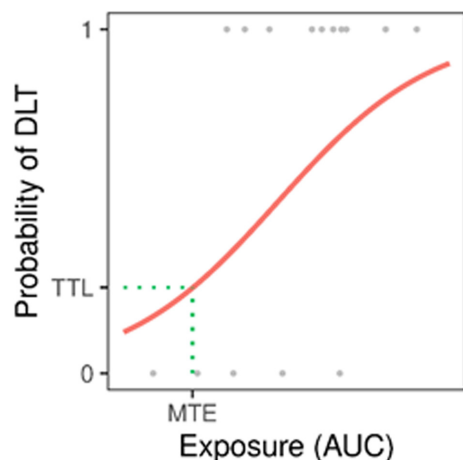

FIGURE 1 Exposure–toxicity model. DLT, dose-limiting toxicity; TTL, target toxicity level; MTE, maximum tolerated exposure (corresponding to maximum tolerated dose via the dose–exposure relationship); AUC, area under the curve.

Area under the curve (AUC) is an exposure metric that measures the total systemic exposure of the investigational agent and is widely used in exposure–response analyses for oncology drugs. Accordingly, we considered it to serve as a useful PK exposure metric for our model. AUC and dose are related by the following function:

$$\text{AUC}_i = \frac{d}{\text{CL}_i} \quad (4)$$

where CL_i is an individual patients' clearance.

By considering exposure as a function of dose and modeling the probability of toxicity as a function of exposure, we arrive at the dose–exposure–toxicity framework proposed.[22]

If clearance CL_i follows a log-normal distribution, we can formulate the dose–exposure–toxicity model as a Bayesian hierarchical model:

$$\text{CL}_i = \text{CL}_{\text{pop}} \times e^\eta \quad (5)$$

$$\text{AUC}_i = \frac{d}{\text{CL}_i} \quad (6)$$

$$y \sim \text{Bernoulli}(p) \quad (7)$$

$$p = P(\text{DLTAUC}_i) = f(\text{AUC}_i) \quad (8)$$

$$\eta \sim N(0, \omega^2) \quad (9)$$

with fixed paramters of population clearance $\text{CL}_{\text{pop}} = 5$, between-subject variability $\text{BSV} \in \{30\%, 60\%, 100\%\}$, and with $\omega^2 = \left(\frac{\text{BSV}}{100}\right)^2$. In Equation (3) we take x to be the observed values of AUC_i.

Dose escalation proceeds as follows:

1. The first patient (or cohort) is assigned to the first dose level.
2. Toxicity outcomes and PK information are collected.
3. The exposure–toxicity model is updated and used to recommend the next dose.
4. The next patient (or cohort) is assigned to the recommended dose.
5. Repeat Steps 2–4 until the stopping rule is fulfilled.
6. Estimate the MTD using the posterior mean toxicity probabilities of the doses.

In our simulations, we use a set of discrete doses. Given a target exposure, we can calculate the closest corresponding dose using the dose–exposure relationship in Equation (4). See Figure 2 for a diagram of the
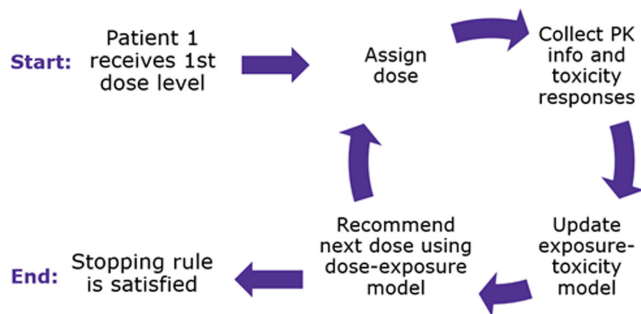
FIGURE 2 Diagram of the dose-escalation procedure using a dose–exposure–toxicity framework. PK, pharmacokinetic.

dose-escalation procedure using a dose–exposure–toxicity framework.

Experiment setup

We compare the 3 + 3, BOIN, BLRM1, BLRM2, and the dose–exposure–toxicity versions of BLRM1 and BLRM2 (which we will refer to as BLRME1 and BLRME2, respectively) on several performance metrics by conducting Monte Carlo simulations.

Following the simulation setup in Zhou et al.,[24] we use the model-free method described by Clertant and O'Quigley[23] to generate data from realistic exposure–toxicity relationships that are monotonic and bounded between 0 and 1. In the model-free method, we generate the relationship between exposure and toxicity using the pseudo-uniform algorithm[23] rather than assuming a parametric logistic regression model. In these simulations, the probability of DLT is determined by exposure (AUC) rather than dose. Therefore, results for all model-based designs are affected by the AUC generated in the simulated data.

This ensures that no model-based design is given an advantage in the simulation study. We compare two rule-based designs, 3 + 3 and BOIN, and two model-based designs, BLRM1 and BLRM2. We consider whether dose escalation in these two model-based designs informed by the emerging exposure–toxicity relationship, which we will call BLRME1 and BLRME2, improves any of the performance metrics that we consider.

Clertant and O'Quigley[23] propose a pseudo-uniform algorithm for generating a dose–toxicity relationship for a set of dose levels, a chosen MTD, and its corresponding TTL. We define the set of doses $\Omega_d = \{10, 20, 35, 50, 70,$ and $90\}$, where the starting dose $d_1 = 10$ is 1/25th the reference dose $d^* = 250$. To generate an exposure–toxicity curve, we calculate the dose levels' corresponding AUCs. In our experiments, we use $CL_{pop} = 5$, $BSV = 30\%$, and TTL $\theta = 0.25$. We run 1000 simulations for each design method. We consider the case where the true MTD is the fourth or fifth dose.

The 3 + 3 design implementation is described in Hansen et al.[6] BOIN is implemented using the R package boin, which uses the standard choices of escalation and de-escalation boundaries for a TTL of $\theta = 0.25$, that is, $\lambda_e = 0.197$ and $\lambda_d = 0.298$. The model-based designs BLRM1, BLRM2, BLRME1, and BLRME2 are implemented with the bcrm package.[29] The bcrm package allows us to apply the often used safety modification of no dose skipping and halt the design after 25 patients with a cohort of size 1. The same priors are used for BLRM1, BLRM2, BLRME1, and BLRME2. For BLRM2 and BLRME2, instead of using the original loss function 1-0-1-2 specified by Neuenschwander et al.,[16] we use a more conservative loss function 1-0-2-3, which prioritizes safety over accuracy. All designs halt after observing 25 patients with cohorts of size 1, except for the 3 + 3 design, which uses cohorts of size 3 and stops according to the rules specified previously.

The designs are compared on metrics that evaluate safety, accuracy, and reliability[24] metrics, and these metrics are described in the next section.

Performance metrics

Safety metrics

Although a key objective of oncology FIH studies has traditionally been to determine the MTD, given the narrow therapeutic range of many oncology therapies, the selected dose-escalation design should ensure patient safety as a priority. To evaluate the safety of a design, we look at the DLT rate over the whole trial (all dose levels), that is, the percentage of patients who experienced a DLT, averaged over all simulated trials.

Accuracy metrics

A design's ability to identify the MTD or at least estimate it with high accuracy is the main goal of these designs. For accuracy, we look at the PCS, that is, the percent of simulated trials in which the design selects the true MTD.

Reliability metrics

Although the average overdose rate of some designs may be comparable, that metric does not account for the variability in the overdose rates of the designs. It is possible that some of those designs may have resulted in simulated trials where most of the patients were overdosed, and these trials' overdose rates were averaged out with the rest

of the trials, whereas other designs had more consistent overdose rates. Zhou et al.[24] introduced metrics for reliability in terms of risking large percentages of patients being overdosed or underdosed in their simulated trials: the overdose risk, that is, the percent of simulated trials with more than 60% of patients treated above the MTD; and the underdose risk, that is, the percent of simulated trials with more than 80% of patients treated below the MTD.

RESULTS

Incorporating exposure into model-based designs

Tables 1 and 2 show the results of two rule-based designs, 3 + 3 and BOIN, and the following four model-based designs: BLRM1, BLRME1 incorporating exposure (BLRME1), BLRM2, and BLRM2 with exposure (BLRME2) when the true MTD is the fourth or fifth dose, respectively. Although BLRM1 and BLRM2 do not take exposure into account in the model analysis, their results are nonetheless affected because the observed DLTs were simulated from

exposure in the data generation. The metrics for comparing the models are calculated using 1000 simulations for each design method.

Our simulation results show that all model-based dose-escalation methods are more accurate than the rule-based dose-escalation methods, although the rule-based BOIN is not far behind. This is consistent with the history of comparisons between rule-based and model-based designs. We also see that incorporating exposure can sometimes improve the accuracy, such as in BLRME1 and BLRME2 when the MTD is the fourth dose (Table 1) and BLRME1 when the MTD is the fifth dose (Table 2).

The overdose and underdose risks characterize the reliability of the design methods by looking at what proportion of simulations overdosed or underdosed the majority of patients. Although the overdose risk is sometimes larger when incorporating exposure into the model, such as for BLRME2, the underdose risk is decreased markedly when exposure is used.

Overall, we can conclude that by incorporating exposure, we see enhancement in reliability. In addition, our results, consistent with previously reported findings, demonstrate that model-based dose-escalation designs are more accurate than rule-based designs.

TABLE 1 Designs compared on performance metrics when MTD is fourth dose, TTL is 0.25, and BSV = 30%

	3 + 3	BOIN	BLRM1	BLRME1	BLRM2	BLRME2
PCS (%)	0.20	0.31	0.31	0.34	0.35	0.36
Overdose risk (%)	0.00	0.09	0.07	0.11	0.06	0.08
Underdose risk (%)	0.46	0.31	0.25	0.15	0.28	0.21
DLT rate (average %)	0.21	0.24	0.24	0.27	0.23	0.25

Note: PCS: the percentage of simulations in which the dose-escalation design correctly selected the true MTD as their estimated MTD. Overdose risk: the percentage of simulations in which more than 60% of patients were treated above the MTD. Underdose risk: the percentage of simulations in which more than 80% of patients were treated below the MTD. DLT rate: the average percentage of patients who experienced a DLT (averaged over the simulations).

Abbreviations: 3 + 3, 3 + 3 design; BLRM1, Bayesian logistic regression model using posterior mean; BLRM2, Bayesian logistic regression model using loss function; BLRME1, Bayesian logistic regression model with exposure using posterior mean; BLRME2, Bayesian logistic regression model with exposure using loss function; BOIN, Bayesian optimal interval design; BSV, between-subject variability; DLT, dose-limiting toxicity; MTD, maximum tolerated dose; PCS, percent correct selection; TTL, target toxicity level.

TABLE 2 Designs compared on performance metrics when MTD is fifth dose, TTL is 0.25, and BSV = 30%

	3 + 3	BOIN	BLRM1	BLRME1	BLRM2	BLRME2
PCS (%)	0.15	0.25	0.18	0.23	0.26	0.26
Overdose risk (%)	0.00	0.06	0.10	0.16	0.07	0.07
Underdose risk (%)	0.63	0.40	0.39	0.22	0.40	0.28
DLT rate (average %)	0.18	0.22	0.22	0.25	0.21	0.23

Note: PCS: the percentage of simulations in which the dose escalation design correctly selected the true MTD as their estimated MTD. Overdose risk: the percentage of simulations in which more than 60% of patients were treated above the MTD. Underdose risk: the percentage of simulations in which more than 80% of patients were treated below the MTD. DLT rate: the average percentage of patients who experienced a DLT (averaged over the simulations).

Abbreviations: 3 + 3, 3 + 3 design; BLRM1, Bayesian logistic regression model using posterior mean; BLRM2, Bayesian logistic regression model using loss function; BLRME1, Bayesian logistic regression model with exposure using posterior mean; BLRME2, Bayesian logistic regression model with exposure using loss function; BOIN, Bayesian optimal interval design; BSV, between-subject variability; DLT, dose-limiting toxicity; MTD, maximum tolerated dose; PCS, percent correct selection; TTL, target toxicity level.

TABLE 3 Impact of pharmacokinetic variability on performance metrics of dose-escalation designs for an illustrative scenario of fourth dose as MTD: BSV = {30%, 60%, 100%}

	BLRM2 30%	BLRME2 30%	BLRM2 60%	BLRME2 60%	BLRM2 100%	BLRME2 100%
PCS (%)	0.35	0.36	0.28	0.30	0.16	0.19
Reached MTD (%)	0.88	0.94	0.84	0.95	0.71	0.93
Overdose risk (%)	0.06	0.08	0.04	0.06	0.03	0.04
Underdose risk (%)	0.28	0.21	0.39	0.22	0.57	0.28
DLT rate (average %)	0.23	0.25	0.24	0.27	0.24	0.28

Note: PCS: the percentage of simulations in which the dose-escalation design correctly selected the true MTD as their estimated MTD. Reached MTD: percentage of simulations in which the MTD was reached. Overdose risk: the percentage of simulations in which more than 60% of patients were treated above the MTD. Underdose risk: the percentage of simulations in which more than 80% of patients were treated below the MTD. DLT rate: the average percentage of patients who experienced a DLT (averaged over the simulations).

Abbreviations: BLRM2, Bayesian logistic regression model using loss function; BLRME2, Bayesian logistic regression model with exposure using loss function; BOIN, Bayesian optimal interval design; BSV, between-subject variability; DLT, dose-limiting toxicity; MTD, maximum tolerated dose; PCS, percent correct selection; TTL, target toxicity level.

Impact of PK variability on performance characteristics of model-based dose-escalation designs

Scenario analyses were conducted by changing the value of BSV in our simulations from 30% to 60% or 100%. See Tables 3 and S1 for results for BLRM2 and BLRME2. We see that as BSV increases, accuracy (PCS) decreases across both designs. However, fewer of the simulated BLRM designs reached the MTD compared with BLRME2, which remains consistent across the different BSV values. This second metric for accuracy suggests that incorporating exposure into the dose-escalation algorithm leads to a higher chance at collecting data at the MTD. Such an outcome is still favorable because a post hoc analysis of the data may be able to determine the correct MTD. Although BLRM2 uses a loss function as a more holistic posterior summary, it is common to declare the MTD at the end of dose escalation by using the MTD posterior mean estimate, such as BLRM1. Overdose risk, underdose risk, and DLT rate all suffer when BSV is large for both BLRM2 and BLRME2.

DISCUSSION

Determination of the MTD in oncology phase I trials relies on estimating the underlying dose–toxicity relationship with adequate confidence. Bayesian dose-escalation studies leverage the totality of data across patients at all studied dose levels in escalation and are increasingly favored in oncology drug development over 3 + 3 designs.[10,16] In our simulations, we witnessed the general improvement in accuracy that model-based designs offer over rule-based designs, consistent with previous reports in the literature, particularly in comparison to the 3 + 3 design.[14] For example, the PCS of model-based methods shows a difference of more than 10% compared with the 3 + 3 design, and although overdose risk is negligible for the 3 + 3 design due to its extremely conservative dose escalation, its underdose risk is sometimes almost double or even triple that of the model-based designs. Meanwhile, the DLT rate does not have such drastic differences in comparison. These improvements are observed despite not incorporating any prior information into the choice of prior parameters. If a study allows for prior information to be incorporated into the choice of prior, then much larger improvements would likely be observed. We chose to present BLRM2 and BLRME2 using the more conservative loss because the results that we observed with the original loss function had slightly higher DLT rates. The choice of a rather conservative loss function in our simulations led to the BLRM models escalating a bit slower. However, by incorporating exposure, BRLME2 maintained conservative properties while reaching the MTD in an efficient manner.

We showed that incorporating exposure into a model-based design leads to a significant decrease in underdose risk, especially compared with the 3 + 3 design, while maintaining accuracy, which is measured in terms of PCS. Although the primary objective of phase I trials is safety characterization in relation to dose coupled with the characterization of PK and PD properties, patients entering FIH oncology phase I trials on investigational new drugs are typically candidates for clinical trials following disease progression on prior therapies. As such, although efficacy is not a primary objective, these trials are not without potential for benefit. In fact, a meta-analysis of phase I oncology trials on molecularly targeted agents suggested that the probability of both overall response and overall survival increased with increasing dose.[30] Similar findings have been observed in another evaluation where an overall response rate of 5% was described across 7330 patients enrolled across 175 contemporary phase I trials, albeit

with some variations in observed dose–response trends by modality/mechanism of action.[31] Accordingly, reduction in the underdose risk may be an important consideration in phase I dose-escalation trials and may be particularly relevant for more traditional mechanisms of action such as small molecule receptor tyrosine kinase inhibitors, directly administered cytotoxic agents, or cytotoxic mechanisms delivered via modalities such as antibody–drug conjugates. Our simulations show substantial improvement in minimizing the underdose risk when exposure is incorporated into a model-based design, especially compared with the 3 + 3 design where the underdose risk is more than double and, in some instances, more than triple that of a BLRME approach.

Of note, for some modalities such as bispecific antibodies, nonmonotonic bell-shaped dose/exposure–response relationships may be expected from first pharmacologic principles and will demand consideration as part of the overall dose optimization program.[32] In this context, it is important to note that the MTD itself is not to be equated with a dose level that is suitable for further expansion or evaluation in phase II studies.[3] Although determination of the MTD informs the upper end of the dose range associated with short-term tolerability, determination of the recommended dose range for further expansion requires the holistic integration of PK, PD, safety, and multicycle tolerability. This should be performed in the context of preclinical prior knowledge, including consideration of the target molecular profile for precision medicines, integrated through translational modeling and simulation applying a totality-of-evidence approach.[33]

Following the simulation study of Zhou et al.,[24] our simulations of each dose-escalation design enlist 25 patients each, except for the 3 + 3 design, which is defined by its stopping rule. There are several valid stopping criteria for model-based (and model-assisted) designs. We avoid choosing a stopping criterion for our simulations because doing so would lead to simulation results that would depend on that choice. Instead, our main purpose in this article is to compare the dose-escalation methods on safety, accuracy, and reliability. Because our PK approach resulted in a greater probability of reaching the true MTD compared with the original BLRM dose-escalation design method, this suggests that fewer patients may be required when our method is used, regardless of what the stopping criterion is as long as the same stopping rule is used for each design method.

Although our simulations use BLRME with AUC as the exposure metric to help guide dose escalation, alternative exposure metrics may also be suitable. For example, maximum concentration (C_{max}) may be important to evaluate as an alternate exposure metric in these models if the commonly observed AEs are anticipated to be C_{max}

driven. Because AUC and C_{max} are typically correlated unless multiple dosing frequencies administering the same total dose are evaluated, we expect that the simulation results using BLRME with C_{max} would be similar to our results that use AUC.

Incorporating exposure into a design will necessitate some logistic considerations from two perspectives. One in terms of the bioanalysis data of the plasma concentrations of the therapeutic being available during the dose-escalation meeting and another in terms of the availability of sufficient data to build a population PK model or other credible approach to describing the dose–exposure relationship and associated variability. As mentioned previously, the collaboration between pharmacometrics, biostatisticians, and clinicians is crucial to increase this efficiency that may offer valuable insights in the study in the longer term to determine the appropriate MTD. In the first few cohorts, there will not be a population PK model available to estimate values such as exposure and clearance, which are needed for the BLRME1 and BLRME2 design methods. However, we may obtain estimates via noncompartmental analysis to provide a starting point for the study. Prior to the availability of clinical PK data, an animal PK model can be extended to humans and used to provide a starting point for dose escalation, as described in early work by Graham and Workman,[34] with many contemporary methods of human PK predictions being widely used currently.[35] As sufficient data are obtained, the exposure metrics from population PK models can be used to guide the selection of doses to reach the MTD. In addition, waiting for the PK data from the last patient in the cohort would delay the safety review to inform dose-escalation decisions. One solution might be to work on lagging PK data available from prior cohorts. In this approach, the previous PK data are used to estimate the latest patient's exposure value and update the exposure–toxicity model for assigning the next patient's dose. Lastly, we envision a hybrid assessment whereby BLRM is used to determine dose assignments for the first few dose levels until a PK model is available, and at that point BLRME can take over the rest of the dose assignments until the dose–toxicity relationship is adequately characterized based on the totality of dose-related safety and PK data. This approach would be especially useful when the first few doses have a large number of PK data points that are below the lower limit of quantitation. Further exploration using a real-time case example should help shed some light into the pros and cons of such a proposed approach.

Ultimately, the design would need to be applied in clinical studies to evaluate the extent of benefit gained in practice from the incorporation of exposure in model-based dose escalation for dosing decisions. Our proposed exposure–toxicity modeling approach

offers the potential to provide valuable insights on the relationship between the drug's systemic exposure and the toxic response and serves as a complement to the dose–toxicity model in model-based dose-escalation methods. Another important finding in our scenario analyses is the adverse impact of increasing PK variability on performance characteristics across designs, notably the accuracy of MTD estimation. This observation reinforces the importance of appropriate control of PK variability in oncology phase I trials both for intrinsic (e.g., renal/hepatic function) and extrinsic factors (e.g., formulation, dosing conditions in relation to food intake or drug–drug interactions) based on absorption, distribution, metabolism, and excretion; formulation scientific considerations; and data available from nonclinical studies ahead of initiation of clinical development.

In summary, the in silico analyses presented herein provide an assessment of the incorporation of variability in systemic exposure in the interpretation of dose–toxicity relationships evaluated in oncology phase I dose-escalation studies. Our findings highlight the importance of cross-functional collaboration across the quantitative disciplines of clinical pharmacology/pharmacometrics and biostatistics to maximize the value of information gained from PK/PD and safety assessments in the early clinical development of oncology therapeutics.

AUTHOR CONTRIBUTIONS

FUNDING INFORMATION

CONFLICT OF INTEREST

ORCID

Armin Schueler https://orcid.org/0000-0003-3833-5757

REFERENCES

1. Le Tourneau C, Lee JJ, Siu LL. Dose escalation methods in phase I cancer clinical trials. *J Natl Cancer Inst*. 2009;101:708-720.

2. Bullock JM, Rahman A, Liu Q. Lessons learned: dose selection of small molecule-targeted oncology drugs. *Clin Cancer Res*. 2016;22:2630-2638.

3. Shah M, Rahman A, Theoret MR, Pazdur R. The drug-dosing conundrum in oncology – when less is more. *N Engl J Med*. 2021;385:1445-1447.

4. Cook N, Hansen AR, Siu LL, Abdul Razak AR. Early phase clinical trials to identify optimal dosing and safety. *Mol Oncol*. 2015;9:997-1007.

5. Iasonos A, O'Quigley J. Adaptive dose-finding studies: a review of model-guided phase I clinical trials. *J Clin Oncol*. 2014;32:2505-2511.

6. Hansen AR, Graham DM, Pond GR, Siu LL. Phase 1 trial design: is 3 + 3 the best? *Cancer Control*. 2014;21:200-208.

7. Storer BE. Design and analysis of phase I clinical trials. *Biometrics*. 1989;45:925-937.

8. Iasonos A, Wilton AS, Riedel ER, Seshan VE, Spriggs DR. A comprehensive comparison of the continual reassessment method to the standard 3 + 3 dose escalation scheme in Phase I dose-finding studies. *Clin Trials*. 2008;5:465-477.

9. Jaki T, Clive S, Weir CJ. Principles of dose finding studies in cancer: a comparison of trial designs. *Cancer Chemother Pharmacol*. 2013;71:1107-1114.

10. Nie L, Rubin EH, Mehrotra N, et al. Rendering the 3 + 3 design to rest: more efficient approaches to oncology dose-finding trials in the era of targeted therapy. *Clin Cancer Res*. 2016;22:2623-2629.

11. Liu S, Yuan Y. Bayesian optimal interval designs for phase I clinical trials. *J R Stat Soc Ser C Appl Stat*. 2015;64:507-523.

12. Yuan Y, Hess KR, Hilsenbeck SG, Gilbert MR. Bayesian optimal interval design: a simple and well-performing design for phase I oncology trials. *Clin Cancer Res*. 2016;22:4291-4301.

13. Conaway MR, Petroni GR. The impact of early-phase trial design in the drug development process. *Clin Cancer Res*. 2019;25:819-827.

14. Le Tourneau C, Gan HK, Razak AR, Paoletti X. Efficiency of new dose escalation designs in dose-finding phase I trials of molecularly targeted agents. *PLoS One*. 2012;7:e51039.

15. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase 1 clinical trials in cancer. *Biometrics*. 1990;46:33-48.

16. Neuenschwander B, Branson M, Gsponer T. Critical aspects of the Bayesian approach to phase I cancer trials. *Stat Med*. 2008;27:2420-2439.

17. Sverdlov O, Wong WK, Ryeznik Y. Adaptive clinical trial designs for phase I cancer studies. *Statistics Surveys*. 2014;8:2-44.

18. Venkatakrishnan K, Pickard MD, von Moltke LL. A quantitative framework and strategies for management and evaluation of metabolic drug-drug interactions in oncology drug development: new molecular entities as object drugs. *Clin Pharmacokinet*. 2010;49:703-727.

19. Yamamoto N, Murakami H, Hayashi H, et al. CYP2C19 genotype-based phase I studies of a c-Met inhibitor tivantinib in combination with erlotinib, in advanced/metastatic non-small cell lung cancer. *Br J Cancer*. 2013;109:2803-2809.

20. Gunhan BK, Weber S, Friede T. A Bayesian time-to-event pharmacokinetic model for phase I dose-escalation trials with multiple schedules. *Stat Med*. 2020;39:3986-4000.

21. Piantadosi S, Liu G. Improved designs for dose escalation studies using pharmacokinetic measurements. *Stat Med*. 1996;15:1605-1618.

22. Ursino M, Zohar S, Lentz F, et al. Dose-finding methods for Phase I clinical trials using pharmacokinetics in small populations. *Biom J*. 2017;59:804-825.

23. Clertant M, O'Quigley J. Semiparametric dose finding methods. *J R Stat Soc Series B Stat Methodology*. 2017;79:1487-1508.

24. Zhou H, Yuan Y, Nie L. Accuracy, safety, and reliability of novel phase I trial designs. *Clin Cancer Res*. 2018;24:4357-4364.

25. International Council for Harmonisation. ICH guideline S9 on nonclinical evaluation for anticancer pharmaceuticals. In: Use CfMPfH, editor; 2010.

26. DeGeorge JJ, Ahn CH, Andrews PA, et al. Regulatory considerations for preclinical development of anticancer drugs. *Cancer Chemother Pharmacol*. 1998;41:173-185.

27. International Council for Harmonisation. Guideline on strategies to identify and mitigate risks for first-in-human and early clinical trials with investigational medicinal products. In: Use CfMPfH, editor; 2017.

28. Elmeliegy M, Udata C, Liao K, Yin D. Considerations on the calculation of the human equivalent dose from toxicology studies for biologic anticancer agents. *Clin Pharmacokinet*. 2021;60:563-567.

29. Sweeting M, Mander A, Sabin T. bcrm: Bayesian continual reassessment method designs for phase I dose-finding trials. *J Stat Softw*. 2013;54:1-26.

30. Gupta S, Hunsberger S, Boerner SA, et al. Meta-analysis of the relationship between dose and benefit in phase I targeted agent trials. *J Natl Cancer Inst*. 2012;104:1860-1866.

31. Hazim A, Mills G, Prasad V, Haslam A, Chen EY. Relationship between response and dose in published, contemporary phase I oncology trials. *J Natl Compr Canc Netw*. 2020;18:428-433.

32. Betts A, van der Graaf PH. Mechanistic quantitative pharmacology strategies for the early clinical development of bispecific antibodies in oncology. *Clin Pharmacol Ther*. 2020;108:528-541.

33. Venkatakrishnan K, Cook J. Driving access to medicines with a totality of evidence mindset: an opportunity for clinical pharmacology. *Clin Pharmacol Ther*. 2018;103:373-375.

34. Graham MA, Workman P. The impact of pharmacokinetically guided dose escalation strategies in phase I clinical trials: critical evaluation and recommendations for future studies. *Ann Oncol*. 1992;3:339-347.

35. Zou P, Yu Y, Zheng N, et al. Applications of human pharmacokinetic prediction in first-in-human dose estimation. *AAPS J*. 2012;14:262-281.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.