R-TPI: Rolling Toxicity Probability Interval Design to Shorten the Duration and Maintain Safety of Phase I Trials

Wentian Guo[1], Yuan Ji [1,2,*], Daniel Li[3,*]

1. Laiya Consulting, Inc., Wilmette, IL, USA

2. NorthShore University HealthSystem, Evanston, IL, USA

3. Juno Therapeutics, Seattle, WA, USA

*: correspondence made to koaeraser@gmail.com, daniel.li@junotherapeutics.com

November 21, 2018

Abstract

Purpose To propose a general statistical design to shorten trial duration and improve safety of phase I trials.

Methods We propose R-TPI, a rolling enrollment design that combines the features in model-based designs such as mTPI-2 (Guo et al., 2017) and rule-based designs such as rolling six (Skolnik et al., 2008). R-TPI employs a novel rolling enrollment scheme, which allows concurrent patient enrollment that is faster than cohort-based enrollment.

Results Bench-marking against rolling six, we found that the R-TPI design is as fast in completing clinical trials but with fewer toxicity events and higher chance of finding the MTD in the single scenario laid out in Skolnik et al. (2008). We also found that in a broad setting involving multiple scenarios, R-TPI was generally faster, safer, and more reliable than standard designs.

Conclusion R-TPI is a general design that can be applied to adult and pediatric phase I trials. It reduces the length of trial duration, leads to safer trials with fewer toxicity events, and maintains relatively a high chance of identifying the MTD.

KEYWORDS: Dose Finding; Incomplete toxicity data; mTPI-2; Safety; Trial duration.

1 Introduction

Phase I oncology dose-finding trials assign cancer patients to ascending doses of a new investigational drug (or drug combinations) and adaptively decide the dose level of newly enrolled patients based on observed binary dose-limiting toxicity (DLT) outcomes. The goal is to determine the maximum tolerated dose (MTD) of the drug(s), defined as the highest dose that has a DLT rate less than or close to a pre-specified target rate. Popular statistical designs, such as the 3+3, CRM, and mTPI-2 typically enroll patients in cohorts and apply sequential decisions that determine the dose level for each cohort based on observed toxicity data. In many cases the cohort size, defined as the number of enrolled patients per cohort, is an integer around 3. Accrual is suspended after enrollment of each cohort of patients until all the patients in the current cohort have observed outcomes, with or without DLTs. This type of cohort-based designs can be inefficient, especially if the trial needs to be frequently suspended. See Skolnik et al. (2008) and Doussau et al. (2016) for discussion. For example, subsequent patients can be turned away during trial suspension, resulting in waste of precious patient resource. In addition, trial duration is prolonged due to between-cohort suspension. Lastly, some cancer population is in high unmet medical needs with rapid disease deterioration and limited alternative therapies, and these clinical trials are often their last treatment options.

Cheung and Chappell (2000) proposed the time-to-event CRM design (TITE-CRM) that allows for continual accrual and dose escalation decisions while some patients' DLT data are pending. TITE-CRM is among the first dose-finding designs that incorporate partial time-to-event informa-

tion from patients who are still being assessed after treatment. This allows TITE-CRM to enroll future patients when some patients do not have observed DLT outcomes, thereby speeding up the trial. However, TITE-CRM, an extension of CRM, inherits some practical issues from CRM. For example, it tends to escalate too quickly leading to unsafe and irrational decisions that might be overruled by safety committees in practice. For a discussion of CRM and TITE-CRM's problem of irrational decisions, see Ji and Yang (2017), Zhou et al. (2018) and Yuan et al. (2018).

Skolnik et al. (2008) proposed the rolling six design (RSD) with the aim to shorten the duration of phase I pediatric oncology trials and reduce the number of accrual suspensions. This design allows up to six patients to be enrolled and followed for DLT evaluation at a dose level without enrollment suspension. Accrual is only suspended when awaiting data from six patients. The use of RSD has gained popularity since its introduction due to its simplicity and transparency similar to the 3+3 design. However, it also inherits the limitation of the 3+3 design (Ji and Wang, 2013, Onar-Thomas and Xiong, 2010, Zhao et al., 2011).First, like 3+3, RSD is only suitable for a target DLT rate around 17% since it defines the MTD as the dose level with one or less DLT out of six patients (Ananthakrishnan et al., 2017, Ji and Wang, 2013). Second, RSD is tailored for pediatric phase I trials in which doses are typically very safe (Lee et al., 2005). For adult trials, RSD may not be suitable since the doses may be overly toxic. We will show in our simulation that RSD indeed leads to trials with higher toxicity rates than the proposed R-TPI design. Third, it has been shown (Doussau et al., 2016) that RSD sometimes is not able to identify the true MTD with high accuracy and assigns many subjects to sub-optimal doses. This is the case if the first few low doses in the trial are quite safe and the true MTD is at a high dose level. Lastly, on average RSD requires more patients in a trial than the 3+3 design due to the enrollment of 6 patients (rather than 3) at a dose level (Sposto and Groshen, 2011).

The mTPI design (Ji et al., 2010, Ji and Wang, 2013) has been demonstrated to be superior to the 3+3 design. More recently, Guo et al. (2017) further improved the mTPI design with the development of the mTPI-2 design. These designs determine the dose level sequentially for each

cohort of patients by calculating the toxicity probability intervals under a Bayesian hierarchical model and optimal decision framework. Being model-based designs, mTPI and mTPI-2 exhibit safer and more reliable operating characteristics than 3+3 and at the same time, provide transparent decision rules prior to trial initiation. Motivated by the RSD and mTPI-2, we propose a rolling toxicity probability interval (R-TPI) design that combines the idea of rolling accrual in RSD with the model-based framework in mTPI-2. Under the R-TPI design, subjects will not be enrolled in cohorts with a fix cohort size and study accrual will rarely be suspended. That is, patients may be enrolled as they become available rather than enrolling in dosing cohorts (Broglio et al., 2015, Skolnik et al., 2008). In addition, R-TPI enjoys the benefits of model-based inference and overcomes the drawback of fixed rules as seen in RSD. For example, R-TPI is suitable for any target DLT rates (such as 10% or 30%) rather than just 17% as in 3+3 or RSD. And more importantly, R-TPI provides transparency to clinicians with precalculated decision tables allowing the decisions to be assessed and implemented easily. Through extensive simulation studies with 60 different scenarios, we assess the performance of R-TPI by using the matched mean sample sizes with 3+3 and RSD designs.

2 Review of Two Designs

2.1 Review of Rolling Six

The rolling six design (Skolnik et al., 2008) extends 3+3 with the aim to to reduce the occurrence of accrual suspension. In 3+3, patient accrual is suspended after the enrollment of 3 patients until the data of these 3 patients are completely observed. During the period of suspension, new eligible patients have to be turned away resulting in a waste of patient resources. The rolling six design reduces this type of suspension by allowing up to 6 patients to be enrolled without breaks at a dose level, i.e. no suspension after the enrollment of the first 3 patients at a dose. Dosing decisions

for a new patient is based on the number of patients enrolled to the current dose , the number of DLTs, and the number of patients still being followed without definitive DLT outcomes. RSD is a rule-based design and all dose assignment rules for the six patients can be prespecified (see Table 1).

Table 1: RSD Decision Table.

	Observed data at dose d			Decision	
#Enrolled	# DLTs	# Non-DLTs	# Pending	MTD Not Exceeded	MTD Exceeded
2	0, 1	any	any	S	-
2	2	0	0	D	-
3	0	0, 1, 2	3, 2, 1	S	-
3	0	3	0	E	-
3	1	0, 1, 2	2, 1, 0	S	-
3	≥ 2	any	any	D	-
4	0	0,1,2,3	4,3,2,1	S	S
4	0	4	0	E	S
4	1	0,1,2,3	3,2,1,0	S	S
4	≥ 2	any	any	D	D
5	0	0,1,2,3,4	5,4,3,2,1	S	S
5	0	5	0	E	S
5	1	0,1,2,3,4	4,3,2,1,0	S	S
5	≥ 2	any	any	D	D
6	0	0,1,2,3,4	6,5,4,3,2	Suspend	Suspend
6	0	5,6	1,0	E	MTD
6	1	0,1,2,3,4	5,4,3,2,1	Suspend	Suspend
6	1	5	0	E	MTD
6	≥ 2	any	any	D	D

2.2 Review of mTPI-2

The mTPI-2 design (Guo et al., 2017) is an extension of the mTPI design, which employs a simple Bayesian beta-binomial hierarchical model and an optimal decision framework. In mTPI, the toxicity rate of a dose can fall into three intervals corresponding to under, proper, and over dosing. The under-dosing interval is defined as $(0, p_T - \epsilon_1)$, the over-dosing interval as $(p_T + \epsilon_2, 1)$, and the equivalence interval as $(p_T - \epsilon_1, p_T + \epsilon_2)$ for proper dosing, where p_T is the target probability of DLT and ϵ_1 and ϵ_2 are used to quantify the tolerance that in real world a dose with similar DLT

rate to p_T can still be considered as the MTD. The three dosing intervals are associated with three different dose-escalation decisions. The under-dosing interval corresponds to a dose escalation (E), the over-dosing interval to a dose de-escalation (D), and the equivalence interval to staying at the current dose (S). The mTPI design selects one of the D, E, S decision for the current cohort of patients by picking the corresponding dosing interval that has the largest unit probability mass (UPM), defined as the probability of the interval divided by the length of the interval. For example, if the under-dosing interval has the largest UPM, decision E will be executed and the next cohort of patients will be treated at the next higher dose level. In addition, two safety rules are proposed in mTPI to protect patients from being exposed to overly toxic doses. These safety rules will be presented later in the proposed R-TPI design.

Due to a statistical phenomenon called the Ockham's razor (Good, 1967, Jefferys and Berger, 1992, Thorburn, 1918), the mTPI design includes some statistically sound but practically debatable decisions. For example, when the target rate $p_T = 0.3$, mTPI chooses S, to stay at the current dose, if 3 out of 6 patients experience DLTs. Because the empirical DLT rate is 3/6 or 50%, this decision may be considered overly aggressive if safety is heavily valued. Guo et al. (2017) showed that the decision is reasonable from a statistical point of view and is caused by Ockham's razor that prefers more parsimonious model. Nonetheless, the authors proposed a new design, called mTPI-2 to mitigate practical concerns due to safety. The mTPI-2 design divides the intervals $(0, p_T - \epsilon_1)$ and $(p_T + \epsilon_2, 1)$ into shorter subintervals with the same length as the equivalence interval $(p_T - \epsilon_1, p_T + \epsilon_2)$. Each subinterval corresponds to one of the three decisions E, S and D that is associated with the original interval. It is shown that the mTPI-2 design is based on a three-level hierarchical model treating each subinterval as a model, each assigned a prior probability. Then the optimal decision in selecting the model coincides with the optimal decision in making a E, S, or D decision. In mTPI-2, the optimality corresponds to a 0-1 loss and minimization of the Bayes risk, i.e., minimizing the posterior expected risk of wrong dose allocations (Guo et al., 2017). In addition, the same safety rules in the mTPI design are also applied to mTPI-2. One of the key features about mTPI and

mTPI-2 is that all the dosing decisions can be precalculated prior to the trial, which makes it transparent and appealing to clinicians.

3 The Proposed R-TPI Design

3.1 Notations

We propose the R-TPI design, a hybrid method combining the mTPI-2 and RSD designs. Consider a toxicity-driven phase I dose-finding trials. Let p_T be the target DLT probability, and p_d be the true and unknown DLT probabilities of dose level d, $d = 1, \ldots, D$, where D denotes the prespecified number of dose levels to be investigated. Generally, we assume that p_d is non-decreasing with dose level, i.e. $p_1 \leq p_2 \leq \cdots \leq p_D$. Assume at a given moment dose d is being used to treat enrolled patients and a total of $(n_d + m_d)$ patients have been assigned to dose d, among whom n_d patients have known outcomes (either with or without DLT) and m_d patients are still being followed without outcomes. Let y_d be the number of patients (among n_d) with DLT, therefore $(n_d - y_d)$ without DLT. The table below describes the breakdowns.

# with DLT	# without DLT	# being followed and no outcomes	Total at dose d
y_d	$(n_d - y_d)$	m_d	$(n_d + m_d)$

Imagine now a new patient becomes eligible for the trial. If $m_d > 0$, i.e., at least one patient is still being followed without definitive outcomes, the new patient is traditionally turned away under cohort-based designs. We consider a rolling enrollment strategy to reduce enrollment suspension and speed up the trial. In the meantime, we aim to maintain the overall patient safety and reliability (probability of selecting the correct MTD) of the trial. The basic idea is to take advantage of the model-based inference in mTPI-2 to avoid ad-hoc rules for dosing decisions, and at the same time, utilize the smart rolling enrollment proposed in RSD. We describe the R-TPI design next.

3.2 Dose-Assignment Decisions

The R-TPI design consists of two sets of enrollment schemes, namely the run-in enrollment and the rolling enrollment.

Run-in Enrollment To begin the trial, R-TPI enrolls the first patient at the starting dose level d_0. In many trials, $d_0 = 1$ is the first dose level but sometimes $d_0 > 1$ can be a higher dose. Subsequently, regardless of the patient's outcome R-TPI keeps enrolling new patients at dose d_0 until either (1) $n_{d_0} > 0$, i.e. there is at least one outcome at d_0, or (2) $n_{d_0} = 0$ and $m_{d_0} = C$, for a pre-determined C value, e.g., $C = 3$. In other words, (2) means that the first C patients have not completed followup at d_0 and are without definitive outcomes. The R-TPI design starts rolling enrollment (specified below) if (1) is the case. In the case of (2) R-TPI first suspends the enrollment until the first outcome at dose d_0 is observed, and then starts the rolling enrollment. Importantly, the run-in enrollment is applied to any new dose level as well when it is first used to treatment patients during the trial. For example, if dose d^* is decided to be the new dose level for treating patients and if d^* has not been used at any time of the trial, the run-in enrollment will be used for the initial patients assigned to dose d^*.

Rolling Enrollment Suppose at a given moment of the trial a new patient becomes eligible for enrollment, and the current dose used for treating patients is d at which $(n_d + m_d)$ patients have been treated. To explain the rolling enrollment, we need to introduce one more notation called k_d, which is best explained with an example as follows. Suppose a total of 6 patients have been treated at dose d. That is $(n_d + m_d)$ equals 6. There exist two ways with which the 6 patients can be assigned to dose d. First, these 6 patients can be enrolled consecutively after dose d is selected to treat patients in the trial. In other words, once dose d is selected for treating patients, the 6 subjects are enrolled one after another and the dose level is never changed. However, it is also possible that the 6 patients are not consecutively enrolled. For example, it is possible that initially 3 patients are

enrolled consecutively at dose level d; based on their DLT outcomes R-TPI changes the dose level to another dose and enrolls patients at the new dose; however, based on the patients DLT outcomes at the new dose R-TPI changes the dose level again, switches back to dose d, and enrolls additional 3 patients. In other words, the 6 patients can be enrolled consecutively or the 6 patients can be enrolled in batches that bracket dose changes. See Figure 1. In the two examples showed in Figure 1, the k_d values equal 6 in the first example where 6 patients are enrolled consecutively and equal 3 in the second example where 3 patients are enrolled in two non-consecutive batches. Therefore, k_d is defined as the number of patients at dose d since it most recently becomes the current dose. Later we will threshold k_d to speed up patient enrollment and reduce the up-and-down switches among doses in R-TPI.

The mTPI-2 design provides a decision $\mathcal{D}_{y,n} \in \{D, E, S\}$ when y out of n patients experience DLTs at a dose. With the notations $\mathcal{D}_{y,n}$, n_d, m_d, k_d, and C, we describe the rolling enrollment of R-TPI as follows. Suppose a new patient is eligible for enrollment. The box below describes the rolling enrollment rules. Remarks about the rules are given after the box.

I. If $m_d = 0$, i.e., all the patients enrolled at dose level d have completed their followup with definitive outcomes, assign the new patient according to $\mathcal{D}_{y_d,n_d}$, the decision of mTPI-2 when y_d out of n_d patients experience DLT outcomes.

II. If $0 < m_d \leq C$, i.e., some patients are still being followed without outcomes, consider three cases:

1. If $\mathcal{D}_{y_d,n_d}$ is D,
   (a) if $\mathcal{D}_{y_d,n_d+m_d}$ is D, de-escalate to dose level $(d-1)$; apply the run-in enrollment if dose $(d-1)$ is a new dose or re-apply the rolling enrollment if it has been used before;
   (b) else, the decision is S and continue patient enrollment at dose d.
2. If $\mathcal{D}_{y_d,n_d}$ is S, consider the following two cases;
   (a) if $\mathcal{D}_{y_d,n_d+m_d}$ is S, assign the new patient to d;
   (b) if $\mathcal{D}_{y_d,n_d+m_d}$ is E,
      i. if $k_d < 3$, enroll the next patient at dose d;
      ii. if $k_d \geq 3$, suspend the enrollment until more patients have observed their outcomes at dose d. Then recalculate the m_d value and re-apply **I** or **II.**
3. If $\mathcal{D}_{y_d,n_d}$ is E, consider the following two cases; ;
   (a) if $\mathcal{D}_{y_d+m_d,n_d+m_d}$ is E, escalate to dose level $(d+1)$; apply the run-in enrollment if dose $(d+1)$ is a new dose or re-apply the rolling enrollment if it has been used before.
   (b) else,
      i. if $k_d < 3$, enroll the next patient to dose d;
      ii. if $k_d \geq 3$, suspend the enrollment until more patients have observed their outcomes at dose d. Then recalculate the m_d value and re-apply **I** or **II.**

- If $m_d > C$, suspend the enrollment until more patients have observed outcomes at dose d.

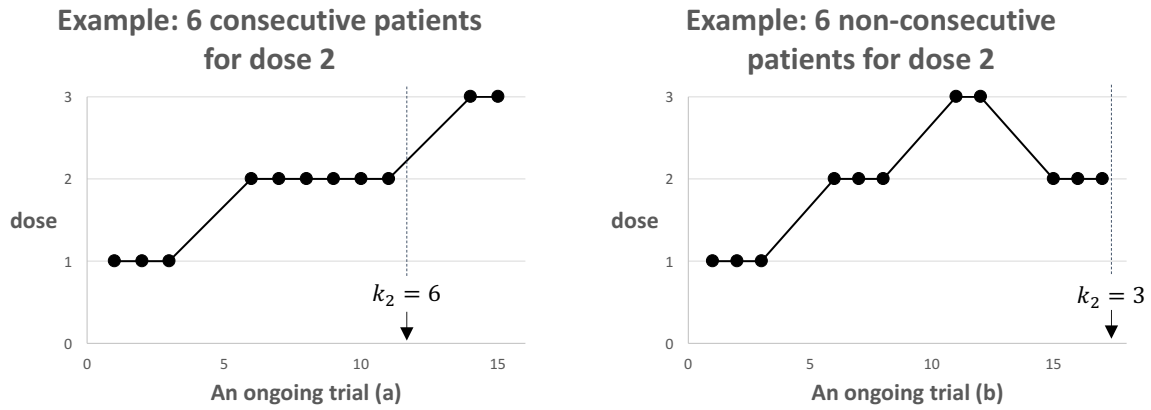

Figure 1: Illustration of two examples of enrollment with different k_d under the same $(n_d + m_d)$ value for dose $d = 2$. Trial (a): six patients are consecutively enrolled to dose 2, and therefore $k_2 = 6$. Trial (b): after enrolling 3 patients at dose 2, the dosing decision assigns the following 2 patients to dose 3; later, dose 2 is revisited and 3 more patients are assigned to dose 2; here k_2=3.

Here C is the maximum number of pending patients without observed outcomes allowed in the trial such that a new patient can be enrolled. The larger value of C the more rapid the enrollment is. Rules 1, 2(a), 2(b)-i, 3(a) and 3(b)-i speed up the patient enrollment since they allow the new patient to be enrolled even when some patients in the trials do not have outcomes yet. Safety is strongly enforced across these rules. For example, rule 3(a) only allows dose escalation if the mTPI-2 design would still escalate even if all of the m_d incompleters experienced DLT outcomes. In Rules 2(b) and 3(b), the value of k_d is the number of patients that has been enrolled at the current dose since the most recent dose change. When k_d is small, say < 3 (Rules 2(b)-i and 3(b)-i), patient enrollment is continued so that more information can be accumulated to help future inference. In contrast, if $k_d \geq 3$ (Rules 2(b)-ii and 3(b)-ii) enrollment is suspended until more patients have observed DLT outcomes. The value of 3 is used to mimic most rule- and model-based designs since

most of them use a cohort size of 3. Of course, it can be modified and calibrated via simulation for a practical trial. To see how k_d can help speed up the trial conduct, consider a trial with the MTD target $p_T = 0.17$. Suppose at dose 1 a total of 3 patients have been enrolled with 0 DLT out of the three patients, and at dose 2 a total of 3 patients have been enrolled and with 1 DLT out the three patients. Suppose the current dose is dose $d = 2$. Since 1 DLT out of 3 patients is observed at dose 2, the decision of R-TPI is D, to de-escalate. Consequently, the next enrolled patient should be assigned to dose 1. That is, the current dose is switched to $d = 1$. Once we switch, we have $k_1 = 0$ by definition. Now according to Rule 2(b)-i, we can enroll 3 consecutive patients at dose 1 without suspension. This is where the trial conduct is sped up. We have conducted sensitivity analyses (results not shown) with different thresholds of k_d. For example, we can insist on suspending the trial whenever a new patient is enrolled at dose 1 until his/her DLT outcome is observed. That is, change the threshold of k_d from 3 to 1. This change slows down the trial so much that the trial length under R-TPI is even longer than that under the 3+3 design. Of course, one can use a value larger than 3 to threshold k_d. While this may further increase trial speed, safety could be a problem if most newly enrolled patients experience DLT events.

For additional patient safety protection, we apply the same safety rules as in the mTPI-2 design. They are 1) if the lowest dose has a high probability of being above the MTD, i.e., $\Pr(p_1 > MTD|y_1, n_1) > 0.95$, the trial terminates before reaching the maximum sample size; 2) if for any dose d, $\Pr(p_d > MTD|y_d, n_d) > 0.95$, then dose d and all the doses higher than d will be suspended from the trial. The only difference here is that at the time a dose is deemed unsafe and suspended, there may be some patients with pending outcomes at this dose level. Once these data are observed later, if the safety rule is no longer violated, this dose could be reopened again for further evaluation. The R-TPI design continues until the maximum sample size is reached or no dose is deemed safe in the trial for evaluation. At the end of the trial, we apply the isotonic regression (Ivanova and Wang, 2006, Ji et al., 2010) to select the final MTD.

3.3 R-TPI Decision Table

The R-TPI design requires users to provide the value of p_T and (ϵ_1, ϵ_2). The p_T value can be easily elicited from the trial clinician. The values of ϵ_1 and ϵ_2 can be set at 0.05 as the default (Ji et al., 2010) or elicited by asking the clinician the lower and higher bound of the DLT rate that would still be considered as close to p_T. Also we need to elicit the value of C to control the speed of patient accrual. The default is $C = 3$. With the provided values of p_T, ϵ_1, ϵ_2, and C, one can generate the R-TPI decision table prior to the trial.

We provide two decision tables of up to seven patients for R-TPI with target DLT rate p_T equal to 0.17 or 0.3, $\epsilon_1 = \epsilon_2 = 0.05$, and $C = 3$. See Tables 2 and 3. One additional column of corresponding decisions of RSD is also provided in Table 2 and 3 for comparison. RSD only provides decisions for up to six patients.

4 Operating Characteristics

4.1 Simulation Setup

To assess and compare the operating characteristics of the R-TPI design with existing designs, we perform extensive simulation studies covering a wide range of real-world scenarios. We compare R-TPI to three designs, 3+3, RSD, and mTPI-2. The 3+3 and mTPI-2 designs used in the simulation studies also adopt the ethics constraint of "decision-in-advance" applicable to the real-life trials. For example, under 3+3 if 2 patients have been enrolled to a newly-used dose d and both of them experience DLTs, stop enrolling the third patient to d and de-escalate to $d - 1$ immediately. We set $C = 3$ or 6. The number of candidate dose levels in a trial is assumed to be $D = 3$, 4, 5, or 6. For each number of dose level D, we consider three target toxicity probabilities, $p_T = 0.10$, 0.17 or 0.30. For each pair of (p_T, D) values, we consider five scenarios. In total, there are 60 different scenarios in the simulation study, each scenario with a unique ascending toxicity profile

Table 2: R-PTI Decision Table with $p_T = 0.17$, $\epsilon_1 = \epsilon_2 = 0.05$ and $C = 3$. The last column is the corresponding decision of RSD (only showing decisions when MTD not exceeded in Table 1).

Observed data at dose d				Decision	
$n_d + m_d$	y_d	n_d	k_d	R-TPI	RSD
1	0	0	1	S	-
1	0	1	1	E	-
1	1	1	1	D	-
2	0	0,1	any	S	S
2	0	2	any	E	S
2	> 0	any	any	D	S or D
3	0	0,1,2	3	Suspend	S
3	0	1,2	< 3	S	S
3	0	3	any	E	E
3	> 0	any	any	D	S or D
4	0	1,2,3	3	Suspend	S
4	0	2,3	< 3	S	S
4	0	4	any	E	E
4	> 0	any	any	D	S or D
5	0	2,3,4	3	Suspend	S
5	0	3,4	< 3	S	S
5	0	5	any	E	E
5	1	any	any	S	S
5	> 1	any	any	D	D
6	0	3,4,5	3	Suspend	Suspend or E
6	0	4,5	< 3	S	-
6	0	6	any	E	E
6	1	any	any	S	Suspend or E
6	> 1	any	any	D	D
7	0	4,5,6	3	Suspend	-
7	0,1	5,6	< 3	S	-
7	0	7	any	E	-
7	1	5,6	3	Suspend	-
7	1	7	any	S	-
7	> 1	any	any	D	-

Table 3: R-PTI Decision Table with $p_T = 0.3$, $\epsilon_1 = \epsilon_2 = 0.05$, and $C = 3$. The last column is the corresponding decision of RSD (only showing decisions when MTD not exceeded in Table 1).

Observed data at dose d				Decision	
$n_d + m_d$	y_d	n_d	k_d	R-TPI	RSD
1	0	0	1	S	-
1	0	1	1	E	-
1	1	1	1	D	-
2	0	0,1	any	S	S
2	0	2	any	E	S
2	> 0	any	any	D	S or D
3	0	0,1,2	3	Suspend	S
3	0	1,2	< 3	S	S
3	0	3	any	E	E
3	1	any	any	S	S
3	> 1	any	any	D	D
4	0	1,2,3	3	Suspend	S
4	0	2,3	< 3	S	S
4	0	4	any	E	E
4	1	any	any	S	S
4	> 1	any	any	D	D
5	0	2,3,4	3	Suspend	S
5	0	3,4	< 3	S	S
5	0,1	5	any	E	E or S
5	1	3,4	≥ 3	Suspend	S
5	1	3,4	< 3	S	S
5	> 1	any	any	D	D
6	0	3,4	3	Suspend	Suspend
6	0	4	< 3	S	-
6	0	5	any	E	E
6	1	3,4,5	3	Suspend	Suspend
6	1	4, 5	< 3	S	-
6	0,1	6	any	E	E
6	2	any	any	S	D
6	> 2	any	any	D	D
7	0	4,5	3	Suspend	-
7	0	5	< 3	S	-
7	0	6	any	E	-
7	1	4,5,6	3	Suspend	-
7	1	5,6	< 3	S	-
7	0,1	7	any	E	-
7	2	6,7	any	S	-
7	> 2	any	any	D	-

(see Appendix A). For each enrolled patient, a random binary DLT/non-DLT outcome is generated with the true probability of toxicity for the corresponding dose at which the patient is assigned. To mimic real-life oncology dose-finding trials, each patient in the simulation study is also assigned an on-study start time (the gap between the time of arrival in the clinic and the starting time of treatment) and the probability of inevaluability (such as drop off). Following Skolnik et al. (2008), we assume that the DLT followup period for each patient is 21 days, the mean on-study start time is 5 days, and the mean inter-patient arrival time is 10 or 5 days, i.e. on average three or six patients per month. To further explore the performance of R-TPI, we also conduct simulation studies under a short mean inter-patient arrival time of 3 days, i.e. on average ten patients per month. Arrival time is sampled from exponential distribution, and on-study start time and time to DLT from uniform distributions. By analyzing twelve real completed children's oncology studies in Skolnik et al. (2008), we fix the inevaluable rate at 11%, and the time to inevaluability of a patient is sampled from a uniform distribution ranging from 0 to the sampled time to event (either DLT or non-DLT) of that patient.

For each scenario, we simulate 1,000 trials and report 1) the probability of correct selection (PCS) of the MTD, which measures the reliability of a design (Ji et al., 2010), 2) the probability of toxicity (POT) for each design, which is calculated as the ratio between the number of patients having DLT and the total number of patients accumulated over all the simulated trials, 3) the average trial duration. As noted in (Ji et al., 2010), the sample size of 3+3 (and RSD) cannot be fixed prior to the trial. For fair comparison, we calibrate the maximum sample size of R-TPI and mTPI-2 to match the average sample sizes of 3+3. The sample size matching procedure is as follows: 1) for each scenario, conduct trial simulation under the 3+3 design; 2) calculate the average sample size for 3+3 across the 1,000 simulated trials; 3) round up the average sample size for 3+3 and use it as the maximum sample size for R-TPI and mTPI-2 in the corresponding scenario. For RSD, the sample size cannot be matched with 3+3 since it cannot be stopped based on prespecified maximum sample size.

4.2 Simulation Results

[**Bench Mark Against RSD**] Skolnik et al. (2008) evaluated RSD under a single scenario with nine dose levels with true probabilities of toxicity (0.05,0.1,0.3,0.5,0.75,0.9,0.95,0.99,0.99). They also allowed a backup dose with 0.02 true probability of toxicity. They reported trial duration of 294(75) days (mean(sd)) based on simulated trials. We were able to replicate their results using our own computer program and obtained a result of 290 $\pm$ 82 days under RSD. We treated the difference between our results and those in their paper as round-off noise. Keeping the exact same simulation settings, we then applied R-TPI to the simulated trials and obtained trial duration of 287 $\pm$ 40 days. Therefore R-TPI results in about the same trial length as RSD using the single scenario in Skolnik et al. (2008). Furthermore, we found that under this scenario, RSD selects dose levels 2 with a probability 0.51, dose level 3 with a probability 0.28, compared to 0.54 and 0.20 for R-TPI. Since dose level 2 is the true MTD and dose level 3 is more toxic than the true MTD, R-TPI seems to perform better in selecting the correct dose than RSD. To examine performance of both designs in broader settings, we conducted simulations based on the aforementioned 60 scenarios. The PCS values are 0.368 (0.106) and 0.426 (0.053) for RSD and R-TPI, respectively. And the mean trial durations are 264 (59) and 249 (54) days for RSD and R-TPI, respectively. Both results favor R-TPI. We report more detailed results next.

[$C = 3$ **for R-TPI**] Figure 2 compares R-TPI, mTPI-2, 3+3 and RSD with matching sample size based on 3+3 when the mean arrival time is 10 days and $C = 3$. From top to bottom Figure 2 shows the PCS, POT, trial duration and sample size for each combination of p_T value and dose number. We can see that the mTPI-2 design has the highest PCS values in scenarios under $p_T = 0.1$ or 0.3. The R-TPI design shows nearly the same PCS values as mTPI-2 when $p_T = 0.1$ and slightly lower values when $p_T = 0.3$. The slight drop in PCS when $p_T = 0.3$ is because sometimes R-TPI escalates the dose slower than mTPI-2. Evidence can be seen using scenario 48 as an example, which has four doses (0.08, 0.16, 0.24, 0.44), $p_T = 0.3$ and the third dose is the true MTD. The PCS values

for mTPI-2 and R-TPI are 0.46 and 0.43, respectively. We find that 66.4% of simulated trials using R-TPI assign 3 or more to the third dose, compared with 81.0% using mTPI-2. Since mTPI-2 is quicker in reaching dose level three and assigns more patients there, it produces slighly larger PCS than R-TPI. However, these results are based on a small sample size of 16 patients. When we increase the sample size to 24 for this scenario, the PCS for mTPI-2 and R-TPI become 0.52 and 0.53, respectively.

The 3+3 design and RSD perform poorly in terms of reliability when $p_T = 0.1$ and 0.3. When $p_T = 0.17$, the 3+3 and RSD designs show great reliability since they are developed for this target toxicity probability value. But we can still observe that R-TPI has similar PCS with 3+3 and RSD even under $p_T = 0.17$.

In terms of safety performance of the designs based on the POT, both mTPI-2 and R-TPI exhibit much lower POT than 3+3 and RSD, especially when $p_T = 0.1$ or 0.17. Across all the simulation scenarios, the R-TPI and RSD designs lead to shorter trials by weeks to months when compared to 3+3 or mTPI-2. We can also see that RSD uses two to seven more subjects on average to complete a trial. The larger sample size of RSD offsets its fast speed of enrollment. That is why R-TPI sometimes can lead to a shorter trial than RSD.

In practice, because safety usually is the most important criterion, a faster design with less safety would not be acceptable. We see that R-TPI exhibits the strongest safety performance among the four designs with a faster enrollment than 3+3 and mTPI-2. In summary, compared to 3+3 and mTPI-2 the R-TPI design is safer and faster, and at the same time maintains high reliability in identifying the true MTD.

Figures 3 and 4 compare four designs with shorter mean patient arrival time of 5 days and 3 days, respectively. In terms of PCS and POT, the findings are similar to the above comparisons. Due to faster enrollment, trial duration is shorter than the previous simulation. The PCS, POT and sample size of mTPI-2 and 3+3 are the same since these two designs do not depend on enrollment speed. The RSD design seems to suffer from the faster enrollment, resulting in inflated sample sizes

(3 to 9 more subjects than 3+3 and R-TPI). In contrast, the R-TPI design exhibits almost identical sample size. Again, the plots show that R-TPI can speed up the trial and still maintain strong safety, making it a more attractive design than the other designs in practice.

[$C = 6$ **for R-TPI**] To illustrate the robustness of R-TPI, we set $C = 6$ for R-TPI and the mean arrival time at 10 days. Overall, the findings are very similar to $C = 3$ (results reported in Appendix B). We recommend $C = 3$ since we found that in some scenarios, with $C = 6$ R-TPI can over-enroll too many patients at a dose and therefore lead to slower escalation and less reliability.

5 Conclusion

The proposed R-TPI design aims to shorten the duration but still maintain safety and desirability of phase I dose-finding clinical trials. Unlike traditional cohort-based designs, the R-TPI design does not require to enroll patients according to a pre-planned cohort size, resulting in undesirable enrollment suspension. On the other hand, R-TPI overcomes many weaknesses of the Rolling Six design by borrowing features from model-based inference in mTPI-2. Therefore, R-TPI reduces the trial suspension between two cohorts, targets different p_T values with wider applications, and at the same time maintains simplicity and transparency in the implementation. This makes R-TPI a general design for both adult and pediatric oncology trials. The R-TPI design compares favorably to popular dose-finding designs, such as the 3+3 design through simulation studies. It results in great reduction in trial duration and accrual suspension without sacrificing safety. It also improves the reliability of finding the true MTD in many scenarios.

To maintain simplicity, we do not model the partial information (such as time-to-DLT) for the incompleters. This is a future direction. As another future direction, we will consider a rolling enrollment scheme for dose-finding trials incorporating efficacy outcomes.

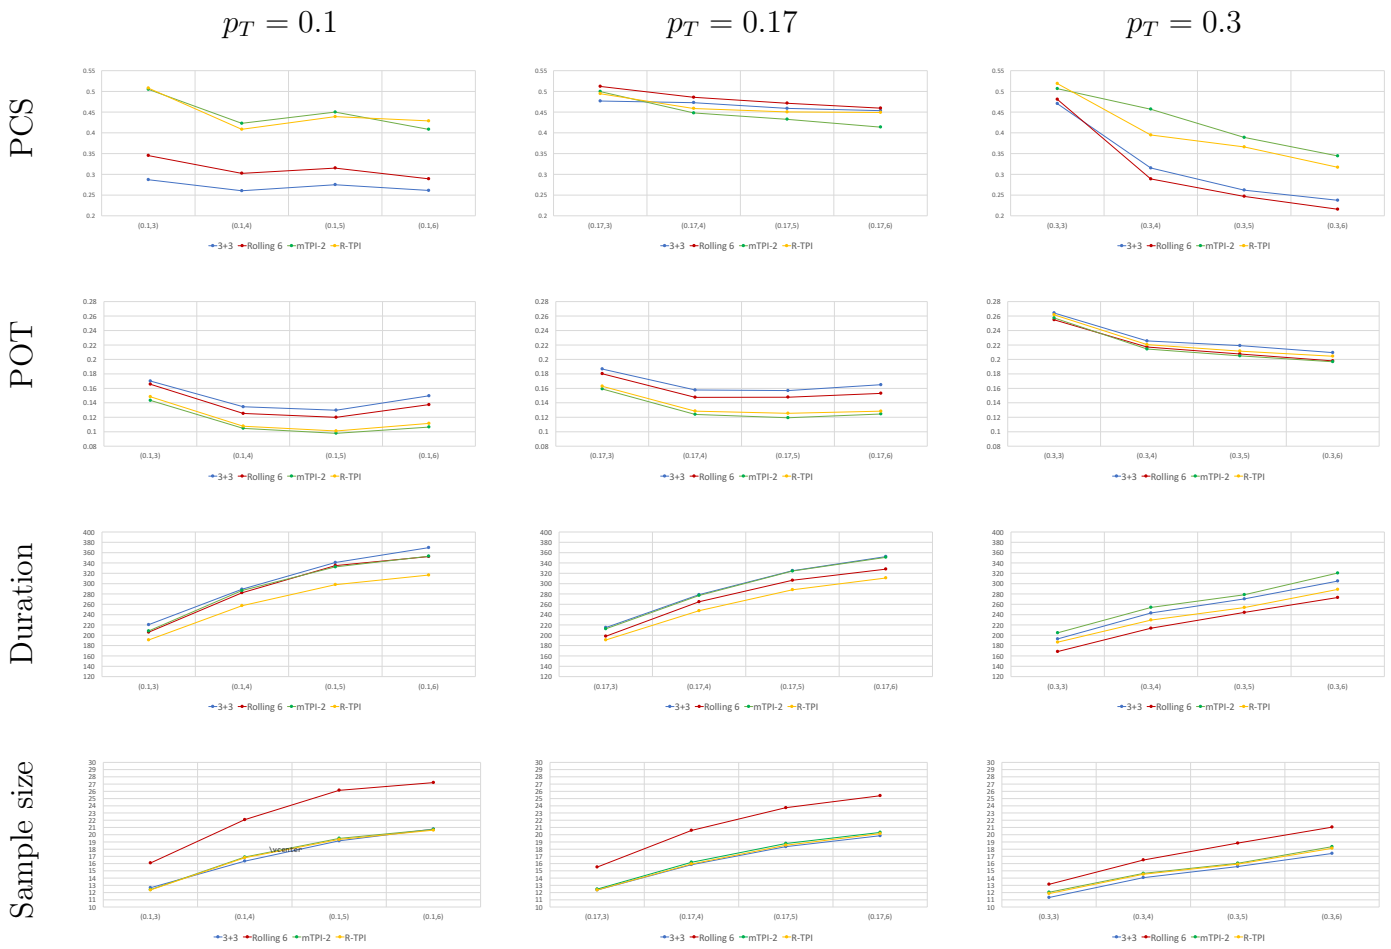

Figure 2: Comparison of the R-TPI ($C = 3$), mTPI-2, RSD and 3+3 designs with matching sample sizes when the mean arrival time is 10 days. Top to bottom we report mean PCS (prob. of correct selection of MTD), mean prob. of toxicity (POT), mean trial duration (days), and mean sample size from the simulation studies. The mean is taken over five scenarios for each pair of the three p_T values (0.1, 0.17, 0.3) and four numbers of dose levels (3-6) in the simulation. The horizontal axis displays combinations of p_T values and number of doses D, denoted by (p_T, D).

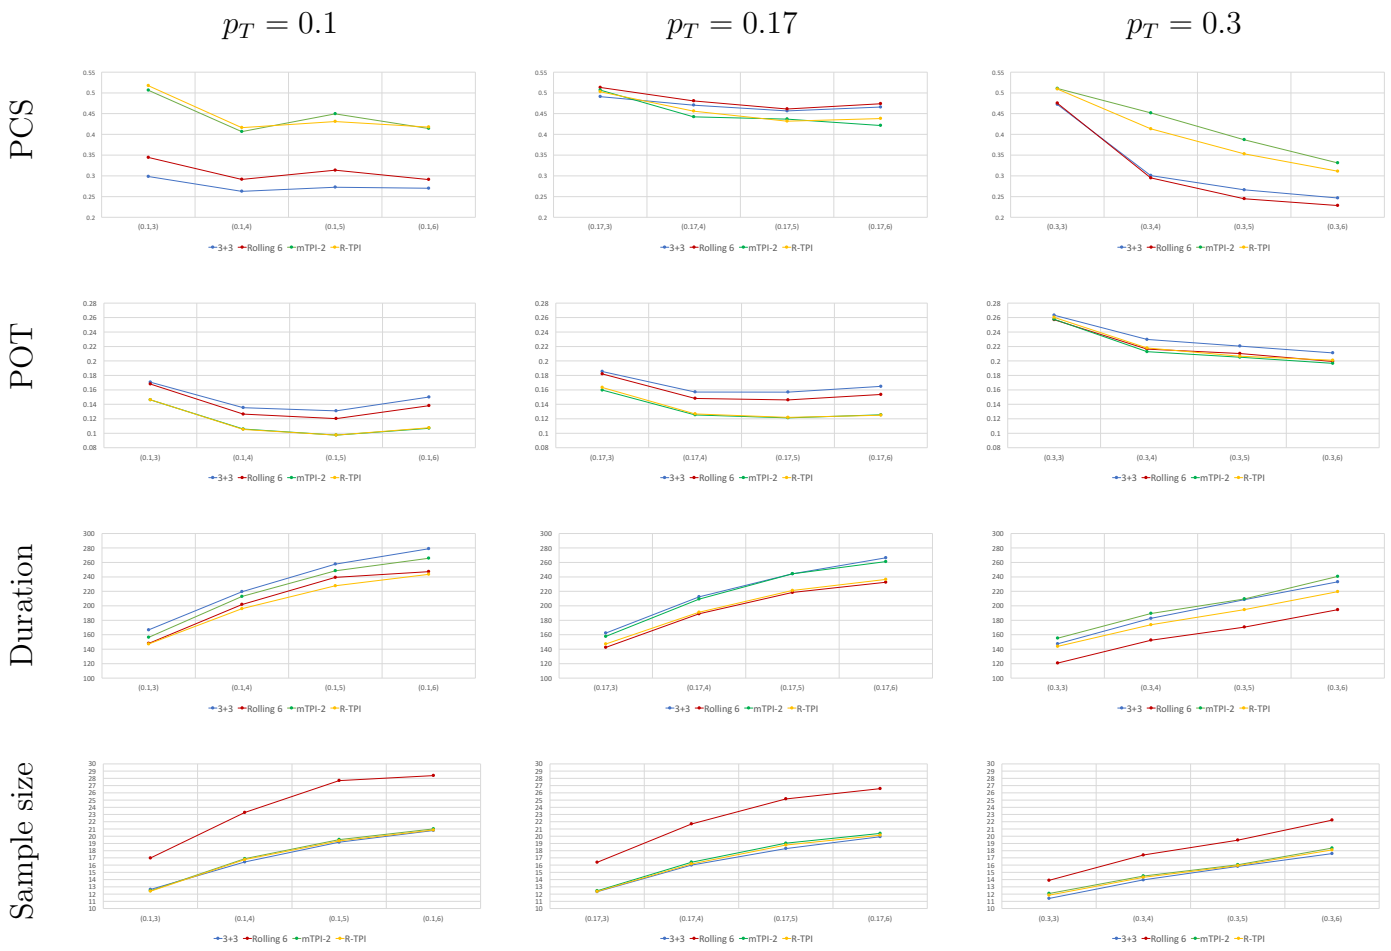

Figure 3: Comparison of the R-TPI ($C = 3$), mTPI-2, RSD and Rolling 6 designs with matching sample size when the mean arrival time is 5 days. Top to bottom we report mean PCS (prob. of correct selection of MTD), mean prob. of toxicity (POT), mean trial duration (days), and mean sample size from the simulation studies. The mean is taken over five scenarios for each pair of the three p_T values (0.1, 0.17, 0.3) and four numbers of dose levels (3-6) in the simulation. The horizontal axis displays combinations of p_T values and number of doses D, denoted by (p_T, D).

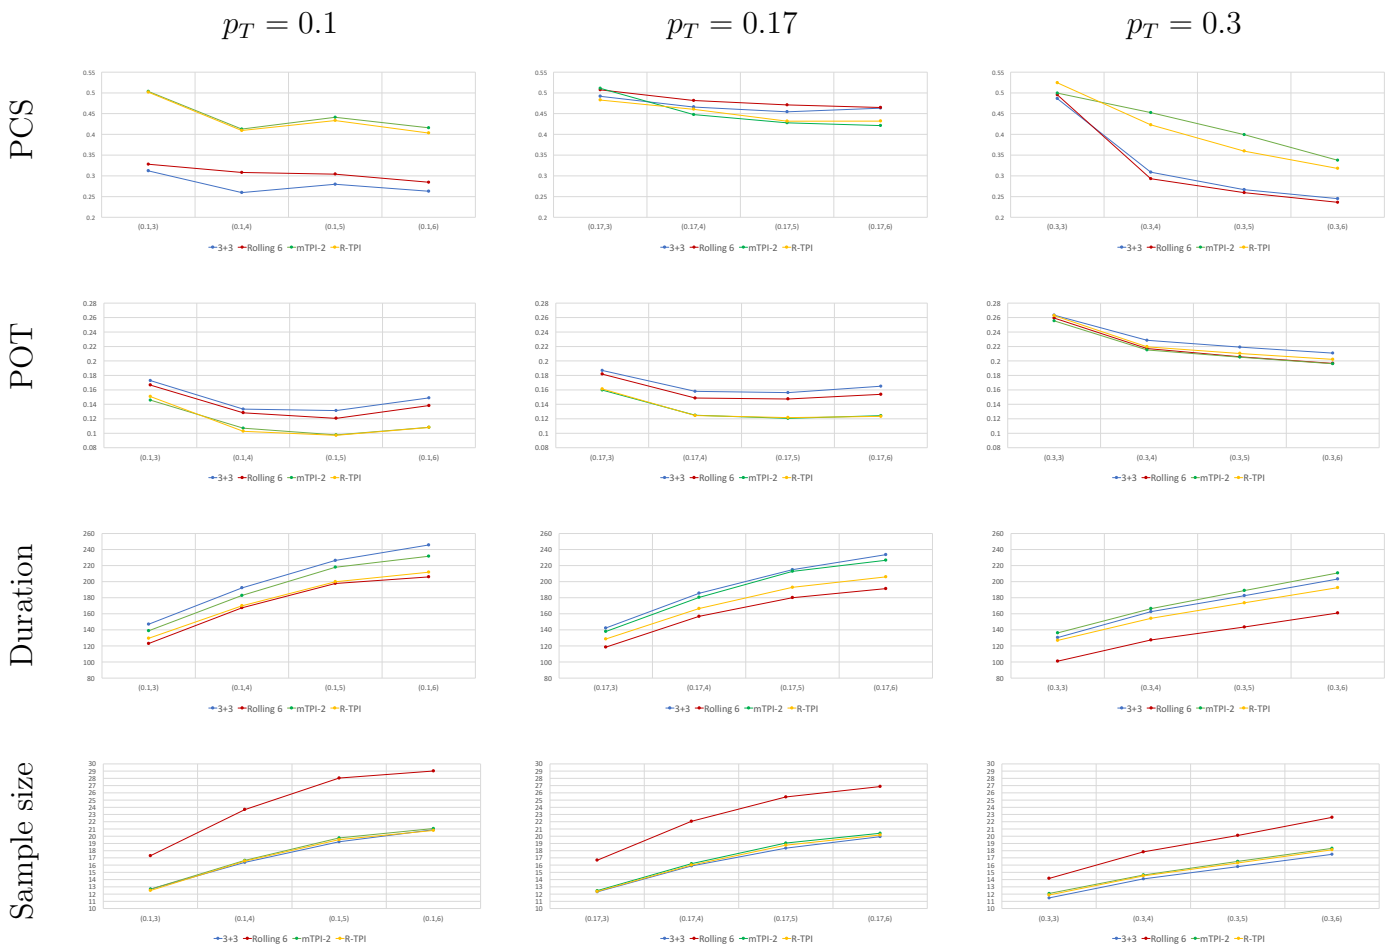

Figure 4: Comparison of the R-TPI ($C = 3$), mTPI-2, RSD and Rolling 6 designs with matching sample size when the mean arrival time is 3 days. Top to bottom we report mean PCS (prob. of correct selection of MTD), mean prob. of toxicity (POT), mean trial duration (days), and mean sample size from the simulation studies. The mean is taken over five scenarios for each pair of the three p_T values (0.1, 0.17, 0.3) and four numbers of dose levels (3-6) in the simulation. The horizontal axis displays combinations of p_T values and number of doses D, denoted by (p_T, D).

References

Ananthakrishnan, R., Green, S., Chang, M., Doros, G., Massaro, J., and LaValley, M. (2017). Systematic comparison of the statistical operating characteristics of various phase i oncology designs. *Contemporary clinical trials communications*, 5:34–48.

Broglio, K. R., Sandalic, L., Albertson, T., and Berry, S. M. (2015). Bayesian dose escalation in oncology with sharing of information between patient populations. *Contemporary clinical trials*, 44:56–63.

Cheung, Y. K. and Chappell, R. (2000). Sequential designs for phase i clinical trials with late-onset toxicities. *Biometrics*, 56(4):1177–1182.

Doussau, A., Geoerger, B., Jiménez, I., and Paoletti, X. (2016). Innovations for phase i dose-finding designs in pediatric oncology clinical trials. *Contemporary clinical trials*, 47:217–227.

Good, I. J. (1967). A bayesian significance test for multinomial distributions. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 399–431.

Guo, W., Wang, S.-J., Yang, S., Lynn, H., and Ji, Y. (2017). A bayesian interval dose-finding design addressingockham's razor: mtpi-2. *Contemporary clinical trials*, 58:23–33.

Ivanova, A. and Wang, K. (2006). Bivariate isotonic design for dose-finding with ordered groups. *Statistics in Medicine*, 25(12):2018–2026.

Jefferys, W. H. and Berger, J. O. (1992). Ockham's razor and bayesian analysis. *American Scientist*, 80(1):64–72.

Ji, Y., Liu, P., Li, Y., and Nebiyou Bekele, B. (2010). A modified toxicity probability interval method for dose-finding trials. *Clinical Trials*, 7(6):653–663.

Ji, Y. and Wang, S.-J. (2013). Modified toxicity probability interval design: a safer and more reliable method than the 3+ 3 design for practical phase i trials. *Journal of Clinical Oncology*, 31(14):1785.

Ji, Y. and Yang, S. (2017). On the Interval-Based Dose-Finding Designs. *ArXiv: 1706.03277 [stat.ME]*.

Lee, D. P., Skolnik, J. M., and Adamson, P. C. (2005). Pediatric phase i trials in oncology: an analysis of study conduct efficiency. *Journal of Clinical Oncology*, 23(33):8431–8441.

Onar-Thomas, A. and Xiong, Z. (2010). A simulation-based comparison of the traditional method, rolling-6 design and a frequentist version of the continual reassessment method with special attention to trial duration in pediatric phase i oncology trials. *Contemporary clinical trials*, 31(3):259–270.

Skolnik, J. M., Barrett, J. S., Jayaraman, B., Patel, D., and Adamson, P. C. (2008). Shortening the timeline of pediatric phase i trials: the rolling six design. *Journal of Clinical Oncology*, 26(2):190–195.

Sposto, R. and Groshen, S. (2011). A wide-spectrum paired comparison of the properties of the rolling 6 and 3+ 3 phase i study designs. *Contemporary clinical trials*, 32(5):694–703.

Thorburn, W. M. (1918). The myth of occam's razor. *Mind*, 27(107):345–353.

Yuan, Y., Lin, R., Li, D., Nie, L., and Warren, K. E. (2018). Time-to-event bayesian optimal interval design to accelerate phase i trials. *Clinical Cancer Research*, pages clincanres–0246.

Zhao, L., Lee, J., Mody, R., and Braun, T. M. (2011). The superiority of the time-to-event continual reassessment method to the rolling six design in pediatric oncology phase i trials. *Clinical Trials*, 8(4):361–369.

Zhou, H., Yuan, Y., and Nie, L. (2018). Accuracy, safety, and reliability of novel phase i trial designs. *Clinical Cancer Research*, pages clincanres–0168.

A The 60 Scenarios Used in Simulation

Table A.1: The true probabilities of toxicity of the 60 scenarios in the simulation study.

Scenario	p_T	D	Dose Level					
			1	2	3	4	5	6
1	0.1	3	0.05	0.1	0.15			
2	0.1	3	0.03	0.06	0.28			
3	0.1	3	0.06	0.2	0.3			
4	0.1	3	0.1	0.2	0.3			
5	0.1	3	0.21	0.32	0.43			
6	0.1	4	0.05	0.1	0.15	0.2		
7	0.1	4	0.03	0.06	0.1	0.15		
8	0.1	4	0.02	0.04	0.06	0.28		
9	0.1	4	0.05	0.1	0.2	0.3		
10	0.1	4	0.2	0.35	0.5	0.65		
11	0.1	5	0.05	0.1	0.15	0.2	0.25	
12	0.1	5	0.02	0.04	0.08	0.1	0.16	
13	0.1	5	0.02	0.04	0.06	0.08	0.28	
14	0.1	5	0.03	0.07	0.12	0.15	0.25	
15	0.1	5	0.19	0.32	0.45	0.58	0.71	
16	0.1	6	0.05	0.1	0.15	0.2	0.25	0.3
17	0.1	6	0.02	0.05	0.08	0.12	0.16	0.2
18	0.1	6	0.02	0.04	0.06	0.08	0.1	0.28
19	0.1	6	0.11	0.15	0.21	0.25	0.3	0.35
20	0.1	6	0.17	0.28	0.39	0.5	0.61	0.72

Table A.1 (continued)

Scenario	p_T	D	Dose Level					
			1	2	3	4	5	6
21	0.17	3	0.08	0.17	0.25			
22	0.17	3	0.06	0.12	0.34			
23	0.17	3	0.1	0.26	0.35			
24	0.17	3	0.04	0.08	0.12			
25	0.17	3	0.27	0.37	0.47			
26	0.17	4	0.08	0.17	0.25	0.33		
27	0.17	4	0.06	0.12	0.17	0.23		
28	0.17	4	0.04	0.08	0.12	0.34		
29	0.17	4	0.03	0.06	0.09	0.12		
30	0.17	4	0.26	0.4	0.54	0.68		
31	0.17	5	0.08	0.17	0.25	0.33	0.41	
32	0.17	5	0.04	0.08	0.12	0.17	0.25	
33	0.17	5	0.03	0.06	0.09	0.12	0.34	
34	0.17	5	0.03	0.06	0.09	0.12	0.15	
35	0.17	5	0.25	0.37	0.49	0.61	0.73	
36	0.17	6	0.08	0.17	0.25	0.33	0.41	0.49
37	0.17	6	0.03	0.1	0.15	0.2	0.25	0.3
38	0.17	6	0.03	0.06	0.09	0.12	0.15	0.34
39	0.17	6	0.04	0.08	0.1	0.12	0.14	0.16
40	0.17	6	0.24	0.34	0.44	0.54	0.64	0.74
41	0.3	3	0.15	0.3	0.45			
42	0.3	3	0.1	0.2	0.44			

Table A.1 (continued)

Scenario	p_T	D	Dose Level					
			1	2	3	4	5	6
43	0.3	3	0.18	0.38	0.46			
44	0.3	3	0.08	0.16	0.24			
45	0.3	3	0.39	0.48	0.57			
46	0.3	4	0.15	0.3	0.45	0.6		
47	0.3	4	0.1	0.2	0.3	0.4		
48	0.3	4	0.08	0.16	0.24	0.44		
49	0.3	4	0.06	0.12	0.18	0.24		
50	0.3	4	0.26	0.38	0.5	0.62		
51	0.3	5	0.15	0.3	0.45	0.6	0.75	
52	0.3	5	0.08	0.16	0.24	0.3	0.38	
53	0.3	5	0.06	0.12	0.18	0.24	0.44	
54	0.3	5	0.05	0.1	0.15	0.2	0.25	
55	0.3	5	0.27	0.37	0.47	0.57	0.67	
56	0.3	6	0.14	0.3	0.44	0.58	0.72	0.86
57	0.3	6	0.06	0.12	0.18	0.24	0.3	0.36
58	0.3	6	0.05	0.1	0.15	0.2	0.25	0.44
59	0.3	6	0.04	0.08	0.12	0.16	0.2	0.24
60	0.3	6	0.27	0.36	0.45	0.54	0.63	0.72

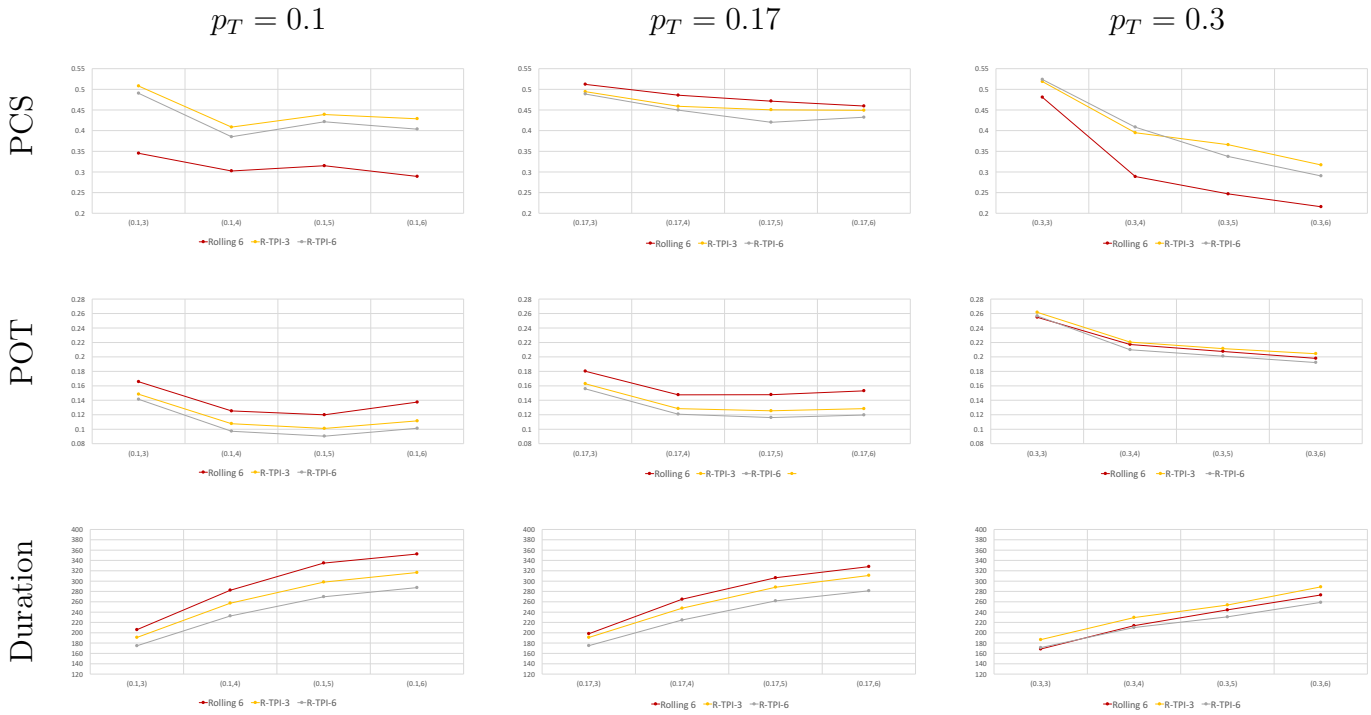

Figure 5: Comparison of the R-TPI ($C = 3$ and $C = 6$), and RSD designs with matching sample sizes. The mean arrival time is 10 days. Top to bottom and left to right we report mean PCS (prob. of correct selection of MTD), mean prob. of toxicity (POT), and mean trial duration (days). The mean is taken over five scenarios for each pair of the three p_T values (0.1, 0.17, 0.3) and four numbers of dose levels (3-6) in the simulation. The horizontal axis displays combinations of p_T values and number of doses D, denoted by (p_T, D).

B Simulation Results Under $C = 6$

We run the simulation with $C = 6$ for R-TPI and set the mean arrival time at 10 days in all scenarios. Results for RSD, R-TPI with $C = 3$, and R-TPI with $C = 6$ are reported in Figure 5. The trial speed of R-TPI with $C = 6$ is faster than R-TPI with $C = 3$ in all scenarios, but R-TPI with $C = 6$ has slightly lower PCS. The POT for R-TPI with $C = 6$ is also smaller than the other two designs.