

RESEARCH ARTICLE

Extending the Continual Reassessment Method to accommodate step-up dosing in Phase I trials

Thomas M. Braun¹ | Francois Mercier²¹Department of Biostatistics, University of Michigan, Ann Arbor, Michigan, USA²Department of Biostatistics, Roche Innovation Centre, Basel, Switzerland**Correspondence**

Thomas M. Braun, Department of Biostatistics, University of Michigan, Ann Arbor, MI, USA.

Email: tombraun@umich.edu

The Continual Reassessment Method (CRM) was developed for Phase I trials to identify a maximum-tolerated dose of an agent using a design in which each participant is treated with a single administration of the agent. We propose an extension of the CRM in which participants receive multiple administrations of an agent using a so-called step-up dosing procedure in which participants receive one or more administrations of lower doses of the agent before they receive their penultimate dose. We use methods developed for the CRM to model the probability of DLT for each administration, which leads to the use of conditional probability models to model the joint probability of DLT across multiple administrations. We compare our approach to two existing methods that use time-to-event modeling methods for modeling the probability of DLT. We demonstrate through simulations that our approach has operating characteristics similar to existing methods, but due to its foundations in the CRM, ours is simpler to implement than existing approaches and is therefore more likely to be adopted in practice.

KEYWORDS

adaptive clinical trial, Bayesian inference, conditional probability, dose escalation, fractionated dosing, up-titration

1 | INTRODUCTION

Phase I trials in oncology are designed to identify safe, and possibly effective, doses of prospective treatments for cancer. A number of Bayesian adaptive designs for Phase I trials have been developed over the past three decades, the first of which was the Continual Reassessment Method (CRM).¹ In these designs, study participants are followed during a fixed period of time for the occurrence of a dose-limiting toxicity (DLT), which is often defined as any treatment-emergent adverse event of toxicity grade 3 or higher according to Common Terminology Criteria for Adverse Events (CTCAE) criteria.

In immunotherapy, DLTs often include cytokine release syndrome (CRS), which is a systemic inflammatory disease characterized by a massive release of cytokine.² CRS can present with a variety of symptoms ranging from those that are mild to those that are life threatening, and sometimes fatal.³ Mild symptoms of CRS include fever, fatigue, nausea, vomiting, headache, rash, arthralgia, myalgia, and malaise. Several case reports have documented CRS in cancer patients treated with immune-checkpoint inhibitors.⁴⁻¹¹

Recently, in the context of CAR T-cells, Stein et al¹² have proposed a model-based approach to retrospectively characterize the kinetics of tisagenlecleucel, and its relationship with the emergence of CRS. In their work, they share

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Statistics in Medicine* published by John Wiley & Sons Ltd.

concerns that co-medication with steroids or tocilizumab may not suffice to mitigate CRS. Instead, in a more recent and physiological-based model, Jiang et al¹³ suggest that a stepwise dosing of blinatumomab in patients with non-Hodgkin's lymphoma may successfully alleviate release of interleukin-6 and prevent CRS adverse events to occur.

To this end, a post-marketing authorization study, listed on ClinicalTrials.gov with identifier NCT01029366, was designed to assess the benefit of fractionated dosing in adult patients with relapsed/refractory acute lymphocytic leukemia treated by CAR T-cell therapy.¹⁴ Patients received either a low dose of tisagenlecleucel as a single infusion (LDS), or a high dose of tisagenlecleucel as a single infusion (HDS), or the high dose was split into fractions (HDF), with 10% of the dose delivered on day 1, 30% on day 2, and the remaining 60% on day 3. In the cohorts treated with a fractionated dosing approach, D2 and/or D3 doses were held if the patient experienced early signs of CRS, including fever. The frequency of grades 4 and 5 CRS was markedly lower in the HDF cohort (1 out of 20 patients) compared to the LDS (2 out of 9) or the HDS (3 out of 6) cohorts. As a result, fractionated or step-up dosing (SUD) has been proposed in other programs¹⁵ and seems to be a promising treatment option to mitigate CRS. However, challenges arise in the design and analysis of DLT data collected during dose escalation studies involving fractionated dosing.

As the goal is now to identify a maximum tolerated dose (MTD) that is preceded by one or more lower doses, which we call a maximum-tolerated schedule (MTS), the first challenge is to deal with probability of a DLT occurring after any of several consecutive dosing events. Both Colin et al¹⁶ and Fernandez et al¹⁷ proposed a logistic regression model with a Markov component for deal with DLTs observed over multiple cycles. Assuming that the drug was administered in consecutive cycles of similar length, an additive time homogeneous transition model could be specified as $\log(1 - p_{ik}) = -\alpha[d_{i1} - \rho d_{i,(k-1)}] - \beta d_{ik}$ where, for patient i receiving cycle k , p_{ik} is the probability of toxicity, d_{ik} is the scaled dose, and $d_{ik} = \sum_k d_{ik}$ is the cumulative administered scaled doses received. Respectively, the parameters α , β , and ρ capture the probability of a DLT on cycle 1, the effect of cumulative dose from previous cycles, and the dependency in short-term toxicity outcomes between cycles (Markov term). A limitation of this model is that there must be sufficient cycles in the DLT assessment period to allow a plausible estimation of ρ and β . Also, Fernandez et al did not consider the setting in which dose varies across dosing events during the DLT assessment period, making their design unsuitable to the situation of fractionated dosing.

A second challenge in the design of fractionated dosing studies occurs because later planned administrations are omitted once a participant experiences a DLT after earlier administrations. Indeed, it is routine practice to discontinue or hold treatment in patients who experience DLT. If the intention is to develop a dose-toxicity model akin to those used in the CRM, the missing administrations and resulting DLT outcomes will inflate the uncertainty associated with the model parameters and make the dose recommended for the next cohort less robust. In addition, by using SUD, we assume the dose of the first administration is lower than or equal to the dose of the second administration, which is lower than or equal to the dose of the third administration. Through this treatment plan, the absence of DLT associated with the second and/or third doses will lead to greater emphasis on DLT outcomes observed at lower doses, resulting in over-conservative dose escalation and loss of the convergence benefits of the CRM.¹⁸

To circumvent this problem, Braun and colleagues^{19,20} modeled DLT as an absorbing state that precludes any further administrations, using time-to-event concepts rather than logistic regression methods. As originally conceived, this work assumed that the hazard of DLT for a single administration increases linearly to a known future point in time and then decreases linearly to zero to a later point in time. The total hazard for several administrations is simply the sum of the hazards of the individual administrations that have been received, and the probability of DLT is based upon the survivor function derived from the cumulative hazard. By modeling the hazard, this approach moves away from a binomial likelihood and instead uses the likelihood of survival models, which accounts for partial follow-up of participants, much like the time-to-event Continual Reassessment Method (TITE-CRM)²¹ was developed to account for partial follow-up in Phase I trials designed with the CRM.

In a similar vein, Gunhan et al²² recently modeled the probability of DLT through a hazard function for each administration like the methods of Braun and colleagues. However, Gunhan et al chose to relate the DLT hazard to a parametric pharmacokinetic (PK) model quantifying the cumulative drug exposure over time. Gerard and colleagues²³ also created a design that integrates PK and pharmacodynamic (PD) models into the design. Given that in a more general setting, it may be unclear how to order consecutive administrations of doses with respect to their cumulative probabilities of DLT, Wages et al²⁴ presented an approach to incorporate partial ordering constraints and used existing CRM modeling for each administration.

Like many adaptive Phase I trial methods, many view the statistical underpinnings of dose and schedule-finding methods as too complex, a view which tends to hinder their implementation in practice. Thus, as an alternative to the time-to-event approaches cited earlier, we have adopted the longitudinal binary outcome view used by Fernandez et al.

However, our model is based on the CRM, which is a widely accepted adaptive design for Phase I trials of single administrations. By using the CRM as a framework, we hope to reach our primary goal with a simpler, more accepted model, that performs as well as more complex methods, but is more likely to be adopted by a wider audience of clinical trialists. The underlying specifics of our design can be found in Section 2. Via simulation, we compare the operating characteristics of our design to the designs of Braun et al and Gunhan et al in Section 3, and we present our concluding thoughts in Section 4.

2 | METHODS

2.1 | Notation

We have a study designed to examine a set of J predefined treatment schedules, each of which is a series of K administrations of an investigational agent. The study will enroll a total of N participants, each of whom will be assigned to one of the J treatment schedules. We denote schedule $j = 1, 2, \dots, J$ as $S_j = \{D_{j1}, D_{j2}, \dots, D_{jK}\}$, in which D_{jk} is the value for the dose of the agent given at administration $k = 1, 2, \dots, K$ in schedule j . We will describe how to select a value for D_{jk} in a later section.

Our primary goal is to identify which treatment schedule has a probability of DLT at the end of follow-up for the entire schedule closest to a targeted probability p^* . For $k < K$, we let τ_k denote the span of time between administrations k and $k - 1$. For each administration, we let $0 \leq t_{ik} \leq \tau_k$ denote the amount of follow-up observed for participant $i = 1, 2, \dots, N$ after receiving administration k , and we let $w_{ik} = t_{ik}/\tau_k$ denote the proportion of completed follow-up.

At the start of administration k , we set the DLT indicator $Y_{ik} = 0$, which changes to $Y_{ik} = 1$ if participant i experiences a DLT before τ_k . Once $Y_{ik} = 1$, we also set $w_{ik} = 1$, that is, we assume complete follow-up occurs, reflecting a recommendation of treatment discontinuation after a patient experiences a DLT. Furthermore, once $Y_{ik} = 1$, participant i receives no further planned administrations, while if $Y_{ik} = 0$ when $t_{ik} = \tau_k$, participant i receives their next administration.

2.2 | Model

For each administration, we model the marginal probability of DLT with a model traditionally used in the CRM, known as the “power” or “empiric” model. First, assuming that patient i is assigned to schedule j , we let $d_{ik} = D_{jk} \in (0, 1)$ be the value assigned to the dose given to participant i at administration k , and we assume

$$\pi_{i1} = Pr(Y_{i1} = 1 | d_{i1}) = d_{i1}^{\exp(\beta)} \quad (1)$$

$$\pi_{i2} = Pr(Y_{i2} = 1 | d_{i2}) = d_{i2}^{\exp(\beta - \theta_2)}$$

$$\pi_{i3} = Pr(Y_{i3} = 1 | d_{i3}) = d_{i3}^{\exp(\beta - \theta_2 - \theta_3)}$$

...

$$\pi_{iK} = Pr(Y_{iK} = 1 | d_{iK}) = d_{iK}^{\exp(\beta - \sum_{k=2}^K \theta_k)}, \quad (2)$$

in which $-\infty \leq \beta \leq \infty$, $\theta_k \geq 0$ for $k = 2, 3, \dots, K$. Note that we place no restriction on the value of β because the DLT probability for the first administration can take any value in $[0, 1]$. However, we do place a non-negativity constraint on each of $\theta_2, \theta_3, \dots, \theta_K$ to enforce an ordering constraint. Specifically, we assume that the cumulative probability of DLT cannot decrease with additional administrations. This restriction could be removed for another setting should this assumption not hold. However, less restriction on the parameters may also impact the ability to sufficiently identify those parameters, requiring continued focus on both appropriate prior distributions and the number of study participants needed to collect sufficient data to estimate those parameters.

Although our model can theoretically accommodate any number of administrations, most practical settings will examine a handful of administrations at most. Furthermore, we highlight that each administration $k \geq 2$ corresponds

to an additional model parameter θ_k . Given the relatively small sample sizes used in dose-finding trials, the number of administrations will have to be limited without strong constraints or assumptions placed on model parameters. As a result, the remainder of this manuscript will focus upon our motivating example that studied $K = 3$ administrations.

We will assume β has a normal prior distribution with mean μ_1 and SD σ , while θ_2 and θ_3 each have exponential prior distributions with respective means μ_2 and μ_3 . Note that β , θ_2 , and θ_3 are *a priori* independent of each other, and we will present a systematic approach for selecting values for the four hyperparameters in Section 2.5.

Because further planned administrations are not given to participants who experience DLT, the observed data cannot be used to directly estimate π_{i2} and π_{i3} . Instead, the observed data allow us to estimate the conditional probabilities of DLT for the second and third administrations, given no DLT was observed in all prior administrations. Nonetheless, the conditional probabilities are easy to generate from our model. Specifically, as derived in the Appendix, we have:

$$\begin{aligned}\phi_{i2} &= Pr(Y_{i2} = 1 \mid Y_{i1} = 0, d_{i1}, d_{i2}) = \frac{\pi_{i2} - \pi_{i1}}{1 - \pi_{i1}}. \\ \phi_{i3} &= Pr(Y_{i3} = 1 \mid Y_{i1} = Y_{i2} = 0, d_{i1}, d_{i2}, d_{i3}) = \frac{\pi_{i3} - \pi_{i2}}{1 - \pi_{i2}}.\end{aligned}$$

2.3 | Likelihood

Thus, at any point in the trial, we have enrolled a total of $M \leq N$ participants, who belong to one of three groups: those who have received the first administration, those who have received the first and second administrations, and those who have received all three administrations. We denote these three groups as \mathcal{G}_1 , \mathcal{G}_2 , and \mathcal{G}_3 , respectively. We let \mathcal{D}_1 denote the collective set of doses, DLT outcomes, and lengths of follow-up for all participants in \mathcal{G}_1 , with corresponding definitions for \mathcal{D}_2 and \mathcal{D}_3 . Thus, each of the three groups has a respective likelihood equal to:

$$\begin{aligned}L_1(\beta \mid \mathcal{D}_1) &= \prod_{i=1}^M \left[\pi_{i1}^{Y_{i1}} (1 - w_{i1} \pi_{i1})^{(1-Y_{i1})} \right]^{I(i \in \mathcal{G}_1)} \\ L_2(\beta, \theta_2 \mid \mathcal{D}_2) &= \prod_{i=1}^M \left[(1 - \pi_{i1}) \phi_{i2}^{Y_{i2}} (1 - w_{i1} \phi_{i2})^{(1-Y_{i2})} \right]^{I(i \in \mathcal{G}_2)} \\ L_3(\beta, \theta_2, \theta_3 \mid \mathcal{D}_3) &= \prod_{i=1}^M \left[(1 - \pi_{i2}) \phi_{i3}^{Y_{i3}} (1 - w_{i1} \phi_{i3})^{(1-Y_{i3})} \right]^{I(i \in \mathcal{G}_3)},\end{aligned}$$

which leads to the overall likelihood

$$L(\beta, \theta_2, \theta_3 \mid \mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3) = L_1(\beta \mid \mathcal{D}_1) \times L_2(\beta, \theta_2 \mid \mathcal{D}_2) \times L_3(\beta, \theta_2, \theta_3 \mid \mathcal{D}_3). \quad (3)$$

Note that the weights used in the likelihood are akin to those used in the time-to-event CRM (TITE-CRM),²¹ which assume that DLTs occur uniformly during the follow-up period, and are related to a cure model for the distribution of times to DLT.²⁵

2.4 | Specifying skeleton

Prior to study start, thought must be given to the numeric value assigned to each dose given at each administration. Given the complexity of selecting appropriate values for all of the $JK = 18$ administrations, we have developed a systematic approach to identifying values that allow for generally good operating characteristics across many settings, as we will show in Section 3.

We start by assigning $D_{11} = \delta$, where δ is a value relatively close to zero and can be seen as an approximate probability of DLT for the lowest dose examined in the study. In our application, we use $\delta = 0.03$; see Table 1. For the remaining first doses of schedules 2, 3, \dots , J , we choose values that increase by an amount defined by an odds ratio OR_b (“b” is for

TABLE 1 Actual dose values, assigned skeleton values, and hypothetical true toxicity probabilities for motivating example presented in Section 3

Schedule	Administration			Administration					
	1	2	3	1	2	3			
	Actual dose values			Skeleton dose values					
1	0.006	0.018	0.018	0.03	0.05	0.05			
2	0.030	0.090	0.090	0.06	0.10	0.10			
3	0.090	0.270	0.270	0.10	0.16	0.16			
4	0.270	0.800	0.800	0.16	0.25	0.25			
5	0.800	2.400	2.400	0.25	0.36	0.36			
6	2.400	7.200	7.200	0.36	0.50	0.50			
	True DLT probabilities								
	Scenario 1			Scenario 2			Scenario 3		
1	0.20	0.21	0.23	0.10	0.11	0.13	0.02	0.03	0.05
2	0.25	0.29	0.34	0.15	0.19	0.24	0.07	0.10	0.14
3	0.31	0.37	0.45	0.21	0.27	0.35	0.12	0.17	0.24
4	0.37	0.45	0.56	0.27	0.35	0.46	0.17	0.24	0.34
5	0.45	0.57	0.73	0.35	0.47	0.63	0.24	0.34	0.48
6	0.56	0.73	0.95	0.46	0.63	0.85	0.34	0.48	0.67
	Scenario 4			Scenario 5			Scenario 6		
1	0.01	0.02	0.03	0.01	0.02	0.03	0.01	0.01	0.02
2	0.03	0.06	0.10	0.04	0.06	0.08	0.03	0.04	0.06
3	0.06	0.10	0.15	0.06	0.10	0.12	0.05	0.07	0.10
4	0.08	0.15	0.24	0.09	0.14	0.15	0.07	0.09	0.13
5	0.11	0.21	0.34	0.13	0.20	0.26	0.10	0.13	0.18
6	0.16	0.30	0.48	0.18	0.28	0.36	0.14	0.19	0.27
	Scenario 7								
1	0.25	0.35	0.45						
2	0.30	0.40	0.50						
3	0.36	0.46	0.56						
4	0.42	0.52	0.62						
5	0.50	0.60	0.70						
6	0.61	0.71	0.81						

Note: Boldface represents the optimal schedule, i.e. the schedule with toxicity probability closest to 0.25 after the third administration.

between-schedules), that is, for $j = 2, 3, \dots, J$,

$$\frac{D_{j,1}}{1 - D_{j,1}} = OR_b \frac{D_{j-1,1}}{1 - D_{j-1,1}}.$$

We then select two additional odds ratios OR_{w1} and OR_{w2} (“w” is for within-schedules) such that for $j = 1, 2, \dots, J$,

$$\frac{D_{j2}}{1 - D_{j2}} = OR_{w1} \frac{D_{j1}}{1 - D_{j1}},$$

and

$$\frac{D_{j3}}{1 - D_{j3}} = OR_{w2} \frac{D_{j2}}{1 - D_{j2}}.$$

Values of the odds ratios should be sufficiently large enough so that successive schedules have DLT probabilities that are distinct enough from each other and promote discrimination between them with the traditional sample sizes used in dose-finding studies. We have found that appropriate values for each of the odds ratios are generally between 1.5 and 2.5, but they need to be determined in conjunction with the hyperparameter values described next. Selection of appropriate values for the odds ratios also can be informed by existing clinical information or previous dose-escalation studies that might inform how doses vary in their DLT probabilities and the cumulative effects of repeated dosing. Nonetheless, like any regression model, estimation of model parameters is facilitated by variability in the predictor variable, which, in our setting, are the values assigned to each administration. Even if one believes, for example, that OR_{w1} should be 1.1, such a study will be hard to implement without an unrealistically large number of participants, and doses with greater variability should be considered.

2.5 | Specifying hyperparameters

In order to determine appropriate prior means for each of β , θ_2 , and θ_3 , we take Equations (1) and (2) and define $\eta_{jk} = Pr(Y_{i1} = 1 | D_{ik} = d_{jk})$, which is the same value for every participant. With regard to the prior mean for β , we identify an appropriate value through the values $D_{11}, D_{21}, \dots, D_{61}$ assigned to the set of first doses among the schedules. Based upon a second-order Taylor series expansion for $\eta_{j1} = d_{j1}^{exp(\beta)}$, we outline in the Appendix that a suitable value for the prior mean is $\mu_1 = \sum_{j=1}^J \log[-\log(D_{j1})] / J$, which is the average of the transformed dose values. We then select $\sigma = \sqrt{\mu_1}$, so that the exponent $exp(\beta)$ in η_{j1} has mean $exp(\mu_1 + 0.5\sigma^2) = exp(1.5\mu_1)$.

For the prior means of θ_2 and θ_3 , we first select a value $k \geq 1$, such that $\eta_{j2} = k\eta_{j1}$ and $\eta_{j3} = k\eta_{j2}$, so that the probability of DLT within-schedule is assumed to increase proportionally with each administration. First focusing on the prior mean for θ_3 , we start with the fact that $\eta_{j3} = \eta_{j2}^{exp(-\theta_3)}$, which leads to $\theta_3 = -\log[\log(\eta_{j3}) / \log(\eta_{j2})]$. If we set $\eta_{j3} = p^*$ and given that $\eta_{j2} = \eta_{j3}/k$, we use the resulting equation for the prior mean of θ_3 , that is,

$$\mu_{\theta_3} = \frac{\log(p^*)}{\log(p^*) - \log(k)}.$$

By similar algebra, we find $\theta_2 = -\log[\log(\eta_{j2}) / \log(\eta_{j1})]$, which leads to the value

$$\mu_{\theta_2} = \frac{\log(p^*) - \log(k)}{\log(p^*) - 2\log(k)}.$$

Using these four hyperparameter values, we generate many draws of β , θ_2 , and θ_3 from their respective prior distributions, which allows us to compute many realizations of the probability of DLT for each administration of each schedule. Averaging over the realizations gives us an expected *a priori* probability of DLT for each administration of each schedule. These averages then allow us to see which schedule is now assumed to be the *a priori* best schedule and whether or not the prior distributions should be modified to produce suitable *a priori* DLT probabilities. For example, we would not want prior distributions that suggest that the first schedule is the best schedule, as this would likely be too informative and lead to poor operating characteristics if the true best schedule were the last schedule. To help the reader apply these ideas in practice, we will explore these thoughts in greater detail in Section 3 when presenting the prior distributions developed for the simulation study.

2.6 | Dose assignment algorithm

Once all necessary study design parameters, including sample size, skeleton dose values, and prior distributions, have been identified, participant assignments are made adaptively as follows:

1. The first participant is assigned to the schedule with the lowest starting dose;
2. Once a new participant $\ell = 2, 3, \dots, N$, can be enrolled, the accrued data for participants $i = 1, 2, \dots, (\ell - 1)$ are used to compute the posterior distributions of probability of DLT at the end of follow-up for each schedule;
3. Per a predefined stopping rule, if the accrued data suggest that the cumulative DLT probability after the follow-up for the last administration in first schedule is unacceptably high, the study ends early and no further accrual occurs because all schedules are considered unsafe.
4. If the stopping rule is not met, the new participant is assigned to the schedule with cumulative posterior mean probability at the end of follow-up closest to a desired target probability p^* , subject to any restrictions on assignments.
5. Repeat steps 2-4 as each new participant is accrued, or until the study is stopped.
6. If the study has not stopped accrual, once participant N has completed their follow-up:
 - (a) Use all of the accrued data to compute the cumulative posterior mean probability of DLT at the end of follow-up for each schedule;
 - (b) Select the schedule with posterior mean probability closest to p^* as the maximum tolerated schedule.

As suggested in step 3 above, Phase I trials often include a stopping rule when excessive toxicity is observed for the first schedule. Although stopping rules can be based upon posterior DLT probabilities, we have chosen instead to use a frequentist-based stopping rule that is independent of the model used for computing DLT probabilities. Specifically, we use the observed number of DLTs seen for participants assigned to the first schedule to compute a one-sided 95% confidence interval for the true DLT probability of the lowest schedule. If the lower bound of this confidence interval is higher than the targeted DLT probability p^* , the trial is stopped and future accrual is terminated. For example, the study would be stopped if four out of five participants assigned to schedule 1 experienced a DLT. This stopping rule is implemented in step 3 of the algorithm above and will be assessed in Section 3 with a setting in which all schedules have excessive probability of DLT.

In step 4 of the algorithm above, many possible restrictions on assignments can occur. For example, it is common to require that a schedule cannot be assigned to a participant until all lower schedules have been assigned to at least N_1 participants, or that participants must be enrolled in cohorts of size C , that is, all members of the same cohort must be given the same assignment. A common choice is $C = 3$, which is adopted from the so-called 3+3 algorithm.²⁶ We could also require that at least N_2 participants must have completed their follow-up on the same schedule before higher schedules can be assigned. Note that our parametric model, through the use of prior distributions, is able to estimate the DLT probabilities of later administrations without the direct observation of individuals who have received those administrations. Discomfort with decision-making on incomplete data is mitigated directly through these restrictions.

Furthermore, one might also consider the entire posterior distribution, beyond the posterior mean, when making assignment decisions. This approach, commonly referred to as Escalation with Overdose Control (EWOC),²⁷ examines how much mass of the posterior distribution lies above a threshold for each schedule's cumulative DLT probability. An upper bound is placed upon how much mass is acceptable and only schedules meeting this criterion are considered for assignment. Thus, it is possible that a schedule may have a posterior mean DLT probability that is closest to the targeted DLT probability, but its entire posterior distribution may be skewed or have too large a variance to be certain that the schedule is safe enough to assign to the next participant.

3 | SIMULATIONS

3.1 | Motivating example

Our motivating example is a study of six schedules ($J = 6$) that follow a SUD plan. Each schedule consists of three administrations ($K = 3$) of an experimental agent, each spaced apart by seven days ($\tau_1 = \tau_2 = 7$), and it is assumed that the dose used in each administration is no more than the dose used in the preceding administration. The actual clinical doses under study are shown in the first six rows of Table 1. Participants are followed for an additional seven days after the third administration, for a total planned follow-up of 21 days given to each participant. After each administration, patients are continually followed for the occurrence of a DLT, defined as any grade 3 or higher adverse event (per National Cancer Institute CTCAE v5.0), or occurrence of CRS grade 3 or above, according to the ASTCT consensus.²⁸ If a participant experiences a DLT, all further planned administrations are withheld for that participant, and they are considered to have complete follow-up.

3.2 | Simulation details

Our study wishes to identify which schedule is associated with a DLT probability by Day 21 closest to $p^* = 0.25$. A maximum of $N = 30$ participants will be enrolled. Participants will be enrolled in cohorts of size $C = 1$ and a schedule cannot be assigned to a participant unless at least $N_1 = 1$ participants have already been assigned to all lower schedules and at least $N_2 = 1$ participants have been fully followed on all lower schedules. Participant enrollment is assumed to follow a Poisson process with mean 21 days, that is, on average, each participant is enrolled after the previous participant has completed their follow-up.

Based upon $D_{11} = 0.03$, and $OR_w = OR_b = 1.5$, the skeleton values assigned to each dose are shown in the first six rows of Table 1. Based upon these skeleton values, we assign β a prior normal distribution with mean $\mu_1 = 0.91$ and SD $\sigma = \sqrt{\mu_1} = 0.95$. Assuming a value $k = 1.6$ defined in Section 2.5, we assign θ_2 a prior exponential distribution with mean $\mu_2 = 0.23$ and θ_3 a prior exponential distribution with mean $\mu_3 = 0.29$.

Figure 1A graphically presents the prior mean and SD for the DLT probabilities for each administration of each schedule that result from the hyperparameter and skeleton values just described. In this figure, we see that all three administrations of Schedules 1 and 2 are *a priori* likely to be safe, while Schedules 4-6 are less likely to be so. Specifically, although each has a mean prior probability of DLT at Day 21 that is below the target of 0.25, all of the latter three schedules have sufficient variability so that much of their prior distributions exceed the target.

To demonstrate how the chosen hyperparameter values impact the resulting prior distributions for DLT probabilities, Figure 1B displays information analogous to Figure 1A, but when the prior mean for β is 2.5 times larger. We now see stronger prior belief that all six schedules are safe, which might be implausible and lead to poor operating characteristics in a setting where Schedules 1 or 2 are the only acceptable schedules. Conversely, Figure 1C displays the resulting prior distributions when the prior means of θ_2 and θ_3 are each 2.5 larger. Now we see that there is greater prior uncertainty placed on all six schedules and we felt this prior would provide insufficient support to an algorithm incorporating data from only 30 participants. It is through these visual examinations that one can develop a suitable set of hyperparameter values before running any simulations.

To assess the operating characteristics of our design, we have simulated 1000 hypothetical trials for each of six scenarios, in which schedule $j = 1, 2, \dots, 6$ has a cumulative DLT rate after the third administration closest to p^* in scenario j . We also include a seventh scenario in which all schedules have Day 21 DLT probabilities above the target. The true DLT rates for each scenario are shown in Table 1. These DLT probabilities are not based on any specific model; values were selected so that neighboring schedules had DLT probabilities after the third administration spaced by approximately 10 percentage points, a difference that is biologically plausible, but is also large enough to allow for discriminating between the optimal schedule and its neighbors.


3.3 | Details for comparator approaches

We will compare the operating characteristics of our design to two other designs, both of which model a hazard function for the time-to-DLT and implicitly generate a cumulative probability of DLT by Day 21. Note that we have modified both methods from what was originally published in order to (i) accommodate use of the skeleton values used with our proposed model, and (ii) allow for dose levels that vary within each schedule. For the first comparator, we generalize the work of Braun et al,^{19,20} assuming each administration of the drug has a hazard function at time u equal to


$$h(u \mid \theta_1, \theta_2, \theta_3) = \begin{cases} \theta_2 \frac{u}{\theta_1} & 0 \leq u \leq \theta_1 \\ \theta_2 \frac{\theta_3 - u}{\theta_3 - \theta_1} & \theta_1 < u \leq \theta_3 \\ 0 & u > \theta_3 \text{ or } u < 0, \end{cases} \quad (4)$$

which increases linearly to a value θ_2 at time θ_1 and then decreases linearly to zero at time θ_3 .


If a participant is assigned to a schedule of doses d_1, d_2, \dots, d_K and dose $d_k, k = 1, 2, \dots, K$ is administered at time t_k , then the total hazard of DLT at future time u is $H(u \mid \theta_1, \theta_2, \theta_3) = \sum_{k=1}^K d_k h(u - t_k \mid \theta_1, \theta_2, \theta_3)$, which leads to a cumulative hazard function equal to $\Lambda(u \mid \theta_1, \theta_2, \theta_3) = \int_0^u H(t \mid \theta_1, \theta_2, \theta_3) dt$ and cumulative probability of DLT $F(u \mid \theta_1, \theta_2, \theta_3) = 1 -$



(A)



(B)



(C)

FIGURE 1 Visual representations of posterior DLT probability distributions resulting from three sets of hyperparameter values examined in Section 3. The height of each bar represents the posterior mean, and the length of each vertical line above the bar indicates the posterior SD. The horizontal dotted line represents targeted DLT probability of 0.25. S1 = schedule 1, S2 = schedule 2, ..., S6 = schedule 6

$\exp\{-\Lambda(u \mid \theta_1, \theta_2, \theta_3)\}$. The likelihood contribution at u is $f(u \mid \theta_1, \theta_2, \theta_3)^{Y(u)} \{1 - F(u \mid \theta_1, \theta_2, \theta_3)\}^{1-Y(u)}$, where $Y(u) = 1$ if a DLT has occurred at u and is zero otherwise, and $f(u \mid \theta_1, \theta_2, \theta_3) = dF(u \mid \theta_1, \theta_2, \theta_3)/du$.

In our simulations, we assume a fixed value of $\theta_3 = 10$, while $\theta_1/10$ has a prior Beta distribution with parameters $a = 5.8$ and $b = 3.9$, and θ_2 has an exponential distribution with mean 0.09. These values were selected so that the expected prior cumulative DLT probability at Day 21 for each schedule was similar to that used with the other methods.

For the second comparator, we generalized the work of Gunhan et al,²² replacing Equation (4) above with the difference of two exponential decay functions

$$h(u \mid \theta_1, \theta_2, \theta_3) = \frac{\theta_1 \theta_2}{\theta_2 - \theta_3} [\exp\{-\theta_3(u - t_k)\} - \exp\{-\theta_2(u - t_k)\}].$$

All remaining equations above for the total and cumulative hazards and the cumulative probability of DLT are unchanged. In the simulations, we have assumed fixed values for both θ_2 and θ_3 , such that $\theta_2 = 0.14$ and $\theta_3 = 0.35$. With regard to θ_1 , we assume $\log(\theta_1)$ has a normal distribution with mean -2 and SD 1, again so that the expected prior cumulative DLT probability of each schedule was similar to that of the other methods.

For both comparators, we use the same stopping rule and limits on escalation as used with our model, so that variations in the operating characteristics presented next are due primarily to the model that is used. All simulations were performed in R, version 4.0.4. Due to a vast savings in computation time, all posterior moments were computed through integral approximation methods, rather than through Markov Chain Monte Carlo methods. All code for our method, as well as the two comparator methods, is available on GitHub at <https://github.com/tombraun1216/CRM-with-Step-Up-Dosing.git>.

3.4 | Operating characteristics

The performance of our design is summarized by three metrics: (i) the proportion of simulations in which each schedule is selected as the MTS at the end of the study; (ii) the average proportion of participants assigned to each schedule during the study, and (iii) the average Day 21 DLT probability of the schedule assigned to each participant. Information regarding metrics (i) and (ii) is shown in Table 2, while information regarding metric (iii) is shown in Figure 2.

TABLE 2 Operating characteristics resulting from simulations described in Section 3.2

Scen	Method	Early term		Schedule 1		Schedule 2		Schedule 3		Schedule 4		Schedule 5		Schedule 6	
		Sel	Assn	Sel	Assn	Sel	Assn	Sel	Assn	Sel	Assn	Sel	Assn	Sel	Assn
1	CRM	0.08	n/a	0.68	0.66	0.20	0.17	0.05	0.09	0.00	0.03	0.00	0.01	0.00	0.00
	GWF	0.07	n/a	0.60	0.51	0.28	0.28	0.05	0.13	0.00	0.03	0.00	0.00	0.00	0.00
	TRI	0.06	n/a	0.72	0.39	0.16	0.27	0.06	0.25	0.00	0.06	0.00	0.01	0.00	0.00
2	CRM	0.01	n/a	0.32	0.38	0.40	0.29	0.23	0.21	0.04	0.09	0.00	0.03	0.00	0.01
	GWF	0.00	n/a	0.21	0.23	0.49	0.40	0.26	0.26	0.04	0.10	0.00	0.02	0.00	0.00
	TRI	0.00	n/a	0.42	0.19	0.28	0.27	0.27	0.38	0.03	0.13	0.00	0.02	0.00	0.00
3	CRM	0.00	n/a	0.02	0.12	0.25	0.22	0.41	0.31	0.27	0.23	0.04	0.09	0.00	0.03
	GWF	0.00	n/a	0.00	0.06	0.23	0.24	0.45	0.35	0.28	0.26	0.04	0.08	0.00	0.01
	TRI	0.00	n/a	0.12	0.08	0.15	0.16	0.46	0.39	0.24	0.29	0.03	0.07	0.00	0.01
4	CRM	0.00	n/a	0.00	0.06	0.03	0.09	0.21	0.19	0.42	0.29	0.30	0.24	0.05	0.13
	GWF	0.00	n/a	0.00	0.04	0.02	0.11	0.25	0.22	0.44	0.34	0.25	0.21	0.05	0.07
	TRI	0.00	n/a	0.04	0.04	0.02	0.10	0.23	0.22	0.42	0.38	0.27	0.20	0.02	0.06
5	CRM	0.00	n/a	0.00	0.05	0.01	0.07	0.06	0.11	0.29	0.22	0.41	0.28	0.24	0.26
	GWF	0.00	n/a	0.00	0.04	0.01	0.09	0.07	0.13	0.31	0.27	0.40	0.28	0.22	0.18
	TRI	0.00	n/a	0.02	0.04	0.01	0.09	0.07	0.12	0.32	0.31	0.43	0.28	0.16	0.16
6	CRM	0.00	n/a	0.00	0.05	0.00	0.05	0.02	0.08	0.12	0.14	0.34	0.24	0.52	0.44
	GWF	0.00	n/a	0.00	0.04	0.00	0.08	0.02	0.09	0.14	0.19	0.33	0.26	0.51	0.34
	TRI	0.00	n/a	0.01	0.04	0.00	0.08	0.02	0.09	0.16	0.22	0.39	0.27	0.42	0.31
7	CRM	0.75	n/a	0.24	0.48	0.00	0.05	0.00	0.02	0.00	0.01	0.00	0.00	0.00	0.00
	GWF	0.71	n/a	0.28	0.47	0.02	0.11	0.00	0.03	0.00	0.01	0.00	0.00	0.00	0.00
	TRI	0.57	n/a	0.41	0.38	0.02	0.22	0.01	0.16	0.00	0.02	0.00	0.00	0.00	0.00

Note: CRM, proposed extension of CRM; GWF, PK hazard method of Reference 22; TRI, triangular hazard method of References 19 and 20; Sel, proportion of simulations in which schedule was selected as best at end of study; Assn, average proportion of participants assigned to schedule during study; early term, no schedule selected due to early termination of study. Boldface text corresponds to operating characteristics corresponding to schedule with true DLT closest to target DLT probability of .25.

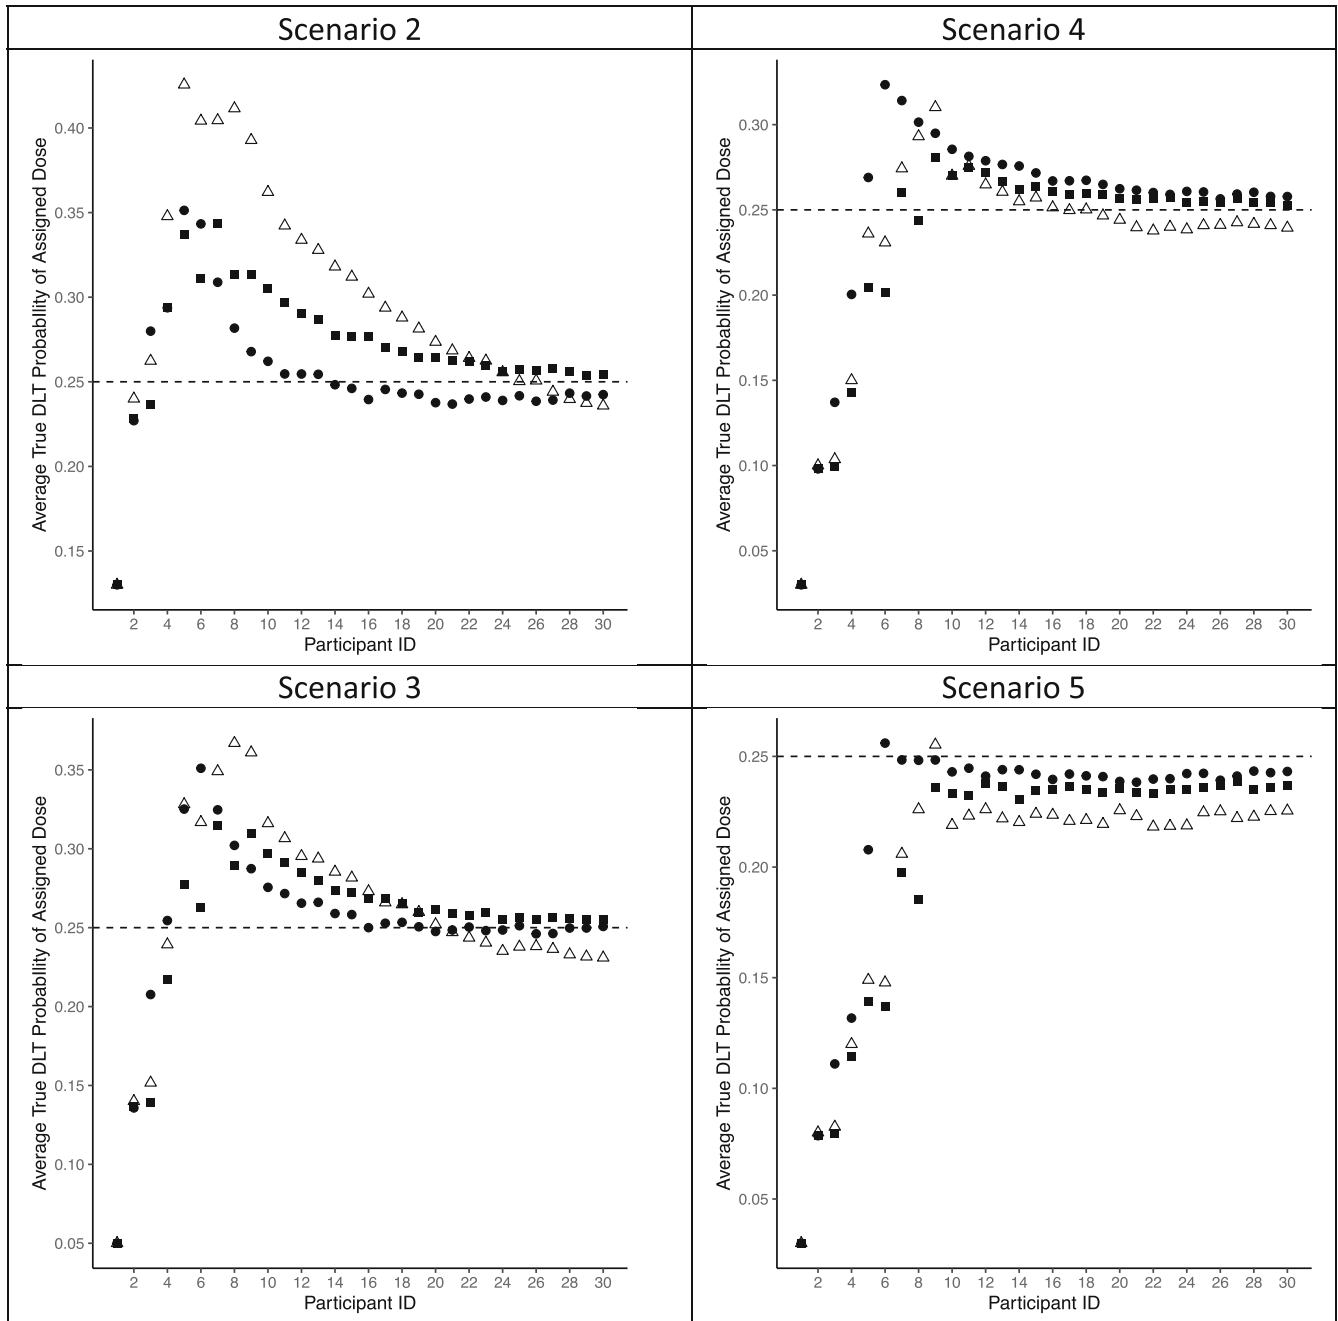


FIGURE 2 Schedule assignment patterns in Scenarios 2-5 of simulations described in Section 3. ●, proposed method using extension of CRM; ■, PK hazard method of Reference 22; △, triangular hazard method of References 19 and 20. Horizontal dashed line represents targeted DLT probability of .25

To start, we focus on the boldfaced values in Table 2, which indicate that final selection or assignment during the study is done at the schedule with Day 21 DLT probability closest to the target. With regard to the decision made at the end of the study, we see that all three methods have fairly similar performance; averaging across all scenarios, we have correct selection proportions of 0.47, 0.48, and 0.46 for our proposed method, the method of Gunhan et al, and the method of Braun et al, respectively. Furthermore, we extended the work of O’Quigley et al,²⁹ and generated, for each of the first six scenarios, a non-parametric optimal upper bound for probability of correct schedule selection at study end. Based upon 10 000 simulations, the realized upper bounds were 0.69, 0.57, 0.52, 0.49, 0.54, and 0.58 for scenarios 1-6, respectively, with an average of 0.57, which supports that all three methods have solid, but not overly optimistic, performance with 30 participants.

With regard to assignments made during the trial, 2 presents fairly comparable performance among the methods, although with greater variation than what was seen for the probability of correct selection at study end. Averaging across all six scenarios, we have average proportions of assignment equal to 0.38, 0.37, and 0.34 for our proposed method, the method of Gunhan et al, and the method of Braun et al, respectively. The greatest discrepancy among the methods appears in Scenario 2, in which both our proposed design and the triangular hazard design both assign fewer participants to schedule 2 than does the PK-based hazard design. This issue is also demonstrated in the upper left-hand panel in Figure 2, in which we see that the proposed design tends to assign Schedule 1 more often than Schedule 2, while the triangular hazard design tends to assign Schedule 3 more than Schedule 2.

To examine this phenomenon in more detail, we examined all six scenarios and tabulated, for each of the three designs, the difference in schedule assignments between a current participant who experiences a DLT and the next participant, that is, determining if the next assignment was at a lower schedule, the same schedule, or a higher schedule. What we found was that across the six scenarios, the probability of de-escalating by two or more schedule levels ranged from 0.05 to 0.13 for our proposed design, 0.01 to 0.03 for the PK hazard design, and 0.01 to 0.04 for the triangular hazard design. Conversely, the probability of escalating one schedule level ranged from 0.01 to 0.02 for our proposed design, 0.01 to 0.03 for the PK hazard design, and 0.07 to 0.11 for the triangular hazard design.

This differential is also supported by the results in Table 2 for scenario 7, in which all six schedules are overly toxic, so that dose assignments should be restricted to the lowest schedules, and the study should terminate before accruing all 30 participants. We see that the proposed method and PK-based hazard method both terminate accrual more often than the triangular hazard approach, which also assigns schedule 3 far more often than the others. Furthermore, early termination occurred after an accrual of 13, 14, and 19 participants, on average, for our method, the PK hazard design, and the triangular hazard design, respectively.

In Figure 2, we have presented the dose assignment patterns for Scenarios 2 to 5 and have omitted Scenarios 1 and 6 because they present redundant information from the other scenarios. We see across all four scenarios that all three methods have assignments that tend to coverage toward the schedule with a DLT probability closest to the target of .25. As expected, all three methods tend to escalate to schedules with DLT probabilities higher than .25, and then respond to observed DLTs and de-escalate in efforts to observe fewer DLTs. The greatest divergence among the three methods is seen for Scenario 2 and supports our earlier discussion of the results presented in Table 2, namely that our proposed method tends to be least aggressive with schedule assignments and the triangular hazard method tends to be most aggressive.

3.5 | Sensitivity analyses

Using the settings from Scenarios 1 to 6 presented in Table 1, we also examined how the operating characteristics of our proposed methods varied with (i) a different prior distribution for θ_2 and θ_3 , (ii) a different skeleton, which implicitly impacts the prior for β , (iii) a larger sample size of 45 participants, and (iv) a shorter average inter-arrival for participants of seven days. The results for aspects (i) and (ii) can be found in Table 3, while the results for aspects (iii) and (iv) can be found in Table S1 in the Appendix.

Assuming an increased value $k = 3.2$ (as described in Section 2.5), we assigned θ_2 a prior exponential distribution with mean $\mu_2 = 0.38$ and θ_3 a prior exponential distribution with mean $\mu_3 = 0.61$. From the first set of six rows in Table 3, we also see no substantial change to the operating characteristics for our design, although with slightly lower performance in scenario 6. The latter set of six rows in Table 3 correspond to operating characteristics when we changed the skeleton by using values $D_{11} = 0.01$, and $OR_w = OR_b = 2.0$ (as described in Section 2.4), which led to skeleton values that were shifted closer to zero. As with the prior distributions for θ_2 and θ_3 , we see little change in operating characteristics for scenarios 1 to 5, although the performance is decreased in scenario 6, which emphasizes that the skeleton values do impact the prior mean for β_1 , which is the likely cause of this result in scenario 6.

In Supplementary Table S1, the first six rows present the operating characteristics when the sample size is increased to 45 participants. Not surprisingly, we see that the probability of correctly selecting the best schedule increases approximately 10 points, with a slightly lower increase in the average proportion of participants assigned to the best schedule. With regard to the average inter-arrival time of patients, when participants arrived an average of every 7 days, rather than every 21 days, the last six rows of Supplementary Table S1 demonstrate no material change in the operating characteristics relative to those in Table 2.

TABLE 3 Simulation results for assessing sensitivity to prior distributions and skeleton

Sensitivity feature	Scen	Sched 1		Sched 2		Sched 3		Sched 4		Sched 5		Sched 6	
		Sel	Assn	Sel	Assn	Sel	Assn	Sel	Assn	Sel	Assn	Sel	Assn
Prior for θ_2 and θ_3	1	0.69	0.69	0.19	0.18	0.05	0.07	0.00	0.02	0.00	0.00	0.00	0.00
	2	0.33	0.42	0.40	0.28	0.22	0.19	0.04	0.08	0.00	0.02	0.00	0.00
	3	0.02	0.15	0.25	0.23	0.42	0.31	0.26	0.21	0.05	0.08	0.00	0.02
	4	0.00	0.08	0.04	0.11	0.23	0.22	0.43	0.29	0.26	0.21	0.04	0.10
	5	0.00	0.07	0.01	0.08	0.06	0.13	0.31	0.24	0.41	0.27	0.20	0.23
	6	0.00	0.05	0.00	0.06	0.03	0.10	0.14	0.16	0.35	0.24	0.48	0.39
Skeleton dose values	1	0.69	0.61	0.21	0.22	0.03	0.09	0.00	0.03	0.00	0.01	0.00	0.00
	2	0.28	0.31	0.47	0.35	0.22	0.23	0.03	0.08	0.00	0.02	0.00	0.00
	3	0.01	0.11	0.23	0.21	0.46	0.34	0.27	0.25	0.03	0.07	0.00	0.01
	4	0.00	0.05	0.02	0.09	0.22	0.20	0.47	0.34	0.27	0.24	0.03	0.08
	5	0.00	0.04	0.01	0.07	0.07	0.13	0.36	0.27	0.44	0.31	0.13	0.17
	6	0.00	0.04	0.00	0.05	0.03	0.10	0.18	0.19	0.42	0.31	0.36	0.32

Note: Sel, proportion of simulations in which schedule was selected as best at end of study; Assn, average proportion of participants assigned to schedule during study; boldface text corresponds to operating characteristics corresponding to schedule with true DLT closest to target DLT probability of .25.

4 | DISCUSSION

In recent Phase I dose-escalation trials assessing the tolerability of new investigational drugs in cancer immunotherapy, high rates of DLT after first dose have been reported. To address this issue, adjusted dosing schemes have been proposed that consist of planned stepwise dose-escalation for each participant at the start of treatment. These dosing schemes, known as SUD or dose fractionation, pose new challenges for the design of such a Phase I dose-escalation trial.

Certainly, identifying the MTD in a Phase I dose-escalation trial is always challenging because the process relies on sparse data. Each patient contributes a single data point, either a 0 if no DLT is observed during the DLT assessment period, or a 1 otherwise. With SUD schemes, the challenge of identifying the MTD is even larger, as a combination of doses have to be identified, each of which induces a probability of toxicity close to the target toxicity level.

To this end, we have proposed an approach for assessing the cumulative probability of DLT for a series of administrations given to participants in a SUD design, whereby participants are first given administrations of a lower dose in hopes of leading to less likelihood of DLT when the final desired dose of the agent is administered. Our design is a simple extension of the CRM, as compared to other approaches founded in time-to-event methods, and our design has operating characteristics comparable to those methods. Our method resembles a piecewise hazard model and is suitable for designs in which participants are treated with a succession of non-decreasing doses. One limitation of our method, is that each administration is assumed to occur within a specific timeframe, so our current model is unable to accommodate treatment delays and resulting DLT information, unless additional assumptions or modifications are made.

The number of parameters in our model also grows with the number of administrations, although we could include approaches to smooth among the parameters, correlate them through their prior distributions, or even assume that the change in DLT probability is the same for each additional administration. Nonetheless, our method does not make an assumption of additivity of hazards among the administrations, as is done with the approaches we used as comparators. And, like all designs using Bayesian methods, any data from a previous study of a related compound, including information on PK and PD patterns, should be considered to assist with the selection of prior distributions and all parameter and dose values.

Our proposed CRM extension requires specifying a skeleton, which is a set of plausible values ascribed to each dose and is often related to the *a priori* probability of DLT for single administration. Typically, the skeleton is chosen to cover a wide range of possible dose-toxicity profiles, and this choice can be difficult given the lack of prior knowledge on dose-toxicity at the study start. To help in this undertaking, we provide a systematic and intuitive approach for choosing the values

used in the skeleton. By simplifying the selection of the initial guesses of the probabilities of DLT in practice, we hope to enhance the use of our CRM-based approach.

Our method is also extremely flexible and allows for nearly any modification desired for a specific trial. For example, although our simulations assigned a dose to each participant individually, such assignments can be done to a collective cohort of participants, such as in groups of three participants, as a safeguard against early escalation to doses later seen to be toxic. Our method also allows for any dose-toxicity model used in the CRM, such as a standard logistic model with known intercept.

Nonetheless, before implementing our or any proposed design, clinical and safety experts should be consulted to evaluate the risk of planning to increase a participant's dose while on study. Additionally, a thoughtful and careful safety monitoring plan should be made and possibly include restrictions on how many participants can be exposed to a dose level, whether a participant is exposed to that dose initially as part of a new study cohort or is exposed during their planned dose escalation. Furthermore, our design is founded on a strong *a priori* belief that SUD is necessary. Although our design does allow for the estimation of DLT probabilities for single administrations of varying doses, as well as the ability to compare the DLT probability of a dose given immediately vs gradually via SUD, investigators must weight scientific interest in this latter comparison with the safety, costs, and time resources that might arise with repeated administrations.

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no new data were created or analyzed.

ORCID

Thomas M. Braun  <https://orcid.org/0000-0002-7113-6998>

Francois Mercier  <https://orcid.org/0000-0002-5685-1408>

REFERENCES

- O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics*. 1990;46:33-48.
- Gödel P, Shimabukuro-Vornhagen A, Bergwelt-Baildon M. Understanding cytokine release syndrome. *Intens Care Med*. 2018;44:371-373.
- Lee DW, Gardner R, Porter DL, et al. Current concepts in the diagnosis and management of cytokine release syndrome. *Blood*. 2014;124:188-195.
- Foran AE, Nadel HR, Lee AF, Savage KJ, Deyell RJ. Nivolumab in the treatment of refractory pediatric Hodgkin lymphoma. *J Pediatr Hematol/Oncol*. 2017;39:e263-e266.
- Rassy EE, Assi T, Rizkallah J, Kattan J. Diffuse edema suggestive of cytokine release syndrome in a metastatic lung carcinoma patient treated with pembrolizumab. *Immunotherapy*. 2017;9:309-311.
- Rotz SJ, Leino D, SS, Mangino JL, Turpin BK, Pressey JG. Severe cytokine release syndrome in a patient receiving PD-1-directed therapy. *Pediatr Blood Cancer*. 2017;64:e26642.
- Zhao L, Yang Y, Li W, Li T, Gao Q. Nivolumab-induced cytokine-release syndrome in relapsed/refractory Hodgkin's lymphoma: a case report and literature review. *Immunotherapy*. 2018;10:913-917.
- Dimitriou F, Matter AV, Mangana J, et al. Cytokine release syndrome during sequential treatment with immune checkpoint inhibitors and kinase inhibitors for metastatic melanoma. *J Immunother*. 2019;42:29-32.
- Honjo O, Kubo T, Sugaya F, Nishizaka T. Severe cytokine release syndrome resulting in purpura fulminans despite successful response to nivolumab therapy in a patient with pleomorphic carcinoma of the lung: a case report. *J Immuno Ther Cancer*. 2019;7:97.
- Kogure Y, Ishii Y, Oki M. Cytokine release syndrome with pseudoprogression in a patient with advanced non-small cell lung cancer treated with pembrolizumab. *J Thor Oncol*. 2019;14:e55-e57.
- Oda H, Ishihara M, Miyahara Y, et al. First case of cytokine release syndrome after nivolumab for gastric cancer. *Case Rep Oncol*. 2019;12:147-156.
- Stein AM, Grupp SA, Levine JE, et al. Tisagenlecleucel model-based cellular kinetic analysis of chimeric antigen receptor-T cells. *CPT Pharmacometr Syst Pharmacol*. 2019;8:285-295.
- Jiang X, Chen X, Jaiprasart P, Carpenter TJ, Zhou R, Wang W. Development of a minimal physiologically-based pharmacokinetic/pharmacodynamic model to characterize target cell depletion and cytokine release for T cell-redirecting bispecific agents in humans. *Eur J Pharm Sci*. 2020;146:105260.
- Frey NV, Shaw PA, Hexner EO, et al. Optimizing chimeric antigen receptor T-cell therapy for adults with acute lymphoblastic leukemia. *J Clin Oncol*. 2020;38:415-422.
- Bartlett NL, Sehn LH, Assouline SE, et al. Managing cytokine release syndrome (CRS) and neurotoxicity with step-up dosing of mosunetuzumab in relapsed/refractory (R/R) B-cell non-Hodgkin lymphoma (NHL). *J Clin Oncol*. 2019;37:7518.
- Colin P, Micallef S, Delattre M, Mancini P, Parent E. Towards using a full spectrum of early clinical trial data: a retrospective analysis to compare potential longitudinal categorical models for molecular targeted therapies in oncology. *Stat Med*. 2015;34:2999-3016.

17. Fernandes LL, Taylor JMG, Murray S. Adaptive phase I clinical trial design using Markov models for conditional probability of toxicity. *J Biopharm Stat.* 2016;26:475-498.
18. Conaway MR, Petroni GR. The impact of early-phase trial design in the drug development process. *Clin Cancer Res.* 2019;25:819-827.
19. Braun TM, Yuan Z, Thall PF. Determining a maximum tolerated schedule of a cytotoxic agent. *Biometrics.* 2005;61:335-343.
20. Braun TM, Thall PF, Nguyen H, Lima M. Simultaneously optimizing dose and schedule of a new cytotoxic agent. *Clin Trials.* 2007;4:113-124.
21. Cheung YK, Chappell R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics.* 2000;56:1177-1182.
22. 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.
23. Gerard E, Zohar S, Thai H-T, Lorenzato C, Riviere M-K, Ursino M. Bayesian dose regimen assessment in early phase oncology incorporating pharmacokinetics and pharmacodynamics. *Biometrics.* 2021;78:300-312. doi:10.1111/biom.13433:1-13
24. Wages NA, O'Quigley J, Conaway MR. Phase I design for completely or partially ordered treatment schedules. *Stat Med.* 2014;33:569-579.
25. Braun TM. Generalizing the TITE-CRM to adapt for early- and late-onset toxicities. *Stat Med.* 2006;25:2071-2083.
26. Braun TM. The current design of oncology phase I clinical trials: progressing from algorithms to statistical models. *Chin Clin Oncol.* 2014;3:1-11.
27. Tighiouart M, Rogatko A. Dose finding with escalation with overdose control (EWOC) in cancer clinical trials. *Stat Sci.* 2010;25:217-226.
28. Lee DW, Santomaso BD, Locke FL, et al. ASTCT consensus grading for cytokine release syndrome and neurologic toxicity associated with immune effector cells. *Biol Blood Marrow Transpl.* 2019;25:625-638.
29. O'Quigley J, Paoletti X, Maccario J. Non-parametric optimal design in dose finding studies. *Biostatistics.* 2002;3:51-56.

SUPPORTING INFORMATION

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

How to cite this article: Braun TM, Mercier F. Extending the Continual Reassessment Method to accommodate step-up dosing in Phase I trials. *Statistics in Medicine.* 2022;41(20):3975-3990. doi: 10.1002/sim.9487

APPENDIX

A.1 Deriving conditional probabilities

To promote readability, all derivations below omit conditioning upon doses.

For the second administration, we have

$$\begin{aligned}\pi_{i2} &= Pr(Y_{i2} = 1) = Pr(Y_{i2} = 1 \mid Y_{i1} = 0) \times Pr(Y_{i1} = 0) \\ &\quad + Pr(Y_{i2} = 1 \mid Y_{i1} = 1) \times Pr(Y_{i1} = 1) \\ &= Pr(Y_{i2} = 1 \mid Y_{i1} = 0) \times (1 - \pi_{i1}) + 1 \times \pi_{i1}.\end{aligned}$$

Solving for $Pr(Y_{i2} = 1 \mid Y_{i1} = 0)$, we have

$$Pr(Y_{i2} = 1 \mid Y_{i1} = 0) = \frac{\pi_{i2} - \pi_{i1}}{1 - \pi_{i1}}.$$

This conditional probability also leads to the joint probability

$$\begin{aligned}Pr(Y_{i2} = 1, Y_{i1} = 0) &= Pr(Y_{i2} = 1 \mid Y_{i1} = 0) \times [1 - Pr(Y_{i1} = 1)] \\ &= \pi_{i2} - \pi_{i1},\end{aligned}$$

which will be used in the following derivation.

For the third administration, we have

$$\begin{aligned}Pr(Y_{i3} = 1) &= Pr(Y_{i3} = 1 \mid Y_{i2} = 1, Y_{i1} = 0) \times Pr(Y_{i2} = 1, Y_{i1} = 0) \\ &\quad + Pr(Y_{i3} = 1 \mid Y_{i2} = 0, Y_{i1} = 0) \times Pr(Y_{i2} = 0, Y_{i1} = 0)\end{aligned}$$

$$\begin{aligned}
& + Pr(Y_{i3} = 1 \mid Y_{i2} = 1, Y_{i1} = 1) \times Pr(Y_{i2} = 1, Y_{i1} = 1) \\
& + Pr(Y_{i3} = 1 \mid Y_{i2} = 0, Y_{i1} = 1) \times Pr(Y_{i2} = 1, Y_{i1} = 1) \\
& = Pr(Y_{i2} = 1, Y_{i1} = 0) \\
& + Pr(Y_{i3} = 1 \mid Y_{i2} = 0, Y_{i1} = 0) \times Pr(Y_{i2} = 0, Y_{i1} = 0) \\
& + Pr(Y_{i2} = 1 \mid Y_{i1} = 1) \times Pr(Y_{i1} = 1) \\
& = Pr(Y_{i2} = 1, Y_{i1} = 0) \\
& + Pr(Y_{i3} = 1 \mid Y_{i2} = 0, Y_{i1} = 0) \times Pr(Y_{i2} = 0, Y_{i1} = 0) \\
& + Pr(Y_{i1} = 1),
\end{aligned}$$

because $Pr(Y_{i3} = 1 \mid Y_{i2} = 1, Y_{i1} = 0)$, $Pr(Y_{i3} = 1 \mid Y_{i2} = 1, Y_{i1} = 1)$, and $Pr(Y_{i2} = 1 \mid Y_{i1} = 1)$ are each exactly equal to 1, and $Pr(Y_{i3} = 1 \mid Y_{i2} = 0, Y_{i1} = 1)$ is exactly equal to 0.

Solving for $Pr(Y_{i3} = 1 \mid Y_{i2} = 0, Y_{i1} = 0)$, we have

$$\begin{aligned}
Pr(Y_{i3} = 1 \mid Y_{i2} = 0, Y_{i1} = 0) &= \frac{Pr(Y_{i3} = 1) - Pr(Y_{i2} = 1, Y_{i1} = 0) - Pr(Y_{i1} = 1)}{Pr(Y_{i2} = 0, Y_{i1} = 0)} \\
&= \frac{\pi_{i3} - (\pi_{i2} - \pi_{i1}) - \pi_{i1}}{Pr(Y_{i2} = 0 \mid Y_{i1} = 0) \times Pr(Y_{i1} = 0)} \\
&= \frac{\pi_{i3} - \pi_{i2}}{[1 - Pr(Y_{i2} = 1 \mid Y_{i1} = 0)] \times [1 - Pr(Y_{i1} = 1)]} \\
&= \frac{\pi_{i3} - \pi_{i2}}{[(1 - \pi_{i2}) / (1 - \pi_{i1})] \times (1 - \pi_{i1})} \\
&= \frac{\pi_{i3} - \pi_{i2}}{1 - \pi_{i2}}.
\end{aligned}$$

A.2 Deriving prior mean for model parameter β

We start with $f(\beta) = \pi_{j1} = d_{j1}^{\exp(\beta)}$. Given first and second derivatives

$$\begin{aligned}
f'(\beta) &= \frac{df}{d\beta} = \log(d_{j1}) \exp(\beta) f(\beta) \\
f''(\beta) &= \frac{d^2f}{d\beta^2} = \log(d_{j1}) \exp(\beta) [f'(\beta) + f(\beta)] \\
&= \log(d_{j1}) \exp(\beta) f(\beta) [\log(d_{j1}) \exp(\beta) + 1]
\end{aligned}$$

we use a second order Taylor expansion of $f(\beta)$ around the prior mean μ_1 to derive

$$\begin{aligned}
E[f(\beta)] &= f(\mu_1) \\
&+ 0.5\sigma^2 \log(d_{j1}) \exp(\mu_1) f(\mu_1) [\log(d_{j1}) \exp(\mu_1) + 1].
\end{aligned}$$

We seek to make the second-order term equal to zero, which occurs when $\log(d_{j1}) \exp(\mu_1) = -1$. Solving for μ_1 , we find the suggested prior mean $\mu_1 = \log[-\log(d_{j1})]$. To collapse over all schedules, we compute the arithmetic average, so that we have

$$\mu_1 = \frac{1}{6} \sum_{j=1}^6 \log[-\log(d_{j1})].$$