A Dose Escalation Design Considering Different Toxicity Onset Time Frames Using a Bayesian Piecewise Proportional Hazard Model

Hongjie Zhu, Ph.D.[1]
[1]BeiGene, 1228 Yanan Middle Road, Jing An Kerry Centre, 20F Tower 3, Shanghai, China 200020

Abstract
When designing the dose escalation trials in oncology, it is recognized that the dose limiting toxicity (DLT) events may have different onset time frames. Focusing only on the events that are likely to occur within a short time period may result in missing some important late-onset events in the dose escalation consideration and thus under-estimating the overall toxicity level of each dose. On the other hand, suspending patient recruitment for a longer DLT observation window delays the dose finding stage of the drug development, especially when numerous dose steps are evaluated. Many existing dose escalation methods have been extended to accommodate the study designs with late-onset DLT events, with which the dose for a new patient or cohort can be recommended even before all existing patients finish their entire DLT observation window. However, these methods are not designed to integrate multiple categories of DLTs defined with different observation windows given their likely onset time frames. We propose to use a time-to-event Bayesian piecewise proportional hazard (TITE-BPPH) model to handle the problem. We provide prior distribution specifications and the overall study design with the dose escalation rules derived based on the model inference. Simulation results are presented to demonstrate the operating characteristics of the method and compare it to the classical 3+3 with DLT status of all patients fully resolved before recruiting a new cohort as a benchmark, and also to an escalation with overdose control proportional hazard (EWOC-PH) method assuming constant hazard. The two model-based approaches reduce the trial duration and usually identify the maximum tolerated dose more accurately than the 3+3. Also the TITE-BPPH method exhibits more consistent performance than the EWOC-PH method among different compositions of the early- and late-onset events and compares favorably with the later method in several scenarios.

Key Words: Dose escalation, oncology, various DLT onset time frames, late-onset events, time-to-event data, piecewise proportional hazard model

1. Introduction

First-in-Human (FIH) studies play an important role in oncology drug development by identifying the proper dose(s) for later phase investigations. To achieve this objective, a dose escalation is usually pursued. Traditionally, the development of cytotoxic therapies focuses on identifying the maximum tolerated dose (MTD) based on the safety profile of different doses of the drug, assuming monotonic relationships of the anti-cancer effect and toxicity level with the doses. In this case, the dose limiting toxicity (DLT) events defined in a predetermined DLT observation window are usually considered for making the dose escalation/de-escalation decisions, although the overall safety data should be examined as well. With the emergence of targeted and immuno- therapies, the drug activities reflected in pharmacodynamic parameters and/or (short-term) efficacy endpoints become important considerations in dose finding, on top of the toxicity signals. However, the safety profile of the drug remains critical in those studies.

Regardless of the therapeutic classes and dose escalation methods, protecting patients' safety in avoiding potential over-dosing is one of the top priorities of these trials. On the other hand, it is important to treat those cancer patients with an effective dose for ethical considerations. Balancing the two criteria along with a commonly limited sample size constitute unique challenges for the FIH studies in oncology.

When identifying the MTD, the 3+3 method was widely used traditionally, although it has several major limitations that have been reported and acknowledged in the literature (Thall and Lee, 2003; Iasonos et al, 2008): 1) it lacks flexibility in handling different cohort sizes and/or targeted DLT probabilities; 2) it tends to treat larger portion of patients at sub-therapeutic doses; and 3) importantly, it usually provides inferior performance due to the lack of solid statistical considerations in the dose escalation rules. Particularly, to illustrate the last point, assume the numbers of DLT events x_d of the d-th dose follow a binomial distribution with the probability mass function $b(x_d; N_c, \theta_d)$, where N_c is the cohort size, θ_d is the corresponding DLT probability, $d = 1, \dots, D$, and D is the total number of candidate dose steps in a pre-specified dose escalation path. Let d_{MTD} represents the index of the true MTD. Then the probability for the 3+3 algorithm to identify the true MTD is $\prod_{d=1}^{d_{MTD}} \Big(b(0; N_c = 3, \theta_d) \cdot \big(1 + b(1; N_c = 3, \theta_d)\big) \Big) \cdot \Big(1 - b\big(0; N_c = 3, \theta_{d_{MTD}+1}\big) \cdot \big(1 + b\big(1; N_c = 3, \theta_{d_{MTD}+1}\big) \big) \Big)$, for any $d_{MTD} < D$, which is capped by the probability of the true MTD fulfilling the 3+3 escalation rule, $b\big(0; N_c = 3, \theta_{d_{MTD}}\big) \cdot \Big(1 + b\big(1; N_c = 3, \theta_{d_{MTD}}\big) \Big)$, i.e., between 0.42 and 0.78 for a $\theta_{d_{MTD}}$ between 1/3 and 1/6. And it can be seen that a flatter dose-toxicity curve would render the task even more challenging for the 3+3 algorithm.

New methodologies with rigorous statistical inference were proposed aiming at improving the MTD identification accuracy while providing more flexibilities to handle various scenarios in real-life clinical trials. One class of them are model-based methods, such as the continual reassessment method (CRM, O'Quigley,et al, 1990), the escalation with overdose control method (EWOC; Babb et al, 1998), and the Bayesian logistic regression method (Neuenschwander, Branson, and Gsponer, 2008), which are built on a parametric dose-toxicity relationship. The continual re-assessment methods have been mentioned in several regulatory guidelines (CHMP, 2007; FDA, 2011, 2019). A previous survey (Love

et al., 2017) also suggests the methods are widely accepted and implemented by the pharmaceutical companies.

Another class of model-assisted methods were invented with the dose escalation rules also derived from statistical inference, which were Likely inspired by the need of transparency and implementation convenience. They are able to explicitly enumerate dose escalation decisions at the trial design stage for all the different data of one dose level that can be observed during the trial conduct, and thus provide excellent transparency to the clinicians. Some examples are the modified target probability interval design (mTPI; Ji et al, 2013), the Bayesian optimal internal design (BOIN, Yuan et al, 2016), the keyboard design (Yan, Mandrekar, and Yuan, 2017), and the recent i3+3 design (Liu, Wang and Ji, 2020). Note that to achieve the desired feature of enumerating all scenarios, the possible scenarios need to be limited. The model-based methods need to use data from all previous doses, making it impractical to enumerate all the combinations. The model-assisted methods are built on the assumption that the toxicities of the different doses are independent (Lin and Yuan, 2018), or they ignore such a relationship if it exists.

However, these methods all face the challenge that some important categories of toxicity events tend to occur in longer time frames. These time frames are typically longer than the traditional DLT observation window, e.g., one or two treatment cycles. In this case, it is important to apply the targeted DLT rate control to an extended DLT observation window. Otherwise, the selected dose is likely to exhibit higher than expected toxicity level in the late phase trials when the patients are followed for a longer time period, which may render an effective drug being intolerable, complicating and/or delaying the overall drug development process. The dilemma here is that the extension of the DLT observation window with the traditional dose escalation methods will also result in a delay, amount of which depends on the patient accrual rate, enrollment method, e.g., rolling or by cohort, and how many doses are to be evaluated.

In order to address the challenge while protecting patients' safety, dose escalation methods without completing all patients' DLT observation window were proposed and evaluated, some of which are direct extensions of the aforementioned methods.

A rolling six design (Skolnik et al, 2008) built on the 3+3 rules and allowing continuous patients enrollment was originally proposed for pediatric trials. The authors enumerated rules for all possible combinations of the numbers of all patients on the current dose level, observed DLT events among them, and patients not finishing the entire DLT observation window yet. Generally, the dose can be escalated/de-escalated when the corresponding 3+3 rules can be confirmed even if there are still patients with data pending; otherwise, a new patient will be recruited at the current dose until six patients are enrolled and then the recruitment will be suspended. Simulation studies show the rolling six design has similar chance of identifying the true MTD as the 3+3 and similar numbers of patients experienced a DLT, but reduces the overall study duration despite using more patients.

On the other hand, various model-based and model-assisted methods have been extended or proposed to deal with the late-onset DLTs and/or fast accrual. In general, three different ways have been used by these methods to handle the incomplete DLT observation window. First, several methods assign weights to patients. For instance, the time-to-event (TITE)-CRM method (Cheung and Chappell, 2000) weight patients by the proportion of completed DLT observation window over the duration of the entire window, which results in a weighted likelihood function. The TITE-EWOC method (Mauguen et al., 2011) adopts a

similar idea. Lin and Yuan (2019) also use a weighting method to extend the class of model-assisted methods built on the binomial distribution. They formulate the weight as the conditional probability that a patient experience a DLT in the proportion of completed DLT observation window conditional on that a DLT will occur, and propose different weight schemes, representing different assumptions. A data-driven adaptive weight is also mentioned, although the transparency advantage of the method class is likely to be lost and the scheme is not recommended by the authors. A second way is to treat the DLT events potentially occurring during the incomplete part of the DLT observation window as missing data and use an imputation method. For instance, Yuan et al., (2018) use a single mean imputation to extend BOIN to its TITE version. The imputation is based on the DLT probability of the same dose, which is estimated among patients with their DLT status fully resolved, as well as an assumed time-to-event distribution. The RED design (Ivanova, Wang, and Foster, 2016) takes a relatively conservative imputation approach and assigns fractional toxicity to the patients with data pending. Some of these methods are compared to the 3+3 and usually report a reduced study duration, improved MTD identification accuracy but with moderate increase in the overdosing probability. A third way is to directly model the time-to-DLT events, as proposed by Tighiouart, Liu, and Rogatko (2014). They use a parametric proportional hazard model, based on which they propose the EWOC-proportional hazard (EWOC-PH) method. Although the authors note that various parametric models can be used, prior specifications are only provided for two parameters, which is used for a constant hazard model and in their simulation studies.

Despite the rich arsenal of methods handling late-onset events, it remains an issue how to integrate the escalation rules based on both the early- and late-onset events, especially that there isn't much prior information regarding which category would eventually dominate the DLT events that will be observed for the new drug, or both of them could occur with non-ignorable probabilities. Note that all of the aforementioned model-based and model-assisted methods above except the EWOC-PH method inherently build their weighting or imputation method on an assumed distribution of the time-to-DLT event, determined *a priori* without using the actual time-to-event data. For instance, the commonly used weight by the proportion of completed DLT observation window assumes a uniform distribution within the observation window. When several categories of DLT events of different onset time frames need to be taken into consideration, different DLT observation windows could be specified for them, since clinically it does not make sense to count the early-onset events with an extended observation window. The investigators may want to evaluate dose for a new cohort when all the patients finish the early-onset DLT window. In this case, we argue that it is more reasonable to assume different DLT hazards at different periods of time and thus propose to use a piecewise model to take into account different categories of events. And thus we propose the TITE-piecewise proportional hazard (TITE-BPPH) dose escalation method.

The subsequent content of the paper will be organized in the following way. Section 2 will describe the proposed dose escalation method in detail including the statistical model, prior distribution specifications, the study design with the dose escalation rules. Section 3 will present simulation studies to evaluate the operating characteristics of the proposed method and compare it to another two methods: the classical 3+3 design, which always complete the longest DLT observation window before recruiting a new cohort, as a benchmark, and the EWOC-PH method assuming constant hazard. We provide conclusions and discussions in Section 4.

2. The Dose Escalation Method

We build on the idea of modeling time-to-DLT of Tighiouart, Liu, and Rogatko (2014) and extend the EWOC-PH method for the scenario of multiple DLT categories of different DLT observation windows. Suppose there are K different categories of DLTs. The categories are determined by their likelihood of onset time frames, which is usually suggested by the clinicians. In order to control model complexity, DLTs expected to have similar onset time frames should be grouped into one category. However, only the number of categories of different observation windows is used by the model. It does not need to distinguish the DLT events by categories. For each category, the corresponding observation window is pre-specified. It is assumed that the DLT status of each category should be fully resolved within the corresponding window. From another perspective, the DLT rate is only controlled with respect to the DLT observation windows for each category, which should be included in the DLT definitions.

2.1 The Piecewise Proportional Hazard Model

We consider the different categories of DLTs as competing risks and combine them together assuming a multi-state competing risk model (Kalbfleisch and Prentice, 2002).

Following the idea of Tighiouart, Liu, and Rogatko (2014), we assume the time-to-event of each category follow a proportional hazard model, except that it is "cause-specific". Given the pre-specified DLT observation window for each category, the hazard of the specific category of DLTs is 0 by design outside the corresponding window. We use the follow cause-specific model for the *k*-th DLT category:

$$h_k(t|d,\mu_{\boldsymbol{k}},\beta)=\begin{cases}h_{0,k}(t|\mu_{\boldsymbol{k}})\cdot\exp(\beta\cdot(\log(d)-\log(d_0))), & 0\le t\le\tau_k\\ 0, & t>\tau_k\end{cases},\quad k=1,..,K,$$

Where d denotes the dose, $d\in\{d_1,\dots,d_D\}$, and $[0,\ \tau_k]$ is DLT observation window for the k-th category of DLTs, $0<\tau_1<\tau_2<\cdots<\tau_K$, $h_{0,k}(t|\mu_{\boldsymbol{k}})$ represents "baseline" hazard of a minimum reference dose d_0 within the observation window with unknown parameter(s) $\mu_{\boldsymbol{k}}$, and β is an unknown parameter capturing dose effect to the DLT hazard. The d_0 usually takes a small value that is believed *a priori* to be definitely lower than the true MTD. This model setting facilitates the prior distribution specification. More details are provided in Section 2.2.

Note that we assume the same dose effect across different DLT categories, and therefore, the "all-cause" hazard model is still a proportional hazard model:

$$h(t|d) = \sum_{k=1}^{K} h_k(t|d, \mu_{\boldsymbol{k}}, \beta)$$
$$= \begin{cases} h^*_{0,1}\left(t|\mu^*_{\boldsymbol{0,1}}\right) \cdot \exp(\beta \cdot (\log(d) - \log(d_0))), & 0 \le t \le \tau_1, \\ h^*_{0,k}\left(t|\mu^*_{\boldsymbol{0,k}}\right) \cdot \exp\left(\beta \cdot (\log(d) - \log(d_0))\right), & \tau_{k-1} < t \le \tau_k, k = 2, \dots, K, \\ 0, & t > \tau_K. \end{cases} \quad (1)$$

where $h^*_{0,k}\left(t|\mu^*_{0,k}\right) = \sum_{j=1}^{k} h_{0,j}\left(t|\mu_{\boldsymbol{j}}\right)$ is the all-cause baseline hazard of d_0. Note that for the dose escalation, we are interested in the overall rate of all categories of DLTs. In other words, we are interested in the all-cause hazard rather than the cause-specific ones. Therefore, we propose to directly model the all-cause hazard. Specifically, we propose to apply a (different) constant baseline hazard for each of the time intervals $[0, \tau_1]$ and $(\tau_{k-1}, \tau_k]'$s for $k > 1$, and argue that this proposal not only addresses the specific design of multiple DLT observation windows, but also strives a balance between model simplicity and the flexibility of allowing inconstant hazard over time for the late-onset categories of DLTs. On one hand, in many cases the DLTs are only expected to occur in a portion of the patients. If we assume a mixture model of patients will or will not develop a DLT and a constant hazard for the patients that will develop DLT, the overall hazard will decrease over time. On the other hand, delayed toxicities are unlikely to occur in a short time period, e.g., within one treatment cycle, indicating a lower hazard during that time. Therefore, a constant hazard assumption may be violated in different directions. The impact of the misspecification could be more profound when the observation window is long and/or the recruitment is fast, such that by the time of recruiting new patient(s), many existing patients are still at their early observation period, resulting in over- or under-estimation of the true toxicity level. Without the need of identifying any cause-specific hazard but only focusing on an all-cause hazard for each time interval, a piecewise constant hazard model takes a data driven approach and can adapt to the different scenarios. The piecewise constant model generated from (1) is

$$h(t|d)$$
$$= \begin{cases} \mu^*_{0,1} \cdot \exp(\beta \cdot (\log(d) - \log(d_0))), & 0 \le t \le \tau_1, \\ \mu^*_{0,k} \cdot \exp\left(\beta \cdot (\log(d) - \log(d_0))\right), & \tau_{k-1} < t \le \tau_k, k = 2, \dots, K, \\ 0, & t > \tau_K. \end{cases} \quad (2)$$

With this model, only one more parameter is needed for each additional category of DLT events. For the Bayesian inference, we propose a simple way to give vague priors for the parameters when there is limited information about the DLT rates by each time point.

At any dose escalation decision point of the trial, assume there are n patients enrolled in the study. Let $T = \{t_i,\ i = 1, \dots, n\}$, where t_i is the time to a DLT event or time to censoring of Patient i, and Ð= $\{\delta_i,\ i = 1, \dots, n\}$, where δ_i equals 1 when a DLT event is

observed for Patient i at t_i and otherwise equals 0. For the ease of notation, we let $\tau_0 \equiv 0$. Then the likelihood function corresponding to (2) is

$$L\left(\beta, \mu_{0,1}^*, \ \ldots, \mu_{0,K}^* \middle| T, \text{ɒ}\right) = \prod_{i=1}^{n} {\mu_{0,k(t_i)}^*}^{\delta_i} \cdot \exp\left(-\left(\sum_{j=1}^{k(t_i)-1} \mu_{0,j}^* \cdot (\tau_j - \tau_{j-1}) \cdot \mathrm{I}(k(t_i) > 1) + \mu_{0,k(t_i)}^* \cdot (t_i - \tau_{k(t_i)-1})\right)\right)$$

Where $k(t_i)$ is the time interval that t_i falls into.

Similar as Tighiouart, Liu, and Rogatko (2014), we can re-parameterize the piecewise constant hazard model for the ease of prior specification: we will parameterize the model by the true MTD, Υ, and the probabilities ρ_k of any DLT occurring within $[0, \ \tau_k]$ for d_0:

$$\Upsilon = d_0 \cdot \left(-\frac{\log(1-\theta)}{\sum_{k=1}^{K} \mu_{0,k}^* \cdot (\tau_k - \tau_{k-1})}\right)^{\frac{1}{\beta}},$$

$$\rho_k = 1 - \exp\left(-\sum_{j=1}^{k} \mu_{0,j}^* \cdot (\tau_j - \tau_{j-1})\right) \ , \mathrm{k} = 1, \ldots, \mathrm{K}$$

Where θ is the targeted DLT rate by the end of the longest DLT observation window τ_K corresponding to the true MTD, e.g., 0.3. Let $\rho_0 \equiv 0$. Then we have,

$$\mu_{0,k}^* = -\frac{1}{\tau_k - \tau_{k-1}} \log\left(\frac{1-\rho_k}{1-\rho_{k-1}}\right), \ k = 1, \ldots, K, \text{ and}$$

$$\beta = \frac{1}{\log(\Upsilon/d_0)} \cdot log\left(\frac{\log(1-\theta)}{\log(1-\rho_K)}\right). \tag{3}$$

2.2 Prior Distributions

The prior distribution for Υ should reflect the prior belief of where the true MTD is located. The relevant information may be obtained from the pre-clinical and other relevant data that was used to determine all the candidate dose levels, including the minimum and maximum doses in the pre-specified dose escalation steps. Tighiouart, Liu, and Rogatko (2014) use a uniform distribution between the minimum and maximum candidate doses, d_1 and d_D. Then by Bayes rule, the posterior probability of $\Upsilon \geq d_\mathrm{D}$ is always 0, and the probability of $\Upsilon \leq d_1$ is always 1, no matter what data are observed during the trial. Since the dose escalation criteria are based on the posterior probabilities as specified in Section 2.3, d_D cannot be recommended and d_1 will not be deemed as overly toxic. From this perspective, it can be considered as an informative prior distribution. Depending on how d_1 and d_D are determined, we can instead determine the prior by specifying prior probabilities for d_1 and d_D to be higher than the true MTD, p_1 and p_2, based on a uniform distribution. Then

$$log(\Upsilon) \sim u(\log(d_1) - \frac{p_1}{p_2 - p_1} \log\left(\frac{d_\mathrm{D}}{d_1}\right), \log(d_1) + \frac{1-p_1}{p_2 - p_1} \log\left(\frac{d_\mathrm{D}}{d_1}\right))$$

Which will reduce to $u(\log(d_1), \log(d_\mathrm{D}))$ when $p_1 = 0$ and $p_2 = 1$.

It is then natural to set d_0 to be the lower bound of the distribution of the MTD , and use a $u(0,\theta)$ as prior distribution for ρ_K. For $\rho_1, \ldots, \rho_{K-1}$, in order to apply the constraint that $\rho_1 \leq \rho_2 \leq \cdots \leq \rho_K$, we give prior to the ratios $\frac{\rho_{k-1}}{\rho_k} \sim beta(0.5,\ 0.5), k = 2, \ldots, K$. If prior knowledge is available, other d_0 and the prior distributions of its DLT probabilities can be used. It can be seen from (3) that as long as the prior probability of d_0 larger than Y is 0 and the upper bound of the prior distribution of ρ_K is no higher than θ, the probability of $\beta > 0$ is 1, forcing a monotonically increasing dose-toxicity relationship. This could be important for a TITE design, where patients on lower dose could complete longer observation period with an DLT observed while the patients on higher dose are still under observation. The prior distributions for ρ_k's are independent from that of β.

2.3 The Model-Based Dose Escalation and the Trial Design

Based on the likelihood function in Section 2.1 and the prior distributions in Section 2.2, we could use an MCMC algorithm to sample from the corresponding posterior distribution of the unknown model parameters $\mu_{0,k}^*$'s and β, based on which we could obtain the posterior distribution of the overall DLT probability of any dose, $prob(DLT|\tau_K, d) = 1 - s(\tau_K|d)$, where s(.) is the survival function. Then the dose escalation rule is applied to the posterior distribution of $prob(DLT|\tau_K, d)$:

$$prob(\ prob(DLT|\tau_K, d) > \theta \mid T, \text{ɒ}) < a\ feasibility\ bound \qquad (4)$$

Next we describe the study design. The design we describe here recruits patients by cohorts of size N_c. It also apply to rolling enrollment by setting $N_c = 1$. A waiting period is pre-specified since the last patient is enrolled in the study before new patient(s) can be recruited. Usually the period is selected among τ_k's. For instance, when there are two categories of DLTs, namely the early- and late-onset events, we can wait till all the patients finish the early-onset observation window.

It is important to only include patients that have good compliance to the study treatment for the safety evaluation of the doses. Although the proposed method can make use of data from patients that discontinue study treatment early due to reasons other than safety, and consider them censored at the time of treatment discontinuation, we do not recommend including patients with a short treatment duration, such as those that do not even finish treatment for the early-onset event observation window.

By the time of each cohort finish the pre-specified waiting time, the BPPH model is fitted with all the data available, and Criterion (4) is evaluated for each candidate dose with respect to a feasibility bound, e.g., 0.5. The highest dose that meets the criterion will be recommended for the new cohort. The process will continue until a pre-specified maximum number of patients are met. Before recommending the final dose, the patients are observed till all the patients have finished all their DLT observation windows, and the statistical model is fitted again to recommend the MTD, under the constraint that the dose must have been evaluated for at least one cohort.

In the simulation studies, we apply a comparable rule as 3+3 to stop trial when the dose has been evaluated for two cohorts and is recommended again. In this case, the patient recruitment will be suspended and all patients will be followed to the end of their observation windows. If the same dose or another dose previously evaluated for two cohorts is recommended, then the currently recommended dose will be the identified MTD and the trial will be stopped. Otherwise, the trial will continue to recruit new patients to evaluate the recommended dose. The results show reasonable operating characteristics of such a design with a cohort size of 3 and not allowing dose skipping.

3. Simulation Studies

3.1 Simulation Settings

We examine the operating characteristics of the TITE-BPPH method in various scenarios and compared it to the 3+3 and the EWOC-PH methods.

We simulate six doses and two categories of DLT events: the early-onset events have a 4-week observation window, while a 12-week window is specified for the late-onset ones. The early- and late-onset events are simulated independently. The true early-and late-onset DLT rate for each dose are provided in Table 1. We have a targeted overall DLT rate of 0.3. The different dose-DLT rate profiles are used to evaluate the methods when different doses are the true MTD. They also deviate from the proportional hazard model used by the EWOC-PH and TITE-BPPH methods. Then in order to compare the two model-based methods that handled hazard over time differently, for each true MTD, we further simulate three scenarios, with the early-onset event rates of the true MTD being 0.01, 0.163, and 0.29, respectively, with the overall DLT rate all being fixed at 0.3. The 0.163 corresponds to the same early- and late-onset event rates. Given the pre-specified DLT rates, the time to early- and late-onset DLT events are simulated from exponential distributions. Patients are recruited with a cohort size of 3. Patients' arrival time is simulated using a homogeneous Poisson process with an average accrual rate of two patients every four weeks. The patient enrollment is suspended after the three patients have been recruited, until new patients can be recruited per the dose escalation method as specified in Section 3.1.1 below.

Although the data used for the different methods cannot be exactly the same since they may recommend different doses during the escalation process, in order to facilitate the comparison, we randomly generated beforehand simulation seeds for patients' arrival time and different seeds for the time-to-event data for a set of cohorts under each dose in each simulation. The different methods share the same set of seeds, so that the patient arrival time will be the same, and as long as the methods recommend the same doses, even if in different orders, the time-to-event data will also be the same.

3.1.1 The implementation of the dose escalation methods

The TITE-BPPH method will recruit new cohort when the last patient of the current cohort completes the 4-week early-onset event observation window. At

each decision making time point, the highest dose with its posterior median DLT probability < 0.3 will be recommended. However, in order to make a fair comparison with 3+3, we do not allow dose skipping during the escalation process. A trial stopping rule described at the end of Section 2.3 is used. The piecewise proportional hazard model assumes different hazards for the [0, 4] and the (4, 12] weeks intervals. The prior distribution for the logarithm of true MTD is specified assuming it is a uniform distribution and the lowest and highest dose steps have 0.15 and 0.85 overdosing probabilities, respectively. This result in a $u(\log(10),\ \log(1094))$. The prior distribution for a minimum reference dose of 9mg is $u(0,\ 0.3)$. The 9mg is chosen to be close to but smaller than the lower bound of the prior MTD distribution. The prior distribution for the ratio of the DLT rates by Week 4 vs. by Week 12 is $beta(0.5,\ 0.5)$.

For the 3+3 method, we always wait for every patient in the current cohort to finish the entire 12-week observation window before recruiting the new cohort. Therefore, it usually has access to more data compared to the other two methods during the dose escalation process.

The EWOC-PH method adopts the same dose escalation and trial stopping rules, and similar prior settings as the TITE-BPPH method, except that a constant overall hazard is assumed during the entire 12-week observation window, and therefore there is no need to specify prior distributions for the ratio of the DLT rates by the different time points.

3.1.2 The operating characteristics

The following operating characteristics are examined to evaluate and compare the different methods:

- Average study duration (in weeks),
- Average study duration when the true MTD was found (in weeks),
- Average number of patients recruited,
- Average number of patients recruited when the true MTD was found,
- Percentage of identifying the true MTD,
- Percentage of recommending a dose higher than the true MTD (overdosing),
- Percentage of patients treated with a dose higher than the true MTD during the trial,
- Percentage of recommending a dose lower than the true MTD (underdosing).

3.2 Simulation Results

Simulation results are shown in Figures 1-8 and discussed in the following sections. For each true MTD, we take the average of the operating characteristics of the 3+3 method among the three scenarios with different early-onset DLT rates, since as excepted, these results are very similar.

3.2.1 Study duration and number of patients recruited

Figure 1 shows the average study duration. It can be clearly seen that the overall study duration is significantly reduced by the EWOC-PH and TITE-BPPH methods, especially for the Scenarios S7-S15, where the true MTD is at the higher dose steps. The average

reduction in study duration ranges from 10 to 50 weeks. The study durations using the EWOC-PH and TITE-BPPH methods are very similar. The study duration among the scenarios of different early-onset DLT rates but with the same MTD is also similar. We further look at the average study duration when the true MTD was found (Figure 2). Similar trend is observed, although the differences between 3+3 and the other two methods are slighted larger.

Although the EWOC-PH and TITE-BPPH methods reduce the overall study duration, they do tend to recruit more patients than the 3+3 (Figure 3). It is partially due to the fact that they have lower underdosing probabilities than the 3+3 (Figure 8). When we look the at the numbers of patients when the true MTD is found, the numbers are closer among the methods (Figure 4). Between EWOC-PH and TITE-BPPH, the numbers are very similar.

3.2.2 The MTD identification accuracy

Figure 5 clearly shows that the EWOC-PH and TITE-BPPH methods provide higher accuracy in identifying the true MTD than the 3+3 for all the simulation scenarios. However, on average, the improvement of the EWOC-PH method over the 3+3 among the scenarios with different true MTDs is reduced from 19% to 16% to 12% when the early-onset DLT rate increases from 0.01 to 0.163 to 0.29, while remarkably, the average improvement of the TITE-BPPH is always around 19% across different early-onset DLT rates, which demonstrates the superiority of the proposed method with the flexibility of dealing with different combinations of early- and late-onset DLT rates. The performance of the TITE-BPPH is consistent across different early-onset DLT rates, and its superiority over the EWOC-PH method appear to be more obvious when the advantage of both of them over 3+3 is more profound.

3.2.3 Overdosing and underdosing probabilities

Figures 6 and 7 indicate that both of the model-based methods increase the likelihood of overdosing. Excluding Scenarios 12-15, where the true MTD is the highest dose, the average overdosing percentage of the EWOC-PH method is 7.3% higher than that of the 3+3 method and the number is 9.6% for the TITE-BPPH method. Note that on average the percentage of patients treated with a dose larger than the true MTD during the trial is only slighted increased by the two methods (1.8% and 3.6%, respectively). The increased overdosing probability compared to the 3+3 is also seen in other TITE methods (Yuan et al., 2018; Lin and Yuan, 2019). It is argued that the overdosing probability should be balanced with precision for identifying the MTD and the chance of underdosing. From Figure 8, on average the EWOC-PH method reduces the percentage of underdosing by 21% from the 3+3, and TITE-BPPH reduces 27%. Clearly the overall MTD misspecification percentage is lowest for the TITE-BPPH method, followed by the EWOC-PH and 3+3 methods.

Table 1: The True DLT Rate Profile of the Simulation Scenarios

True MTD (mg)	*Scenario*	*Rate of early-onset DLT events of the true MTD at Week 4*	*Rate of late-onset DLT events of the true MTD at Week 12*	*Overall DLT rate of the true MTD at Week 12*	*Early-onset DLT event rate by Week 4*						*Overall DLT event rate by Week 12*					
					20	*60*	*120*	*240*	*360*	*540*	*20*	*60*	*120*	*240*	*360*	*540*
	S1	0.01	0.29	0.3	<0.01	0.01	0.02	0.03	0.05	0.06	0.12	0.3	0.5	0.7	0.85	0.9
60	S2	0.163	0.163	0.3	0.06	0.16	0.29	0.45	0.61	0.68	0.12	0.3	0.5	0.7	0.85	0.9
	S3	0.29	0.01	0.3	0.12	0.29	0.49	0.69	0.84	0.89	0.12	0.3	0.5	0.7	0.85	0.9
	S4	0.01	0.29	0.3	<0.01	<0.01	0.01	0.02	0.03	0.05	0.05	0.16	0.3	0.52	0.67	0.81
120	S5	0.163	0.163	0.3	0.03	0.08	0.16	0.31	0.42	0.56	0.05	0.16	0.3	0.52	0.67	0.81
	S6	0.29	0.01	0.3	0.05	0.15	0.29	0.51	0.66	0.80	0.05	0.16	0.3	0.52	0.67	0.81
	S7	0.01	0.29	0.3	<0.01	<0.01	<0.01	0.01	0.02	0.03	0.02	0.07	0.14	0.3	0.45	0.6
240	S8	0.163	0.163	0.3	0.01	0.04	0.07	0.16	0.26	0.37	0.02	0.07	0.14	0.3	0.45	0.6
	S9	0.29	0.01	0.3	0.02	0.07	0.13	0.29	0.44	0.59	0.02	0.07	0.14	0.3	0.45	0.6
	S10	0.01	0.29	0.3	<0.01	<0.01	<0.01	<0.01	0.01	0.02	0.01	0.02	0.05	0.16	0.3	0.5
360	S11	0.163	0.163	0.3	<0.01	0.01	0.03	0.08	0.16	0.29	0.01	0.02	0.05	0.16	0.3	0.5
	S12	0.29	0.01	0.3	<0.01	0.02	0.05	0.15	0.29	0.49	0.01	0.02	0.05	0.16	0.3	0.5
	S13	0.01	0.29	0.3	<0.01	<0.01	<0.01	<0.01	<0.01	0.01	0.01	0.02	0.03	0.04	0.1	0.3
540	S14	0.163	0.163	0.3	<0.01	0.01	0.02	0.02	0.05	0.16	0.01	0.02	0.03	0.04	0.1	0.3
	S15	0.29	0.01	0.3	<0.01	0.02	0.03	0.04	0.10	0.29	0.01	0.02	0.03	0.04	0.1	0.3

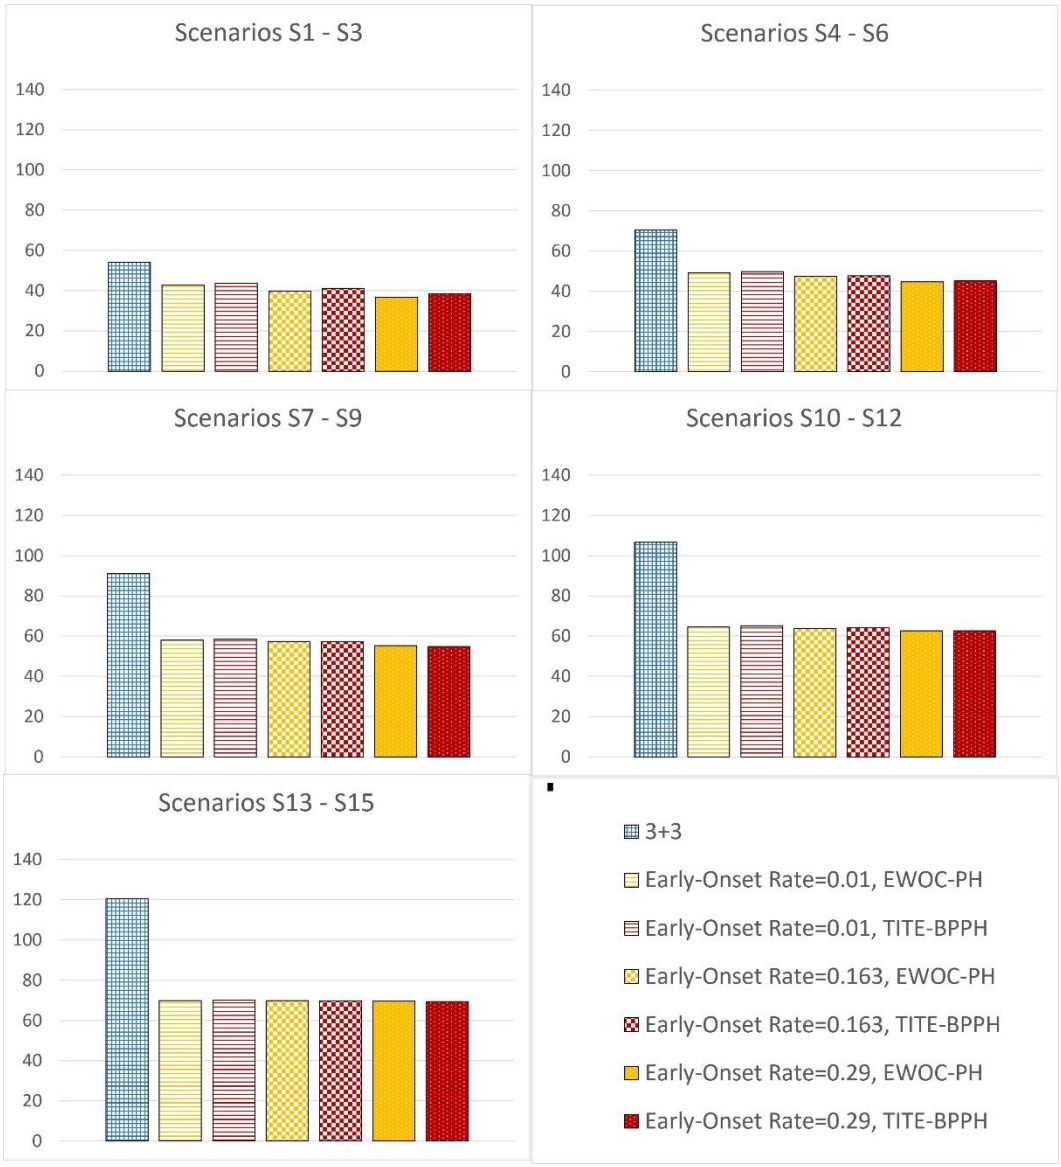

Figure 1: Simulation results: average study duration (in weeks)

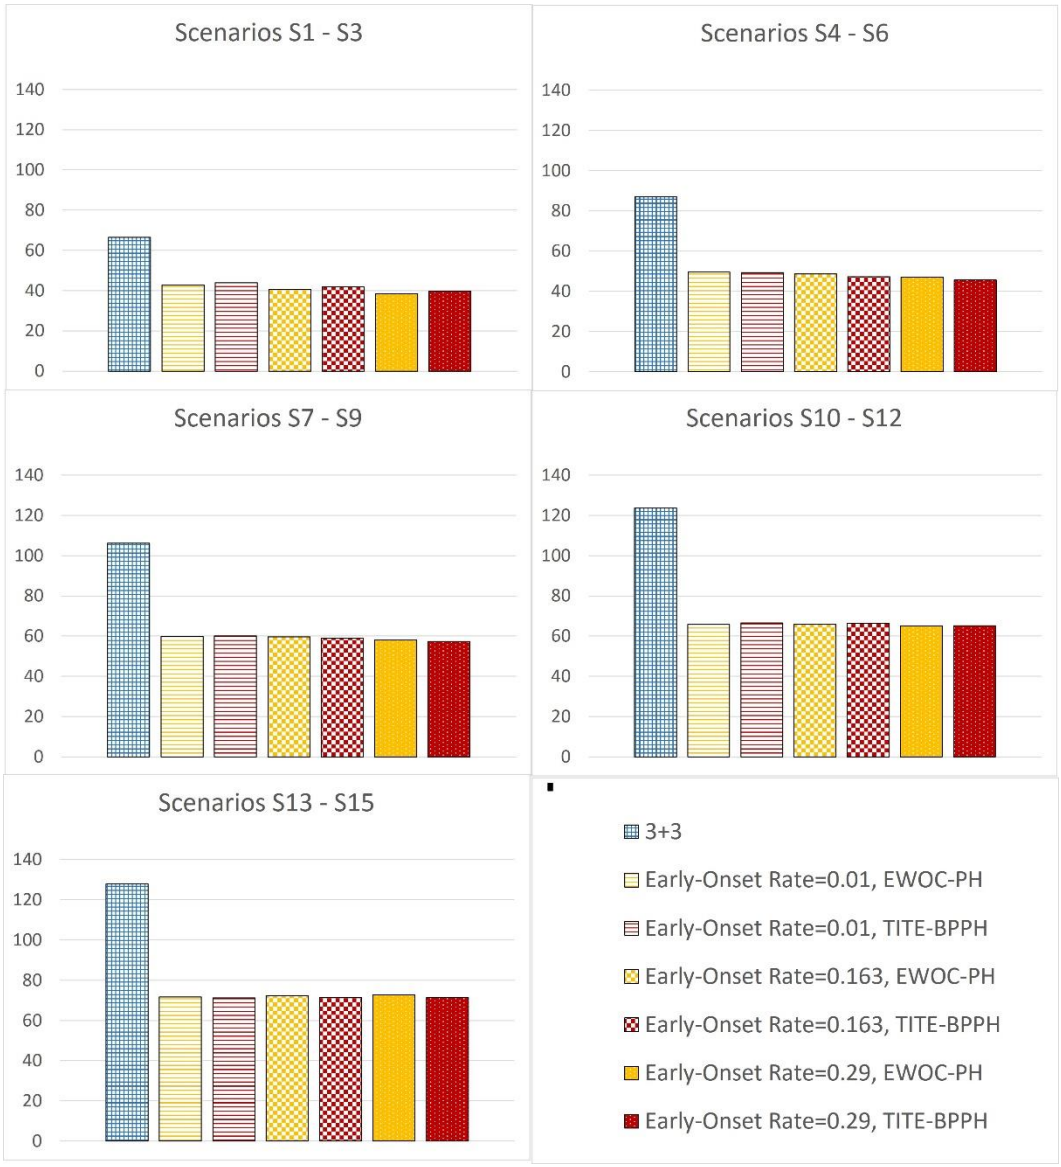

Figure 2: Simulation results: Average study duration when the true MTD was found (in weeks)

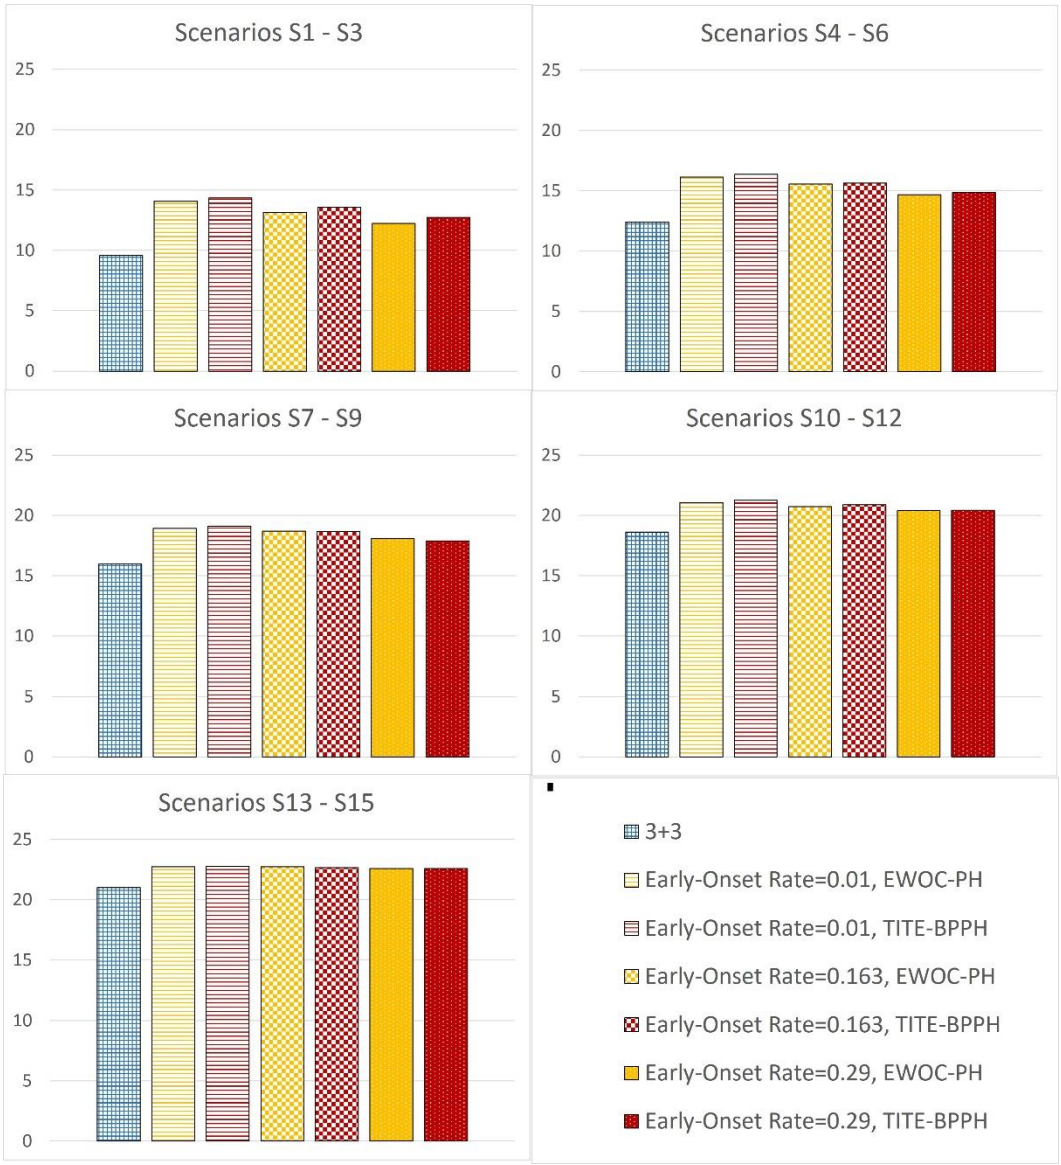

Figure 3: Simulation results: Average number of patients recruited

Figure 4: Simulation results: Average number of patients recruited when the true MTD was found

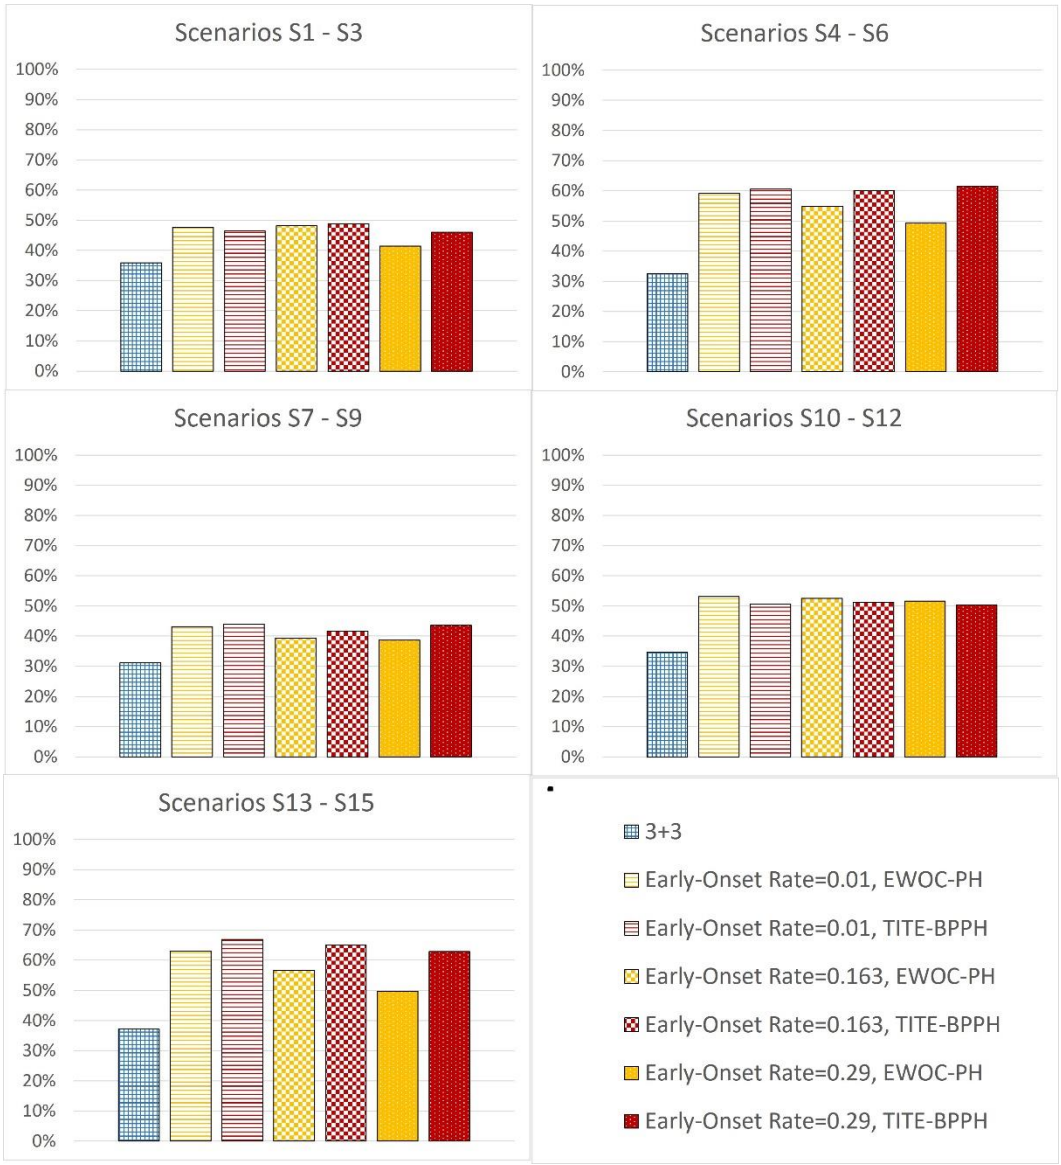

Figure 5: Simulation results: Percentage of identifying the true MTD

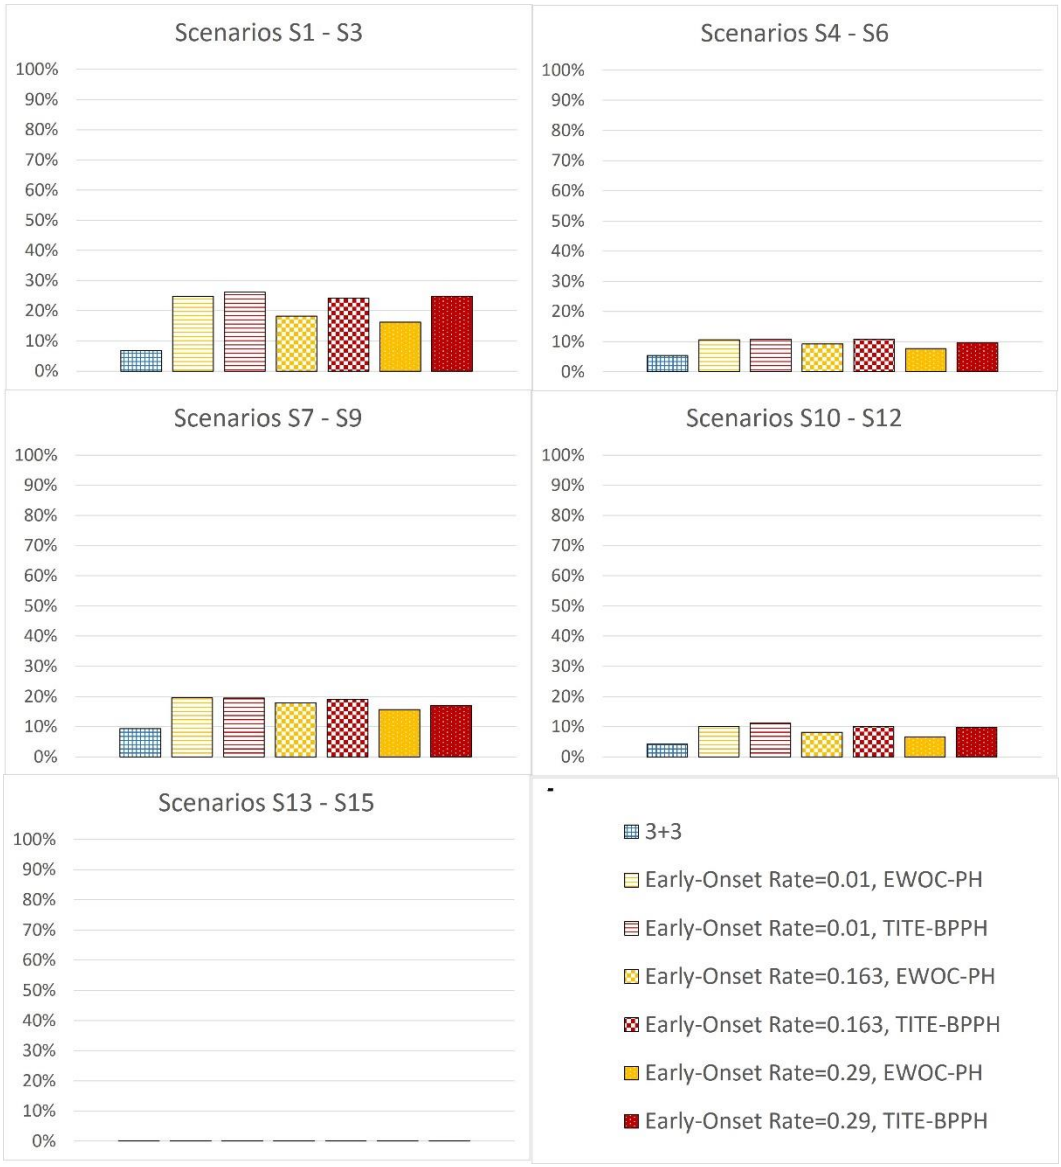

Figure 6: Simulation results: Percentage of recommending a dose higher than the true MTD (overdosing)

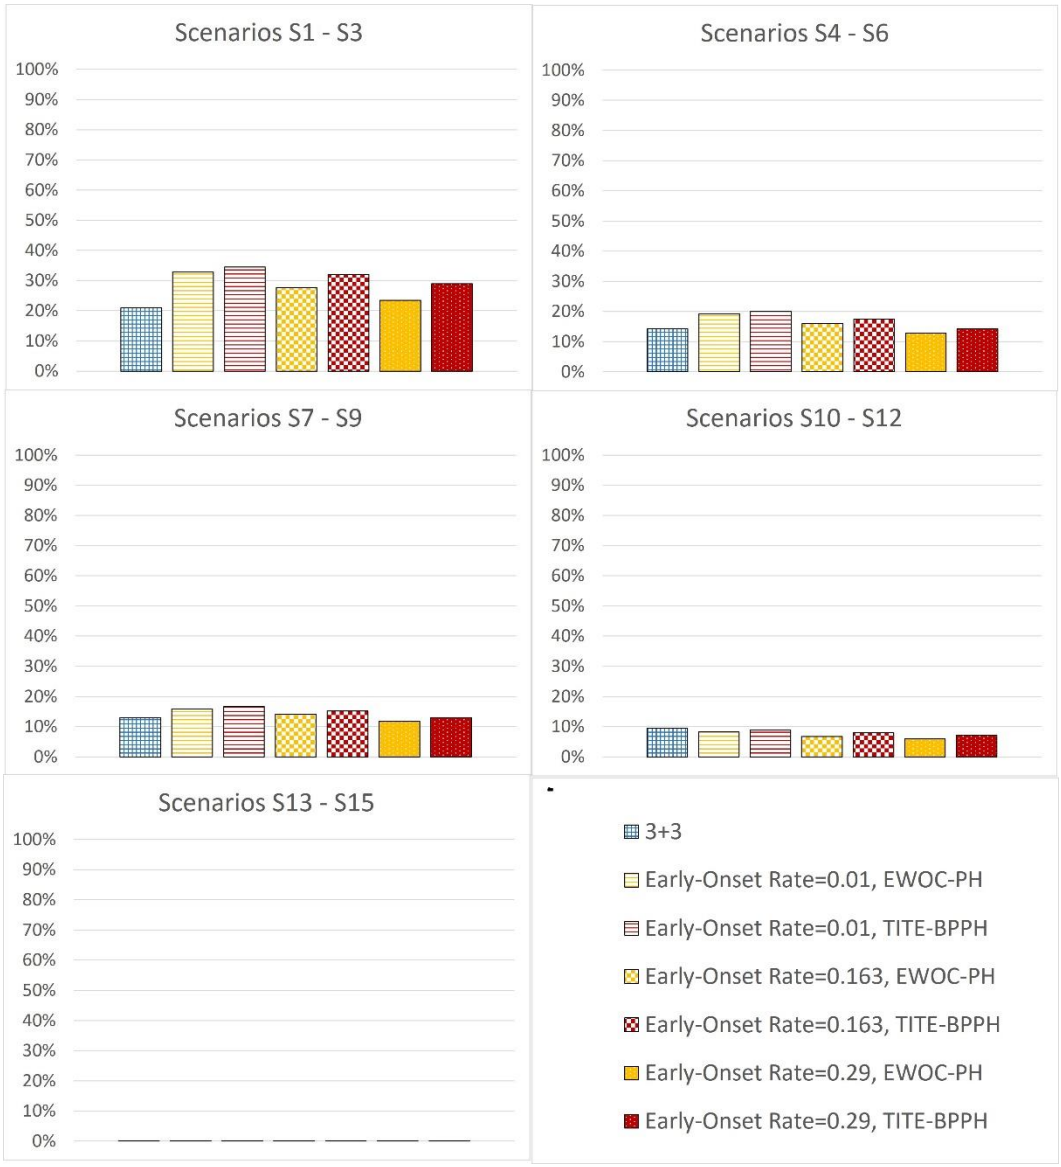

Figure 7: Simulation results: Percentage of patients treated with a dose higher than the true MTD during the trial

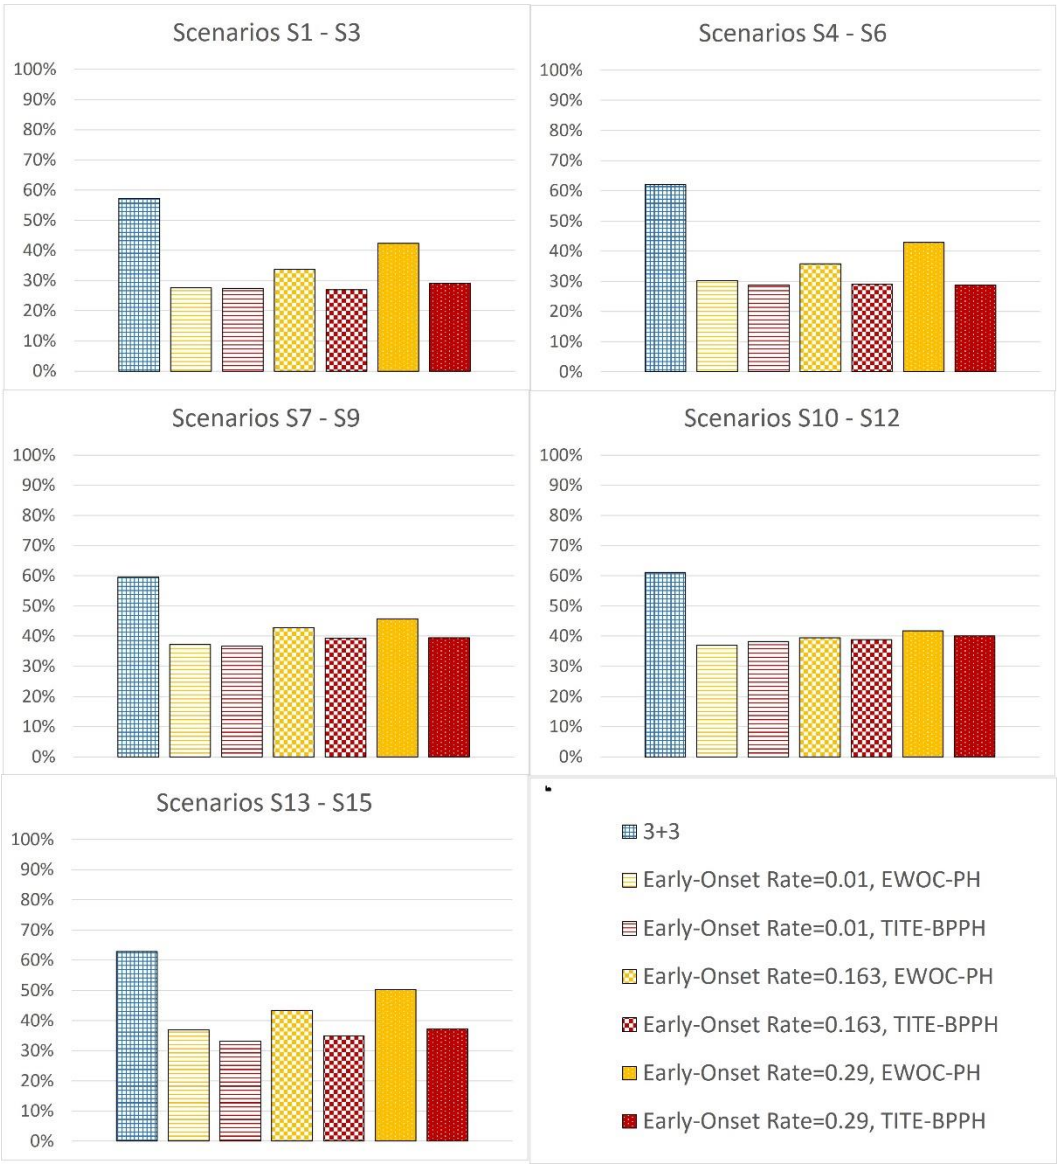

Figure 8: Simulation results: Percentage of recommending a dose lower than the true MTD (underdosing)

4. Conclusion and Discussion

In this paper, we propose a TITE-BPPH method with the essence of using a Bayesian piecewise proportional hazard model to support dose escalation decision when there are multiple categories of DLTs with different observation windows. In simulations studies, the proposed method demonstrates consistent performance across scenarios of different early- and late-onset DLT rates. It always outperforms the 3+3 in identifying the true MTD and also tends to perform better than a EWOC-PH method assumed constant hazard over time, when the overall DLT hazard over time deviates from the constant hazard assumption. Both of the methods reduce the overall study duration compared to a classical 3+3 method that requires all patients to complete the DLT observation window before

recruiting new patients. Both methods tend to increase the overdosing probability, which can be better controlled by adopting a more stringent feasibility bound in (4) than 0.5, if overdosing outweighs underdosing when balancing the two risks.

As mentioned in the introduction, various model-assisted methods are readily available to handle late-onset toxicities, although the methods are not designed to handle the specific situation of combining the different DLT categories without prior knowledge of their relative rates. Also we note that for these methods with some patients not completing the entire DLT observation window, although all dose escalation and de-escalation scenarios can still can laid out, they become more complicated, as the missing observation period needs to be taken into the decision rules. The underlying statistical models may still need to be explained in order to fully appreciate the rules, which, otherwise, may be regarded to be arbitrary. The proposed method falls into the model-based category, for which it is usually impractical to tabulate all scenarios especially with the TITE-type design. On the other hand, it fit a dose-toxicity curve, one benefit of which is that it can be used to suggest if an optional dose should be explored between two standard doses.

The proposed design is flexible in the waiting time. It can be made variable at different escalation steps and extended when a given dose on the escalation path start to be evaluated. This method can be used when the first doses in the escalation path are considered relatively safe.

Finally, we note that the method can be naturally extended for combo-therapy dose escalation, which will be described in a separate article.

References

Cheung YK, Chappell R (2000) Sequential designs for phase I clinical trials with late-onset toxicities. Biometrics 56: 1177–1182.

CHMP (2007) Guideline on clinical trials in small populations. https://www.ema.europa.eu/en/clinical-trials-small-populations#current-effective-version-section

FDA (2011) Guidance for Industry: Clinical considerations for therapeutic cancer vaccines. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-considerations-therapeutic-cancer-vaccines

FDA (2019) Guidance for Industry: Adaptive designs for clinical trials of drugs and biologics. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/adaptive-design-clinical-trials-drugs-and-biologics-guidance-industry

Iasonos A, Wilton AS, Riedel ER, Seshan VE, Spriggs DR.(2008) A comprehensive comparison of the continual reassessment method to the standard 3 + 3 dose escalation scheme in Phase I dose-finding studies. Clin Trials. 5(5):465-77.

Ivanova A, Wang Y, and Foster MC (2016) The rapid enrollment design for phase I clinical trials. Stat Med. 35(15): 2516–2524.

Kalbfleisch, J and Prentice, R (2002), The Statistical Analysis of Failure Time Data. Second Edition, Wiley, Hoboken.

Lin R and Yuan Y (2019) Time-to-event model-assisted designs for dose-finding trials with delayed toxicity. Biostatistics doi: 10.1093/biostatistics/kxz007.

Liu, M, Wang, SJ and Ji, Y (2020). The i3+ 3 Design for Phase I Clinical Trials. J Biopharm Stat. 30(2):294-304 .

Love SB, et al. (2017) Embracing model-based designs for dose-finding trials. Br J Cancer. 117(3):332-339.

Mauguen A, Le Deley MC, Zohar S (2011) Dose-finding approach for dose escalation with overdose control considering incomplete observations. Stat Med 30: 1584–1594.

Neuenschwander, Branson, Gsponer (2008). Critical aspects of the Bayesian approach to phase I cancer trials. Statistics In Medicine, 27:2420-2439

Normolle D and Lawrence T. (2006) Designing Dose-Escalation Trials With Late-Onset Toxicities Using the Time-to-Event Continual Reassessment Method. J Clin Oncol. 24(27):4426-33

O'Quigley J, Pepe M, Fisher L (1990) Continual Reassessment Method: A Practical Design for Phase 1 Clinical Trials in Cancer. Biometrics, 46(1), 33–48

Sharma RA, et al. (2016) Clinical development of new drug–radiotherapy combinations. Nat Rev Clin Oncol. (10):627-42

Skolnik JM, Barrett JS, Jayaraman B, Patel D, Adamson PC (2008) Shortening the timeline of pediatric phase I trials: the rolling six design. J Clin Oncol 26:190–5.

Tighiouart M, Liu Y, Rogatko A. (2014) Escalation with overdose control using time to toxicity for cancer phase I clinical trials. PLoS One. 9(3):e93070.

Thall PF, Lee SJ. (2003) Practical model-based dose-finding in phase I clinical trials: methods based on toxicity. Int J Gynecol Cancer. 13(3):251-61.

Yan, F., Mandrekar, S. J., and Yuan, Y. (2017). Keyboard: a novel Bayesian toxicity probability interval design for phase I clinical trials. Clinical Cancer Research 23, 3994-4003.

Yuan Y, Lin R, Li D., Nie L. and Warren KE (2018) Time-to-Event Bayesian Optimal Interval Design to Accelerate Phase I Trials. Clin Cancer Res 24(20):4921-4930.