original report

Bayesian Adaptive Design for Finding the Maximum Tolerated Sequence of Doses in Multicycle Dose-Finding Clinical Trials

abstract

Purpose **Statistical designs for traditional phase I dose-finding trials consider dose-limiting toxicity in the first cycle of treatment. In reality, patients often go through multiple cycles of treatment and may experience toxicity events in more than one cycle. Therefore, it is desirable to identify the maximum tolerated sequence of three doses across three cycles of treatment.**

Methods **Motivated by a three-cycle dose-finding clinical trial for a rare cancer with a JAK inhibitor, we proposed and implemented a simple Bayesian adaptive dose-cycle finding (BaSyc) design that allows intercycle and intrapatient dose modification. Because of the patient-specific dosing strategy over cycles, the BaSyc design is suited as a method in precision oncology.**

Results **BaSyc is simple and transparent because its algorithm can be summarized as two tabulated decision rules before the trial starts, allowing physicians to visually examine these rules. In addition, BaSyc employs a time-saving enrollment scheme that speeds up the trial. Extensive simulation studies show that BaSyc has desirable operating characteristics in identifying the maximum tolerated sequence.**

Conclusion **The BaSyc design provides a first-of-kind multicycle approach for dose finding and will likely lead to better and safer patient care and drug development.**

INTRODUCTION

Phase I oncology clinical trials are typically first-in-human studies characterized by a small number of patients with a goal of determining the maximum tolerated dose of an experimental drug. The maximum tolerated dose is defined as the highest dose at which the probability of dose-limiting toxicity (DLT) is no more than, or close to, a prespecified target response rate p_T (eg, $p_T = 0.3$). Numerous dose-finding designs have been developed, such as the 3 + 3 design,[1] the continual reassessment method,[2] and the modified toxicity probability interval (mTPI; mTPI-2) methods,[3-5] as recently reviewed by Sverdlov et al.[6] Most of the dose-finding trials consider DLT within the first cycle of treatment as the primary outcome and do not allow multicycle intrapatient dose modification. In recent drug development and patient care efforts, it has become more common to consider a regimen that involves multiple drugs over multiple cycles. For example, a new developmental drug may be considered as an additional treatment in chemotherapy that includes other standard drugs over multiple cycles.

Our work was motivated by and applied to such a trial at The University of Chicago for a rare form of leukemia. In this trial, a Janus kinase inhibitor, ruxolitinib, was added to a three-cycle US Food and Drug Administration (FDA)–approved chemotherapy backbone that includes cycle-specific drugs (Table 1). Although the dose of chemotherapy is fixed at the FDA-approved level, the investigational drug was administered at varying doses for each patient across cycles. Therefore, intrapatient dose modifications were required over three cycles to protect patient safety and learn the toxicity profile of different dose sequences. In other words, the trial could be considered as a single-agent multicycle

Jiaying Lyu
Emily Curran
Yuan Ji

Author affiliations and support information (if applicable) appear at the end of this article.

Corresponding author: Yuan Ji, PhD, The University of Chicago, 5841 South Maryland Ave MC 2000, Chicago, IL 60637; e-mail: koaeraser@gmail.com.

Table 1. Scheme of the Dose-Cycle Finding Clinical Trial

Induction (course 1)	Consolidation (course 2)	Interim Maintenance (course 3)	Delayed Intensification (course 4)	Maintenance (course 5)
Four-drug BFM-based regimen	Cyclophosphamide Vincristine Dexamethasome Pegylated asparaginase Cytarabine Mercaptopurine Methotrexate (IT)	Methotrexate (IV) Vincristine Pegylated asparaginase Methotrexate (IT)	Doxorubicin Cyclophosphamide Vincristine Dexamethasome Pegylated asparaginase Cytarabine Thioguanine Methotrexate (IT)	Vincristine Dexamethasone Mercaptopurine Methotrexate (PO) Methotrexate (IT)
Screening	Ruxolitinib	Ruxolitinib	Ruxolitinib	Monitoring
4 weeks	4-week treatment, 4-week washout	4-week treatment, 4-week washout	4-week treatment, 4-week washout	2-3 years

NOTE. The chemotherapy backbone is based on the C10403 regimen.[7] Ruxolitinib is added to courses 2 to 4 as a three-cycle treatment, with potentially varying doses across cycles.

Abbreviations: BFM, Berlin-Frankfurt-Munster; IT, intrathecally; IV, intravenously; PO, orally.

dose-finding study, which aimed to identify the maximum tolerated sequence (MTS) of ruxolitinib dose levels across three cycles. The MTS is defined as the highest sequence of three doses at which the probability of toxicity per cycle (TPC) is no more than or close to a prespecified target response rate p_T, given that the probability of toxicity at each cycle is also no more than or close to p_T. The probability of TPC of a sequence is defined as the average of three cycle-specific toxicity probabilities at the three dose levels of the sequence. Sequence A is higher than sequence B if the sum of the three dose levels of A is larger than the sum of B. Because of the low prevalence of the disease, the study sample size was small, up to 15 patients. Three doses of ruxolitinib were considered for each of the three cycles; therefore, there were a total of 27 possible dose-cycle combinations for three cycles of treatment.

The complication of the trial was the intrapatient dose modifications over cycles. A formal statistical inference would involve a dose-cycle-response model and optimal sequential decision making for each patient. Such approaches usually work well when the sample size of the trial is sufficiently large, as seen in the dynamic treatment regimen methods,[8] reinforcement learning,[9] and sequential multiple assignment randomized trials.[10] These approaches aim to better reflect the intrinsically multistage, adaptive structure of medical decisions, in both trial design and analysis of the trial data. However, they are not tailored to optimize patient doses over multiple cycles in a small dose-finding trial. Some works in the literature have investigated multiple-cycle treatment strategies, most of which have not considered intrapatient dose modification.[11-16] There are two exceptions. The first is Lee et al,[17] who proposed an innovative phase I/II design to adaptively and dynamically optimize a patient's dose in each of two cycles of therapy based on the joint binary cycle-specific efficacy and toxicity outcomes. Although this method allows intrapatient dose modification, it is only compatible with two cycles, requires relatively large sample sizes, and is complicated and difficult to implement. The second exception is Fernandes et al,[18] who developed an approach to model DLT in multiple treatment cycles based on Markov chains. However, this design is tailored for a specific trial setting and is difficult to generalize to trials in different settings.

BAYESIAN ADAPTIVE DOSE-CYCLE FINDING

Overview

The Bayesian adaptive dose-cycle finding (BaSyc) design consists of decisions to determine the appropriate dose in each cycle for enrolled patients and to select the MTS once the trial is completed. Patients are usually enrolled in cohorts, (eg, three patients per cohort). If a cohort of patients is either just entering the trial or completes treatment in a cycle, the BaSyc design applies two decision rules to the cohort: 1) the dose continuation rule (DCR) that determines whether the cohort should receive treatment in the next cycle for additional treatment

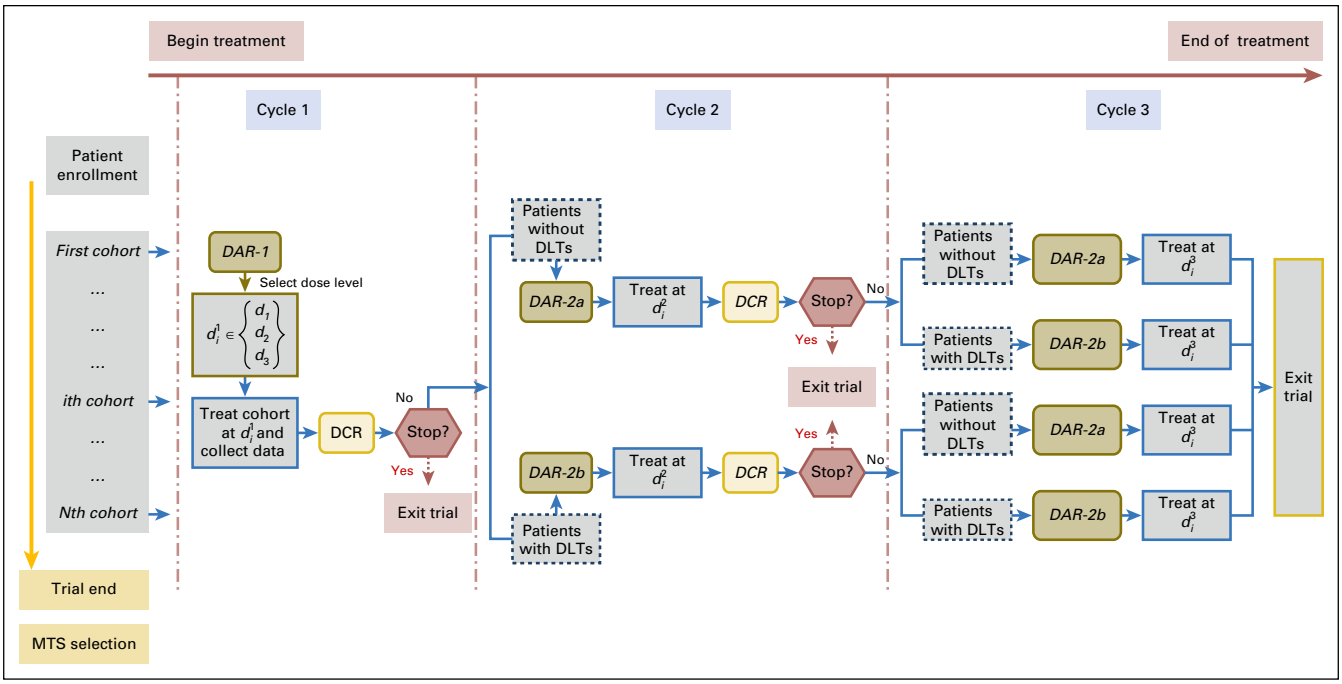

Fig 1. The Bayesian adaptive dose-cycle finding (BaSyc) design scheme for the motivating dose-cycle trial with three doses and three cycles of treatment. BaSyc consists of decisions to determine the dose in each cycle for the enrolled patients and to select the maximum tolerated sequence (MTS) when the trial is completed. In each cycle, the dose allocation rule (DAR) selects the doses for treating patients, in which DAR-1 selects the dose for cycle 1 and DAR-2 selects the dose for subsequent cycles. DAR-2 is composed of two subcases: DAR-2a and DAR-2b, which are applied to the two subcohorts—patients without dose-limiting toxicities (DLTs) and patients with DLTs, respectively. Also, in each cycle, the dose continuous rule (DCR) decides whether patients will move on to the next cycle for additional treatment. A patient follows the DAR and DCR throughout the trial and may end up with different decisions depending on the observed data. For example, a patient might exit the trial after one cycle of treatment (potentially because of toxicity [indicated by Stop]) based on the DCR applied in cycle 1. d_1, d_2, and d_3, three dose levels available; d_i^k, dose level treated at cycle k for patient i; N, maximum sample size.

(unless the cohort has completed all three cycles), and 2) the dose allocation rule (DAR) that decides the dose level of the next cycle if the DCR allows the patients to continue to the next cycle. The DCR and DAR are continuously applied to patients who are actively receiving treatment in any cycle of the treatment. At the end of the trial, BaSyc estimates the MTS based on a model-based inference accounting for the order of the dose sequences. Figure 1 illustrates the general scheme of the BaSyc design for the motivating trial with three doses and three cycles of treatment.

DCR and DAR

The basic concept of the DCR is that if the probability of toxicity of the lowest dose in the next cycle is considered lower than the target rate p_T, the cohort can move to the next cycle for additional treatment (ie, the decision is a go). This is because the cohort can always be treated at the lowest dose in the next cycle if the DCR is satisfied. Here, we assumed that there are intercycle washouts and palliative care that allow patients to recover from potential adverse effects from previous treatment. Technical details of the DCR are given in the Appendix.

If the DCR is a go for a cohort, the cohort will enter the next cycle for additional treatment. Then the DAR decides the dose level of the next cycle. There are three cases where the DAR is applied. First, the DAR always assigns the first cohort patients at the lowest dose level in all the cycles of treatment, which is a default rule to ensure the safety of the lowest dose in all three cycles. Second, for a newly enrolled cohort, the DAR decides the dose level for the cohort in the first cycle, labeled DAR-1. In this case, DAR-1 follows the same decision rule as in the mTPI-2 design.[5] Specifically, the dose level is decided based on the observed toxicity data from previously enrolled patients treated at the most recent dose in the first cycle. Given the observed data, the mTPI-2 design will decide the dose for the newly enrolled cohort. Third, the DAR is needed

when an enrolled cohort completes the existing cycle of treatment. Within a cohort, after a cycle of treatment, some patients may experience DLTs, and others may not. To personalize the decision for the patients in the same cohort, we proposed splitting the cohort into two subcohorts—patients without DLTs and those with DLTs. We used the DAR to specify the dose level of the next cycle for each subcohort separately, denoted as DAR-2a and DAR-2b.

DAR-2a. For the subcohort of patients without DLTs observed in the existing cycle, DAR-2a decides a dose based on the integration of the toxicity data of the cohort in the existing cycle and the toxicity data of the previously enrolled cohort that has already been treated in the next cycle.

DAR-2b. For the subcohort of patients with DLTs, DAR-2b assigns a lower dose in the next cycle that is deemed acceptable. This ad hoc rule is needed to protect patient safety.

The two subcohorts undergo the DCR and DAR separately hereafter until all three cycles of treatment are completed or until they exit the trial. Technical details are given in the Appendix.

Given the observed toxicity data, all the DCR and DAR decisions can be enumerated and summarized into tabular format. Figure 2 shows an example. In the first table, the DCR is summarized as go/no-go decisions. This table is applied to the lowest dose of the next cycle at which the current cohort would be treated. The second table provides the decisions based on the DAR. Motivated by the mTPI-2 design,[5] dose allocation is presented in the form of dose escalation, stay, or de-escalation to give the candidate dose level at the next cycle relevant to the dose level of the cycle we currently consider. This table is applied to the dose at which the current cohort is treated and the dose at which the previous cohort is most recently treated in the next cycle. Therefore, it might be used multiple times for the DAR. The two tables can be generated and predetermined before the trial starts for examination by physicians.

In the DAR, we used the mTPI-2 design as the backbone for decision making. To apply the decision rules of mTPI-2, we needed to provide the p_T value and an equivalence interval (EI) in the form of $(p_T - \varepsilon_1, p_T + \varepsilon_2)$. Here, the EI specifies the range of toxicity probabilities that is considered close to the target probability p_T by the mTPI-2 design. The default values of ε_1 and ε_2 are 0.05.

Dose-Finding Algorithm

To begin a trial, the BaSyc design treats the first cohort of patients at the lowest dose level in all three cycles. Then, after the toxicity outcomes from the first cycle are observed, the DCR decides whether the cohort will go to the next cycle, and the DAR decides which dose level to use if the DCR decision is go. Also, a new cohort is open for enrollment at this point at the first cycle as the current cohort moves onto the next cycle or exits the trial. BaSyc determines a dose level based on the DAR in the first cycle for this newly enrolled cohort of patients. This process is repeated until all cohorts finish all cycles of treatment unless the trial is terminated early because of excessive toxicity (safety rules are given in the Appendix). As a default rule, the cohort at a given cycle must wait for the cohort at the next cycle to complete follow-up before it can enter the next cycle.

We demonstrate in Table 2 how the BaSyc design guides the motivating dose-cycle trial with three doses and three cycles. In this trial, each cycle lasted 8 weeks consisting of a 4-week treatment followed by 4-week washout and recovery. A walk-through of the table is provided in the Appendix. For safety, the BaSyc design applied additional safety rules[3,19] throughout the entire trial. Details are provided in the Appendix.

Cohort Enrollment

It is worth noting that the BaSyc design uses a fast cohort enrollment scheme. Specifically, once the current cohort completes the first cycle of treatment with observed outcome data, the next cohort can be enrolled and assigned a dose level in the first cycle. The current cohort will either exit the trial or move to the next cycle. For example, in Table 2, in week 24, after cohort 3 completed the first-cycle treatment at dose level 2, the next cohort (cohort 4) was enrolled and treated at dose level 1 as its first-cycle therapy. Cohort 3 would continue to the next cycle for additional treatment. This enrollment scheme is important to speed up the trial with a small sample size targeting a rare disease.

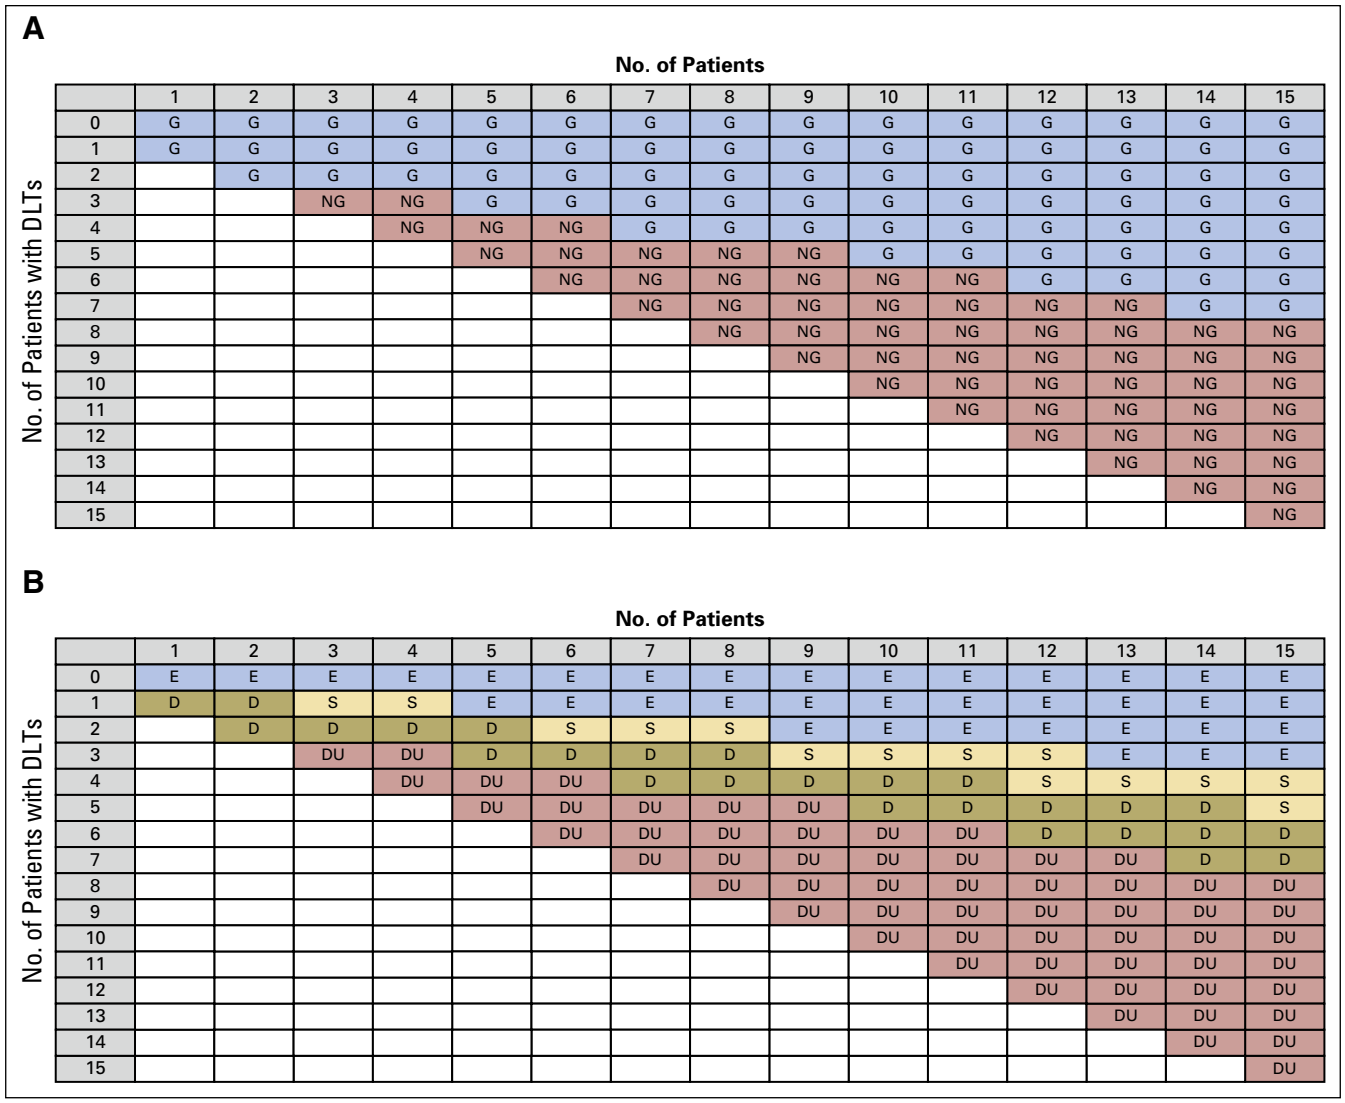

A	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	G	G	G	G	G	G	G	G	G	G	G	G	G	G	G
1	G	G	G	G	G	G	G	G	G	G	G	G	G	G	G
2		G	G	G	G	G	G	G	G	G	G	G	G	G	G
3			NG	NG	G	G	G	G	G	G	G	G	G	G	G
4				NG	NG	NG	G	G	G	G	G	G	G	G	G
5					NG	NG	NG	NG	NG	G	G	G	G	G	G
6						NG	NG	NG	NG	NG	NG	G	G	G	G
7							NG	NG	NG	NG	NG	NG	NG	G	G
8								NG	NG	NG	NG	NG	NG	NG	NG
9									NG	NG	NG	NG	NG	NG	NG
10										NG	NG	NG	NG	NG	NG
11											NG	NG	NG	NG	NG
12												NG	NG	NG	NG
13													NG	NG	NG
14														NG	NG
15															NG

B	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E
1	D	D	S	S	E	E	E	E	E	E	E	E	E	E	E
2		D	D	D	D	S	S	S	E	E	E	E	E	E	E
3			DU	DU	D	D	D	D	S	S	S	S	E	E	E
4				DU	DU	DU	D	D	D	D	D	S	S	S	S
5					DU	DU	DU	DU	DU	D	D	D	D	D	S
6						DU	DU	DU	DU	DU	DU	D	D	D	D
7							DU	DU	DU	DU	DU	DU	DU	D	D
8								DU	DU	DU	DU	DU	DU	DU	DU
9									DU	DU	DU	DU	DU	DU	DU
10										DU	DU	DU	DU	DU	DU
11											DU	DU	DU	DU	DU
12												DU	DU	DU	DU
13													DU	DU	DU
14														DU	DU
15															DU

Fig 2. Two decision tables of the Bayesian adaptive dose-cycle finding (BaSyc) design for the motivation trial. The sample size was 15, target toxicity probability $p_T = 0.3$, and equivalence margins $\varepsilon_1 = \varepsilon_2 = 0.05$. Each column number represents the number of patients treated, and each row number represents the number of patients experiencing toxicity, at a given dose where the decision is applied. (A) Go (G)/no-go (NG) decisions according to the dose continuous rule (DCR), in which the letters in blue and red represent G and NG decisions, respectively. This table is applied to the lowest dose of the next cycle at which the current cohort might be treated. For example, if the current cohort has completed treatment in cycle 1, the table is applied to dose level 1 in cycle 2. (B) Dose allocation decisions according to the dose allocation rule (DAR), in which the letters in blue, gold, green, and red represent dose escalation (E), staying at the current dose (S), dose de-escalation (D), and dose de-escalation and exclusion (DU) because of toxicity, respectively. This table is applied to the dose at which the current cohort is treated and the dose at which the previous cohort is treated in the next cycle.

MTS Selection

At the end of the trial, among all the candidate sequences of dose-cycle combinations, a sequence of three doses for three cycles would be selected as the MTS. We develop a method based on the isotonic transformation[20] of the estimated probability of TPC for each sequence to determine the MTS. Technical details are provided in the Appendix.

SIMULATION STUDIES

Simulation Setup

We evaluate the reliability (MTS recommendation) and safety of the BaSyc design using computer-simulated trials based on the setting of the motivating trial. There are three prespecified doses and three cycles. The maximum sample size is 15, and patients are enrolled in cohorts of three. The MTS target toxicity probability is

Table 2. Hypothetical Dose-Cycle Trial With Three Doses and Three Cycles Based on the BaSyc Design

			Cycle 1				Cycle 2				Cycle 3			
Week	**Cohort**	**Patient ID**	**D1**	**D2**	**D3**	**DCR**	**D1**	**D2**	**D3**	**DCR**	**D1**	**D2**	**D3**	**DCR**
8	1	1	**O**			Go to cycle 2	√							
		2	**O**				√							
		3	**O**				√							
	2 (enroll)	4-6		√										
	Total		0 of 3	0 of 0	0 of 0		0 of 0	0 of 0	0 of 0		0 of 0	0 of 0	0 of 0	
16	1	1	O				**O**			Go to cycle 3	√			
		2	O				**O**				√			
		3	O				**O**				√			
	2	4		**X**		Go to cycle 2	√							
		5		**O**				√						
		6		**O**				√						
	3 (enroll)	7-9		√										
	Total		0 of 3	1 of 3	0 of 0		0 of 3	0 of 0	0 of 0		0 of 0	0 of 0	0 of 0	
24	1	1	O				O				**O**			Exit
		2	O				O				**X**			
		3	O				O				**O**			
	2	4		X			**O**			Go to cycle 3	√			
		5		O				**O**			√			
		6		O				**X**			√			
	3	7		**O**		Go to cycle 2	√							
		8		**X**			√							
		9		**X**			√							
	4 (enroll)	10-12	√											
	Total		0 of 3	3 of 6	0 of 0		0 of 4	1 of 2	0 of 0		1 of 3	0 of 0	0 of 0	

(Continued on following page)

Table 2. Hypothetical Dose-Cycle Trial With Three Doses and Three Cycles Based on the BaSyc Design (Continued)

Week	Cohort	Patient ID	Cycle 1				Cycle 2				Cycle 3			
			D1	D2	D3	DCR	D1	D2	D3	DCR	D1	D2	D3	DCR
32	1*	1-3*	0 of 3*				0 of 3*				1 of 3*			Exit*
	2	4		X			O				**O**			Exit
		5		O				O			**O**			
		6		O				X			**O**			
	3	7		O			**O**			Go to cycle 3		√		
		8		X			**O**					√		
		9		X			**O**					√		
	4	10	**O**			Go to cycle 2		√						
		11	**O**					√						
		12	**X**				√							
	5 (enroll)	13-15		√										
	Total		1 of 6	3 of 6	0 of 0		0 of 7	1 of 2	0 of 0		1 of 6	0 of 0	0 of 0	
40	1*	1-3*	0 of 3*				0 of 3*				1 of 3*			Exit*
	2*	4-6*		1 of 3*			0 of 1*	1 of 2*			0 of 3*			Exit*
	3	7		O			O					**X**		Exit
		8		X			O					**O**		
		9		X			O					**O**		
	4	10	O					**O**		Go to cycle 3		√		
		11	O					**O**				√		
		12	X				**O**					√		
	5	13		**O**		Go to cycle 2	√							
		14		**O**			√							
		15		**X**			√							
	Total		1 of 6	4 of 9	0 of 0		0 of 8	1 of 4	0 of 0		1 of 6	1 of 6		

(Continued on following page)

Table 2. Hypothetical Dose-Cycle Trial With Three Doses and Three Cycles Based on the BaSyc Design (Continued)

Week	Cohort	Patient ID	Cycle 1				Cycle 2				Cycle 3			
			D1	D2	D3	DCR	D1	D2	D3	DCR	D1	D2	D3	DCR
48	1*	1-3*	0 of 3*				0 of 3*				1 of 3*			Exit*
	2*	4-6*		1 of 3*			0 of 1*	1 of 2*			0 of 3*			Exit*
	3*	7-9*		2 of 3*			0 of 3*					1 of 3*		Exit*
	4	10	O					O				**O**		Exit
		11	O					O				**X**		
		12	X				O					**O**		
	5	13		O			**X**			Go to cycle 3	√			
		14		O			**O**					√		
		15		X			**O**					√		
	Total		1 of 6	4 of 9	0 of 0		1 of 11	1 of 4	0 of 0		1 of 6	2 of 6		
56	1*	1-3*	0 of 3*				0 of 3*				1 of 3*			Exit*
	2*	4-6*		1 of 3*			0 of 1*	1 of 2*			0 of 3*			Exit*
	3*	7-9*		2 of 3*			0 of 3*					1 of 3*		Exit*
	4*	10-12*	1 of 3*				0 of 1*	0 of 2*				1 of 3*		Exit*
	5	13		O			X				**X**			Exit
		14		O			O					**X**		
		15		X			O					**O**		
	Total		1 of 6	4 of 9	0 of 0		1 of 11	1 of 4	0 of 0		2 of 7	3 of 8	0 of 0	

NOTE. Sample size is 15, and cohort size is three. The first column tracks the duration of the trial in weeks. Within each cycle, the dose allocation rule (DAR) decides the dose level of the next cycle if the dose continuation rule (DCR) allows the patients to continue to the next cycle as well as the dose continuation rule/decision (eg, go to cycle 2). Also, within each section, √ indicates the DAR decision for the corresponding cohort, showing the dose level to be assigned to the cohort in that cycle. Bold font highlights the current dose level being used in each cycle. The final row within each week shows the cumulative No. of dose limiting toxicities experienced by the No. of patients.

Abbreviations: BaSyc, Bayesian adaptive dose-cycle finding; D1, dose level 1; D2, dose level 2; D3, dose level 3; DCR, dose continuation rule; O, no dose-limiting toxicity; X, dose-limiting toxicity; √, the dose allocation rule decision for the corresponding cohort.

*Cohorts that have completed all three cycles or exit the trial.

p_T = 0.3, and the choices of ε_1 and ε_2 are both 0.05. The probability threshold for the additional safety rules is ξ = 0.95. We investigate six different scenarios in the simulation studies, as summarized in Table 3. We provided the true toxicity probabilities of the three dose levels in the first cycle and assumed that the toxicity probability would increase by 10% per cycle at the same dose level. For example, in scenario 1, the probability of toxicity for dose level 1 in cycle 1 is equal to 0.15. Then, in cycles 2 and 3, the true probability is assumed to increase to 0.165 and 0.1815, respectively. For each scenario, 10,000 simulated trials are conducted on computer.

Operating Characteristics

The simulation results for scenarios 1 to 6 are summarized in Table 3, with four sections per scenario. Section 1 lists the true toxicity probabilities of three prespecified dose levels in the first cycle and the target toxicity probability p_T. Sections 2 to 5 represent four different criteria that summarize the performance of the BaSyc design. Section 2 presents the frequency of MTS selection. The higher the frequency, the better the performance. Recall that the MTS is defined as the highest sequence of three safe doses with the true probability of TPC (calculated as the average of the three true toxicity probabilities corresponding to the three dose levels) closest to p_T and falling in the EI $(p_T - \varepsilon_1, p_T + \varepsilon_2)$. If a scenario has no dose sequence falling in the EI, the highest sequence of three safe doses with the true probability of TPC less than p_T is considered the true MTS. Here, a safe dose is a dose with toxicity probability less than $(p_T + \varepsilon_2)$. In our simulation, EI = (0.25, 0.35), and we highlight the true MTS with bold face. Section 3 provides the average number of toxicities and patients treated at each dose cycle. A desirable design should allocate most patients at doses with a probability of toxicity below p_T. Section 4 presents the percentage of patients treated above the MTS, and we would want to see a small value. Section 5 provides the percentage of early termination according to rule 1. This percentage should be high if the lowest sequence is overly toxic. The BaSyc design performs well, exhibiting a strong safety profile and desirable power in selecting the MTS. Simulation results for a larger sample size are summarized in Table 4. Details of the results and a sensitivity analysis of the sample size are provided in the Appendix.

DISCUSSION

We propose the BaSyc design to fulfill the need for statistical designs for multiple-cycle dose-finding trials that allow for intrapatient modification. No practical methods are available, and the two existing designs[17,18] in the literature are either highly complex or tailored for specific trials.

Perhaps the most attractive feature of the BaSyc design is its simplicity and transparency, as shown in Figure 2. In addition, the performance of the BaSyc design is promising. In particular, the BaSyc design exhibits strong safety, because it tends to assign patients to the doses with safe toxicity rates. Appendix Table A1 summarizes the dose allocation decisions for the BaSyc design across 10,000 simulated trials. We can see that escalation and de-escalation will never happen if the true toxicity probability of a dose is located in the higher interval $(p_T + \varepsilon_2, 1)$ and the lower interval $(0, p_T + \varepsilon_1)$, respectively.

In this report, we implemented BaSyc for a real-world trial with three dose levels for three cycles of treatment. This trial is pending after receiving approval for the investigational new drug application by the FDA for review. The BaSyc design can be easily applied to different numbers of doses and cycles with almost no modification. However, in practice, having many doses will quickly lead to an unmanageable number of possible dose-cycle combinations. For example, a total of 5^3 = 125 sequences are available for trials with five dose levels and three cycles.

BaSyc uses independent beta/binomial models as the working model to estimate the toxicity probability of each dose in each cycle based on all patients' data at that dose/cycle. This naïve construction leads to simplicity and transparency. Alternatively, one could consider a dependent model, such as the Markov model in Fernandes et al,[18] to estimate the toxicity probabilities.

In BaSyc, we focused on the dose-cycle trials with binary outcomes. Because BaSyc is a hybrid of a model-free algorithm and a model-based method, our proposed design can be extended to ordinal outcomes. The dose-finding methods that account for toxicity categorized according to their severity[21,22] could replace mTPI-2, which is used for the DCR and DAR in BaSyc.

Regarding the sample size of the trial, Ji et al[4] suggested that the sample size not be less than

Table 3. Simulation Results for Scenarios 1 to 6 With a Sample Size of 15 Patients and Three Patients per Cohort

Scenario	Dose Sequence Selection					No. of DLT of No. Patients			Patients Treated Above MTS (%)	Probability of Early Termination (%)
	Dose Sequence	Probability of TPC	Selection (> 5%)	No. of DLTs, mean (SD)	Allocation (%)	Cycle 1	Cycle 2	Cycle 3		
1*									5.1	3.32
Dose No.										
1	1, 1, 1	0.166	0.242	0.544 (0.671)	0.304	1.000 of 6.591	1.638 of 9.854	1.983 of 10.844		
2	2, 1, 1	0.216	0.197	0.902 (0.769)	0.164	1.784 of 5.847	1.486 of 4.518	1.315 of 3.649		
3	**2, 2, 1**	0.270	0.190	0.645 (0.625)	0.089	1.061 of 2.373	0.177 of 0.348	0.078 of 0.141		
4	2, 2, 2	0.331	0.131	0.363 (0.481)	0.034					
5	3, 2, 1	0.320	0.070	1.173 (0.800)	0.063					
6	3, 1, 1	0.266	0.056	0.869 (0.756)	0.035					
2†									6.8	3.75
Dose No.										
1	1, 1, 1	0.166	0.209	0.534 (0.666)	0.288	0.936 of 6.246	1.546 of 9.369	1.883 of 10.434		
2	2, 2, 1	0.250	0.185	0.629 (0.627)	0.090	1.543 of 5.763	1.439 of 4.818	1.289 of 3.930		
3	2, 1, 1	0.206	0.174	0.869 (0.771)	0.145	1.247 of 2.766	0.240 of 0.483	0.114 of 0.207		
4	**2, 2, 2**	0.298	0.148	0.328 (0.469)	0.042					
5	3, 2, 1	0.310	0.082	1.158 (0.818)	0.070					
6	3, 2, 2	0.358	0.065	0.874 (0.682)	0.032					
7	3, 1, 1	0.266	0.057	0.883 (0.761)	0.039					
3‡									0	0.43
Dose No.										
1	3, 3, 2	0.233	0.209	0.597 (0.630)	0.080	0.340 of 4.325	0.518 of 5.961	0.670 of 6.899		
2	**3, 3, 3**	0.265	0.204	0.289 (0.453)	0.044	0.809 of 5.075	1.136 of 6.468	1.210 of 6.310		
3	3, 2, 2	0.203	0.163	0.579 (0.627)	0.083	1.334 of 5.572	0.664 of 2.531	0.501 of 1.738		
4	2, 2, 2	0.177	0.140	0.196 (0.397)	0.094					
5	3, 2, 1	0.171	0.066	0.975 (0.782)	0.068					
6	1, 1, 1	0.088	0.061	0.276 (0.503)	0.215					
7	2, 2, 1	0.144	0.059	0.534 (0.579)	0.074					

(Continued on following page)

Table 3. Simulation Results for Scenarios 1 to 6 With a Sample Size of 15 Patients and Three Patients per Cohort (Continued)

Scenario	Dose Sequence Selection					No. of DLT of No. Patients			Patients Treated Above MTS (%)	Probability of Early Termination (%)
	Dose Sequence	Probability of TPC	Selection (> 5%)	No. of DLTs, mean (SD)	Allocation (%)	Cycle 1	Cycle 2	Cycle 3		
4§									22.6	99.21
Dose No.										
1	1, 1, 1	0.53	0.008‖	1.390 (0.855)	0.202	3.795 of 7.849	3.118 of 5.876	1.962 of 3.400		
2	2, 1, 1	0.59	< 0.001‖	1.687 (0.861)	0.021	0.527 of 0.800	0.039 of 0.054	0.012 of 0.016		
3						0.018 of 0.021	0 of 0	0 of 0		
5¶									6.3	0.44
Dose No.										
1	**2, 2, 2**	0.177	0.397	0.195 (0.396)	0.129	0.354 of 4.364	0.540 of 6.222	0.689 of 7.133		
2	3, 2, 2	0.317	0.183	0.889 (0.616)	0.107	1.060 of 6.583	1.313 of 7.511	1.363 of 6.977		
3	2, 2, 1	0.144	0.136	0.552 (0.585)	0.092	2.334 of 4.023	0.785 of 1.223	0.584 of 0.836		
4	1, 1, 1	0.088	0.080	0.276 (0.502)	0.214					
5	3, 2, 1	0.284	0.073	1.250 (0.765)	0.071					
6	2, 1, 1	0.115	0.052	0.758 (0.728)	0.052					
6#									21.4	4.01
Dose No.										
1	**1, 1, 1**	0.166	0.450	0.543 (0.669)	0.373	1.203 of 7.997	1.884 of 11.432	2.205 of 12.12		
2	2, 1, 1	0.269	0.283	0.987 (0.743)	0.225	2.665 of 5.810	1.601 of 3.196	1.351 of 2.435		
3	2, 2, 1	0.383	0.136	0.758 (0.625)	0.062	0.514 of 0.970	0.024 of 0.041	0.005 of 0.009		

NOTE. Bold font indicates the true maximum tolerated sequence for each scenario. A total of 10,000 trials were simulated for each scenario. The true toxicity probabilities of three prespecified dose levels in the first cycle are presented for each scenario. The toxicity probability increases by 10% per cycle at the same dose level. Probability of toxicity per cycle is the average of the toxicity probabilities of the three doses in the sequence.
Abbreviations: DLT, dose-limiting toxicity; MTS, the highest sequence of three doses at which the probability of toxicity per cycle is no more than or close to a pre-specified target response rate p_1, given that the probability of toxicity at each cycle is also no more than or close to p_1; TPC, toxicity per cycle.

*True toxicity probability of first cycle: 0.15, 0.3, 0.45; $p_T = 0.3$.

†True toxicity probability of first cycle: 0.15, 0.27, 0.45; $p_T = 0.3$.

‡True toxicity probability of first cycle: 0.08, 0.16, 0.24; $p_T = 0.3$.

§True toxicity probability of first cycle: 0.48, 0.66, 0.84; $p_T = 0.3$.

‖Particularly for scenario 4, because of the over toxicity of its first dose, all the sequences selected at the end of the simulated trial are listed.

¶True toxicity probability of first cycle: 0.08, 0.16, 0.58; $p_T = 0.3$.

#True toxicity probability of first cycle: 0.15, 0.46, 0.53; $p_T = 0.3$

Table 4. Simulation Results for Scenarios 1 to 6 With a Sample Size of 30 Patients and Three Patients per Cohort

Scenario	Dose Sequence Selection					No. of DLT of No. Patients			Patients Treated Above MTS (%)	Probability of Early Termination (%)
	Dose Sequence	Probability of TPC	Selection (> 5%)	No. of DLTs, mean (SD)	Allocation (%)	Cycle 1	Cycle 2	Cycle 3		
1*									5.9	4.35
Dose No.										
1	**2, 2, 1**	0.270	0.199	0.643 (0.630)	0.103	1.489 of 9.908	2.962 of 17.945	3.810 of 20.929		
2	2, 2, 2	0.331	0.194	0.365 (0.481)	0.039	3.918 of 13.066	3.355 of 10.224	2.798 of 7.712		
3	2, 1, 1	0.216	0.166	0.894 (0.769)	0.179	2.783 of 6.175	0.421 of 0.853	0.141 of 0.265		
4	1, 1, 1	0.166	0.152	0.548 (0.676)	0.215					
5	3, 2, 1	0.320	0.083	1.141 (0.810)	0.084					
6	3, 1, 1	0.266	0.064	0.878 (0.767)	0.049					
7	3, 2, 2	0.381	0.062	0.890 (0.691)	0.032					
2†									8.6	4.05
Dose No.										
1	**2, 2, 2**	0.298	0.242	0.329 (0.470)	0.051	1.345 of 8.947	2.765 of 16.679	3.600 of 19.847		
2	2, 2, 1	0.250	0.181	0.633 (0.632)	0.108	3.578 of 13.205	3.327 of 11.209	2.844 of 8.677		
3	2, 1, 1	0.206	0.127	0.887 (0.773)	0.161	3.154 of 7.039	0.590 of 1.182	0.234 of 0.437		
4	1, 1, 1	0.166	0.116	0.547 (0.677)	0.193					
5	3, 2, 2	0.358	0.091	0.849 (0.685)	0.044					
6	3, 2, 1	0.310	0.090	1.141 (0.811)	0.092					
7	3, 1, 1	0.266	0.057	0.881 (0.765)	0.049					
3‡									0	0.50
Dose No.										
1	3, 3, 3	0.265	0.412	0.290 (0.454)	0.082	0.372 of 4.677	0.668 of 7.507	0.992 of 10.226		
2	3, 3, 2	0.233	0.247	0.606 (0.631)	0.142	1.225 of 7.741	2.437 of 13.786	2.745 of 14.173		
3	3, 2, 2	0.203	0.148	0.604 (0.631)	0.128	4.191 of 17.471	2.244 of 8.581	1.585 of 5.462		
4	2, 2, 2	0.177	0.091	0.192 (0.394)	0.071					
4§									11.9	99.98
Dose No.										
1	1, 1, 1	0.530	< 0.001‖	1.391 (0.85)	0.107	4.024 of 8.377	3.263 of 6.185	2.079 of 3.585		
2						0.554 of 0.841	0.044 of 0.062	0.012 of 0.016		
3						0.019 of 0.022	0 of 0	0 of 0		

(Continued on following page)

Table 4. Simulation Results for Scenarios 1 to 6 With a Sample Size of 30 Patients and Three Patients per Cohort (Continued)

Scenario	Dose Sequence Selection					No. of DLT of No. Patients			Patients Treated Above MTS (%)	Probability of Early Termination (%)
	Dose Sequence	Probability of TPC	Selection (> 5%)	No. of DLTs, mean (SD)	Allocation (%)	Cycle 1	Cycle 2	Cycle 3		
5¶									5.0	0.45
Dose No.										
1	**2, 2, 2**	0.177	0.673	0.193 (0.395)	0.229	0.374 of 4.726	0.816 of 9.194	1.082 of 11.220		
2	3, 2, 2	0.317	0.154	0.850 (0.620)	0.125	2.752 of 17.189	3.210 of 18.251	3.342 of 17.262		
3	2, 2, 1	0.144	0.051	0.644 (0.58)	0.113	4.622 of 7.987	1.553 of 2.444	0.973 of 1.394		
6#									16.8	4.35
Dose No.										
1	**1, 1, 1**	0.166	0.502	0.536 (0.665)	0.401	2.404 of 16.105	3.880 of 23.47	4.611 of 25.223		
2	2, 1, 1	0.269	0.276	0.941 (0.746)	0.245	5.126 of 11.149	2.768 of 5.499	2.031 of 3.675		
3	2, 2, 1	0.383	0.108	0.742 (0.628)	0.049	1.002 of .898	0.034 of 0.056	0.008 of 0.012		

NOTE. Bold font indicates the true maximum tolerated sequence for each scenario. A total of 10,000 trials are simulated for each scenario. The true toxicity probabilities of three prespecified dose levels in the first cycle are listed are presented for each scenario. The toxicity probability increases by 10% per cycle at the same dose level. Probability of toxicity per cycle is the average of the toxicity probabilities of the three doses in the sequence.
Abbreviations: DLT, dose-limiting toxicity; MTS, the highest sequence of three doses at which the probability of toxicity per cycle is no more than or close to a pre-specified target response rate p_T, given that the probability of toxicity at each cycle is also no more than or close to p_T; TPC, toxicity per cycle.

*True toxicity probability of first cycle: 0.15, 0.3, 0.45; $p_T = 0.3$.

†True toxicity probability of first cycle: 0.15, 0.27, 0.45; $p_T = 0.3$.

‡True toxicity probability of first cycle: 0.08, 0.16, 0.24; $p_T = 0.3$.

§True toxicity probability of first cycle: 0.48, 0.66, 0.84; $p_T = 0.3$.

‖Particularly for scenario 4, because of the over toxicity of its first dose, all the sequences selected at the end of the simulated trial are listed.

¶True toxicity probibility of first cycle: 0.08, 0.16, 0.58; $p_T = 0.3$

#True toxicity probability of the first cycle: 0.15, 0.46, 0.53; $p_T = 0.3$

the cohort size times (the number of doses + 1). The sample size of the motivating trial is limited because of the rareness of the disease and the desire to complete the trial in 2 years. With larger sample sizes, one can adopt dose-insertion methods[23,24] to allow insertion of new doses when necessary.

Lastly, a formal personalized dosing policy would follow the spirit of Lee et al[17] and use model-based inference on multiple cycles. To maintain simplicity in practice, BaSyc estimates an MTS with fixed dose levels at the end of the trial. The estimated MTS could be used as a guiding point to develop a final policy. For example, one may not allow a higher dose level than that in MTS for any cycle of the treatment in the policy.

DOI: https://doi.org/10.1200/PO.18.00020
Published online on ascopubs.org/journal/po on October 9, 2018.

AUTHOR CONTRIBUTIONS

Conception and design: Emily Curran, Yuan Ji

Data analysis and interpretation: Jiaying Lyu, Yuan Ji

Manuscript writing: All authors

Final approval of manuscript: All authors

AUTHORS' DISCLOSURES OF POTENTIAL CONFLICTS OF INTEREST

The following represents disclosure information provided by authors of this manuscript. All relationships are considered compensated. Relationships are self-held unless noted. I = Immediate Family Member, Inst = My Institution. Relationships may not relate to the subject matter of this manuscript. For more information about ASCO's conflict of interest policy, please refer to www.asco.org/rwc or ascopubs.org/po/author-center.

Jiaying Lyu
No relationship to disclose

Emily Curran
Consulting or Advisory Role: Shire
Research Funding: Gilead Sciences, Merck
Travel, Accommodations, Expenses: Gilead Sciences

Yuan Ji
Stock and Other Ownership Interests: Laiya Consulting

Affiliations

Jiaying Lyu, School of Public Health, Fudan University, Shanghai, People's Republic of China; **Emily Curran** and **Yuan Ji**, The University of Chicago, Chicago; and **Yuan Ji**, NorthShore University HealthSystem, Evanston, IL.

REFERENCES

1. Storer BE: Design and analysis of phase I clinical trials. Biometrics 45:925-937, 1989
2. O'Quigley J, Pepe M, Fisher L: Continual reassessment method: A practical design for phase 1 clinical trials in cancer. Biometrics 46:33-48, 1990
3. Ji Y, Liu P, Li Y, et al: A modified toxicity probability interval method for dose-finding trials. Clin Trials 7:653-663, 2010
4. Ji Y, Wang SJ: Modified toxicity probability interval design: A safer and more reliable method than the 3 + 3 design for practical phase I trials. J Clin Oncol 31:1785-1791, 2013
5. Guo W, Wang SJ, Yang S, et al: A Bayesian interval dose-finding design addressing Ockham's razor: mTPI-2. Contemp Clin Trials 58:23-33, 2017
6. Sverdlov O, Wong WK, Ryeznik Y: Adaptive clinical trial designs for phase I cancer studies. Stat Surv 8:2-44, 2014
7. Stock W, Luger SM, Advani AS, et al: Favorable outcomes for older adolescents and young adults (AYA) with acute lymphoblastic leukemia (ALL): Early results of US Intergroup trial C10403. Blood 124, 2014
8. Wahed AS, Tsiatis AA: Optimal estimator for the survival distribution and related quantities for treatment policies in two-stage randomization designs in clinical trials. Biometrics 60:124-133, 2004

9. Zhao Y, Zeng D, Socinski MA, et al: Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. Biometrics 67:1422-1433, 2011
10. Murphy SA, Bingham D: Screening experiments for developing dynamic treatment regimes. J Am Stat Assoc 104:391-408, 2009
11. Simon R, Freidlin B, Rubinstein L, et al: Accelerated titration designs for phase I clinical trials in oncology. J Natl Cancer Inst 89:1138-1147, 1997
12. Legedza AT, Ibrahim JG: Longitudinal design for phase I clinical trials using the continual reassessment method. Control Clin Trials 21:574-588, 2000
13. Zhang J, Braun TM: A phase I Bayesian adaptive design to simultaneously optimize dose and schedule assignments both between and within patients. J Am Stat Assoc 108:892-901, 2013
14. Doussau A, Asselain B, Le Deley MC, et al: Dose-finding designs in pediatric phase I clinical trials: comparison by simulations in a realistic timeline framework. Contemp Clin Trials 33:657-665, 2012
15. Doussau A, Thiébaut R, Paoletti X: Dose-finding design using mixed-effect proportional odds model for longitudinal graded toxicity data in phase I oncology clinical trials. Stat Med 32:5430-5447, 2013
16. Yin J, Qin R, Ezzalfani M, et al: A Bayesian dose-finding design incorporating toxicity data from multiple treatment cycles. Stat Med 36:67-80, 2017
17. Lee J, Thall PF, Ji Y, et al: Bayesian dose-finding in two treatment cycles based on the joint utility of efficacy and toxicity. J Am Stat Assoc 110:711-722, 2015
18. Fernandes LL, Taylor JM, Murray S: Adaptive phase I clinical trial design using Markov models for conditional probability of toxicity. J Biopharm Stat 26:475-498, 2016
19. Ji Y, Li Y, Nebiyou Bekele B: Dose-finding in phase I clinical trials based on toxicity probability intervals. Clin Trials 4:235-244, 2007
20. Goodman SN, Zahurak ML, Piantadosi S: Some practical improvements in the continual reassessment method for phase I studies. Stat Med 14:1149-1161, 1995
21. Bekele BN, Li Y, Ji Y: Risk-group-specific dose finding based on an average toxicity score. Biometrics 66:541-548, 2010
22. Bekele BN, Thall PF: Dose-finding based on multiple toxicities in a soft tissue sarcoma trial. J Am Stat Assoc 99:26-35, 2004
23. Hu B, Bekele BN, Ji Y: Adaptive dose insertion in early phase clinical trials. Clin Trials 10:216-224, 2013
24. Guo W, Ni Y, Ji Y: TEAMS: Toxicity- and efficacy-based dose insertion design with adaptive model selection for phase I/II dose-escalation trials in oncology. Stat Biosci 7:432-459, 2015

Appendix

Notation

Suppose under consideration is a total of I candidate dose levels and J treatment cycles, indexed by $i = 1,2,\ldots,I; j = 1,2,\ldots,J$. Denote p_{ij} the unknown probability of toxicity associated with the *i-th* dose in *j-th* cycle. The toxicity is assumed to be monotonic increasing with the dose level: $p_{1j} \le p_{2j} \le \ldots \le p_{Ij}$, for $j = 1,2,\ldots,J$. Let x_{ij} and n_{ij} be the number of toxicity events and total patients treated at dose i in cycle j, respectively. Let k index the cohort, $k = 1,2,\ldots,K$, where K is the maximum number of cohorts that will be recruited to the trial.

During dose-cycle finding (not for maximum tolerated sequence [MTS] selection), we use simple and independent beta priors across doses and cycles and an independent binomial likelihood as a working model for the estimation of p_{ij}. This inference is conducted to guide dose-cycle finding during the course of patient enrollment and assignment. The application in Bayesian adaptive dose-cycle finding (BaSyc) is detailed here.

Dose Continuation Rule

For cohort k that has been treated at dose i in cycle j, the dose continuation rule (DCR) for cohort continuation (whether the cohort can continue treatment at cycle $(j + 1)$) is described as follows:

- If cohort k has finished the last cycle of therapy (ie, $j = J$), cohort k exits the trial.
- If cohort k has not finished all cycles of therapy (ie, $j < J$), and if $\Pr\{p_{1,j+1} > p_T | data\} > \xi$, where ξ is a prespecified cutoff probability and often set to be close to 1 (eg, $\xi = 0.95$), stop cohort k for any additional treatment, and patients in cohort k exit the trial.
- Otherwise, continue this cohort k to the next cycle $(j + 1)$.

This algorithm states that when the probability that toxicity of the lowest dose level (dose level 1) in cycle $(j + 1)$ is higher than the target p_T does not exceed a large value, cohort continuation is allowed. The posterior probability $\Pr\{p_{1,j+1} > p_T | data\}$ is calculated based on a simple beta/binomial model. Let $x_{1,j+1}$ and $n_{1,j+1}$ be the cumulative number of patients (across all the previous and current cohorts) with dose-limiting toxicity (DLT) events and treated at dose level 1 at cycle $(j + 1)$, respectively. Then assume $x_{1,j+1} \mid n_{1,j+1} \sim Binomial(n_{1,j+1}, p_{1,j+1})$, $p_{1,j+1} \sim Beta(1,1)$ and the posterior distribution $p_{1,j+1} \mid data \sim Beta(1 + x_{1,j+1}, 1 + n_{1,j+1} - x_{1,j+1})$. This beta distribution is used to calculate the posterior probability.

Dose Allocation Rule

If cohort k in cycle j can proceed to the next cycle, the dose assigned to the next cycle $(j + 1)$ is dependent on the cumulative data from both the current cycle j and the next cycle $(j + 1)$. We consider three subrules, denoted as dose allocation rule 1 (DAR-1), DAR-2a, and DAR-2b. Details are as follows.

- As a safety and default rule, the patients in the first cohort (ie, $k = 1$) are always treated at the lowest dose $(i = 1)$ in all J cycles.
- DAR-1: For a newly enrolled cohort (except the first cohort), DAR-1 provides the dose level in the first cycle and follows the same decision from the modified toxicity probability interval 2 (mTPI-2) design.[5] That is, the dose level is decided based on the toxicity data from the current dose used to treat patients in the first cycle. Such data accumulate all the toxicity outcomes from previously treated cohorts at this dose.

Then, consider a cohort of patients that has just completed treatment in the first cycle at a dose level. Suppose the cohort will move to the second cycle for additional treatment. Note that some patients in the cohort might experience DLT outcomes in the first cycle, and the others might not. As a personalized strategy and to protect patient safety, we split the cohort into two subcohorts: the patients with DLTs and those without. The dose allocation decisions are made separately for these two subcohorts, using the following DAR-2a and DAR-2b rules:

- DAR-2a: For the subcohort of patients without DLTs after the first cycle treatment, the dose level of the second cycle for this (sub)cohort depends on two types of data: 1) the toxicity data ($data_1$) at the dose ($dose_1$) in cycle 1 where the cohort has been treated, and 2) the toxicity data ($data_2$) at the dose ($dose_2$) in cycle 2 where the previous cohort has been treated. Here, the previous cohort refers to the cohort of patients enrolled just before the current cohort, and the toxicity data or the data (also hereinafter) refer to the data from all enrolled patients treated at the corresponding dose. DAR-2a applies the mTPI-2 design[5] to decide the dose level in cycle 2. Specifically, given $data_1$ and $dose_1$, the mTPI-2 design will output a candidate dose level i_1, and given $data_2$ and $dose_2$, mTPI-2 will output another candidate dose i_2. Then DAR-2a selects the dose $i = min\{i_1,i_2\}$, the minimum of i_1 and i_2, as the final dose for treating the (sub)cohort in cycle 2.

- DAR-2b: For the second subcohort of patients with DLTs, de-escalate to a lower dose in the second cycle as follows. If the current dose is the lowest dose, we still stay at the same dose; otherwise, decrement by one level. Specifically, given the dose ($dose_1$) in cycle 1 and the candidate dose i_2 that mTPI-2 gives based on $data_2$ and $dose_2$ from cycle 2, DAR-2b selects the dose $i = min\{\max\{dose_1 - 1,1\},i_2\}$, the minimum of the dose 1 level lower than the current dose in cycle 1, and i_2 as the final dose for treating the cohort in cycle 2.

Note that because of the cohort split, the previously enrolled subcohorts may have been treated in cycle 2 at two dose levels, say $dose_{21}$ and $dose_{22}$. The outcome data from both subcohorts, denoted as $data_{21}$ and $data_{22}$, will be used to inform the cycle 2 dose for the subsequent cohort that will start from the first cycle. In particular, for each of the $data_{21}$ and $data_{22}$, the mTPI-2 design will give two candidate doses i'_{21} and i'_{22} as the potential dose for the subsequent cohort in cycle 2. Then, the minimum of i'_{21} and i'_{22} will be provided as the candidate dose, denoted as $i_2 = min\{i'_{21},i'_{22}\}$, for the subsequent cohort as the cycle 2 candidate dose. The two subcohorts will undergo the DCR and DAR separately hereafter until they either complete all three cycles of treatment or exit the trial. The same principle is applied to decide the dose level in cycle 3 if needed.

Walkthrough of the Hypothetic Trial

Table 2 summarizes how the BaSyc design guides the motivating dose-cycle trial with three doses and three cycles. In this trial, each cycle lasts 8 weeks, consisting of a 4-week treatment followed by 4-week washout and recovery. We walk through the first few cohorts in Table 2. First, according to the default safety rule, cohort 1 is treated at dose 1 (D1) in all three cycles. Therefore, at week 0 (not shown), cohort 1 is enrolled and assigned to D1. At week 8, cohort 1 completes treatment at D1 with 0 toxicity outcome (0 DLT of three patients). The DCR decision is to go to cycle 2. At this point, cohort 2 is enrolled and assigned to the next higher dose, dose 2 (D2), according to the DAR decision, because cohort 1 has not exhibited any toxicity. At week 16, cohort 1 completes treatment at D1 in cycle 2 with 0 toxicity outcome, and cohort 2 completes treatment at D2 in cycle 1 with 1 DLT outcome. For cohort 1, the DCR decision is to go to cycle 3 and to be treated at D1 (again as the default rule). For cohort 2, the DCR decision is to go to cycle 2, because the lowest dose in cycle 2 is considered safe. Cohort 2 is then split into two subcohorts according to the DAR. The single patient with DLT, patient 4, is de-escalated to D1, and the remaining two patients without DLTs, patients 5 and 6, stay at D2. Also, at week 16, cohort 3 is enrolled in the first cycle and treated at D2 according to the DAR. D2 is selected for cohort 3 because the three patients in cohort 2 exhibited one DLT event when treated at D2. Therefore, the DAR decides to continuously treat cohort 3 at the same dose, D2. This process continues for subsequent cohorts and cycles at weeks 24, 32, and so on, until all 15 patients (or five cohorts) have completed treatment at week 56.

Safety Rule Applied in BaSyc

For safety, the BaSyc design applies two additional safety rules throughout the entire trial:

- Rule 1 (early termination): For any cycle, if the toxicity probability of the lowest dose is higher than the target p_T with a large probability—that is, if $\exists j^* \in \{1,2,\ldots,J\}$ such that $\Pr\{p_{1j^*} > p_T | data\} > \xi$, where ξ is close to 1 (say = $\xi = 0.95$)—terminate the trial early.

- Rule 2 (dose exclusion): If the toxicity probability of dose i in cycle j is higher than the target p_T with a large probability, $\Pr\{p_{ij} > p_T | data\} > \xi$, exclude doses i and the higher doses from cycle j and higher cycles—that is, doses $\{i,i+1,\ldots,I\}$ will never be used in the trial in cycles $\{j,j+1,\ldots,\tilde{J}\}$ again.

The posterior probabilities in both rules are calculated based on the simple beta/binomial calculation proposed in the mTPI design.[3] Details are provided in Appendix under Dose Continuation Rule.

MTS Definition and Selection

After the enrollment and dose-cycle finding are completed, in the MTS selection, we use a frequentist and nonparametric probability model to make inference. Specifically, we essentially want to estimate the toxicity per cycle (TPC) q_s, which is the average of the three toxicities of the doses in a sequence s (described under Definition of MTS). We then construct an unbiased estimate of TPC $\tilde{q}_s$ (described under Estimation of $\tilde{q}_s$) and order-transform them based on isotonic regression (described under MTS Selection). Lastly, we select the MTSs in which the isotonically transformed $\tilde{q}_s$ are closest to p_T among all the cycles.

Definition of MTS.

In the motivation trial, there are three doses and three cycles, which amounts to a total of $S = 27$ sequences of dose cycles. We index a dose sequence by s, $s = 1,2,\ldots,S$, and label each sequence as (d_{s1}, d_{s2}, d_{s3}), where d_{s1}, d_{s2}, and d_{s3} denote the dose levels of sequence s in cycles 1, 2, and 3, respectively. For each sequence s, we define the cumulative dose level as the sum of the dose levels across three cycles (ie, $d^*_s = d_{s1} + d_{s2} + d_{s3}$). We define the ordering of the S sequences by the magnitude of the d^*_s values, with the lowest order being $d^*_1 = 3$ for sequence $s = 1$ of 1,1,1 and the highest order being $d^*_{27} = 9$ for sequence $s = 27$ of 3,3,3. For sequences with the same cumulative dose level, the sequence with higher dose in an early cycle is deemed lower. For example, sequence 2,1,1 is considered lower than 1,2,1. This tiebreaker is arbitrary and can be modified based on physician's input. For each sequence s, denote $p_{d_{sj}j}$ the probability of toxicity for the dose in cycle j of sequence s, $j = 1, 2, 3$. Let $q_s = \sum_{j=1}^{3} p_{d_{sj}j}/3$ be the average probability of toxicity of the sequence across three cycles (ie, the probability of TPC). Then, the MTS is the highest sequence whose q_s is closest to p_T. Also, for additional safety, we require that $p_{d_{sj}j}$ in the MTS for each cycle j must be no higher than p_T as well.

Estimation of q_s.

First, we select MTS among sequences that have been used in the trial and are not unacceptable. A sequence s is unacceptable if the toxicity probability of each dose in each cycle is higher than p_T with a large probability, defined as $\exists j = 1,2,3$, $\Pr\{p_{d_{sj}j} > p_T | Data\} > \xi$, where d_{sj} denotes dose levels {1,2,3} and ξ is close to 1, say $\xi = 0.95$. The posterior probability $\Pr\{p_{d_{sj}j} > p_T | Data\}$ is based on a simple beta/binomial hierarchic model, similar to the calculation under Dose Continuation Rule.

Among all the safe sequences that satisfy the condition above, we want to select the highest sequence with overall q_s close to p_T, subject to the toxicity probability of the same dose level nondecreasing across cycles. Suppose Data = $\{(x_{sj}, n_{sj})$, $s = 1,2,\ldots,S$, $j = 1,2,3\}$ and denote x_{sj} the number of patients with DLTs among the total number of patients n_{sj} treated at dose d_{sj} in cycle j of sequence s. An unbiased estimate of q_s is the average toxicity rate defined as $\tilde{q}_s = \sum_{j=1}^{3} x_{sj} / \sum_{j=1}^{3} n_{sj}$.

It is obvious that $E(\tilde{q}_s) = q_s$ when n_{sj} is the same across cycles j. Recall n_{sj} is the number of patients completing all three cycles of treatment after the dose sequence s. So n_{sj} is equal across cycles for a sequence. Therefore, we use $\tilde{q}_s$ as an unbiased estimator to estimate q_s.

MTS selection.

One could select the sequence that has the smallest value of $|\tilde{q}_s - p_T|$. However, the problem with that approach is that because of random chance, a sequence with a higher order (and therefore higher toxicity) might have a smaller $\tilde{q}_s$ than a lower-ordered sequence. For example, it is possible that the toxicity data for sequence 1,1,1 include only one toxicity event in all three cycles, whereas the data for sequence 2,1,1 include no toxicity event. Then, if $n_{sj} = 3$, the $\tilde{q}_s$ value for sequence 1,1,1 is 1/9 and 0 for sequence 2,1,1. However, sequence 2,1,1 is in theory a more toxic sequence, because it uses dose level 2 in cycle 2 instead of dose level 1. Therefore, selecting the sequence based on the distance of $|\tilde{q}_s - p_T|$ is not valid, because we would select 1,1,1 as the MTS, not 2,1,1 in this case.

To address this issue, we propose an isotonic transformation and calculate transformed $\tilde{q}_s$ values using the pool adjacent violator algorithm (PAVA)[20] just like in the mTPI design[3,4] and mTPI-2 design.[5] Here, we still use the sequence-specific data for the isotonic transformation. However, the sequences are partially ordered with ties, and therefore, PAVA cannot be directly applied, because it requires simple ordering of the sequences. For example, sequences 1,2,1, and 2,1,1 have the same order because both of their d^*_s values equal 4. To solve this issue, we propose a simple solution. When there is a tie between m sequences, there are $m!$ possible simple orderings among them. Then, assuming with probability $1/m!$ that one of the orderings, say ordering w, is true, we apply PAVA to the following ordering w of $\tilde{q}_s$, $w = 1,2,\ldots,m!$. We obtain $\hat{q}_{sw}$ as the PAVA-transformed value of $\tilde{q}_s$ given the w-th simple ordering. We then select the dose sequence with the smallest difference $|\hat{q}_{sw} - p_T|$ as the estimated MTS for the w-th simple ordering. Let us denote the estimated MTS D_w. Then out of all the $m!$ possible simple orderings, we obtain $m!$ estimated MTS, denoted as $\{D_1, D_2, \ldots, D_{m!}\}$. It is likely that these $D_{m!}$ MTSs also have ties, because the same sequence can be selected under different simple orderings. Therefore, let $\{D_1^*, D_2^*, \ldots, D_L^*\}$ $(L \le m!)$ be the unique sequences that are selected as the estimated MTS, and let $\{f_1, f_2, \ldots f_L\}$ be the corresponding frequencies of selection, with f_l being the number of times sequence l is selected as the estimate MTS, $l = 1,2,\ldots,L$. Then, finally, the sequence $D_{l^*}^*$ that has nonincreasing dose levels across three cycles and the largest frequency f_{l^*} is selected as the estimated MTS. We require that the dose levels of MTS are nonincreasing because of the assumption that the toxicity rate of the same dose increases over cycles. Therefore, the dose levels of the MTS cannot be increasing, or else the lower dose in the early cycle can always be replaced by the higher dose in the later cycle. When there are ties between f_{l^*} for multiple sequences, select the sequence with the lowest order as a conservative choice. If the orders are tied, then select the sequence with the lowest dose in the first cycle, the lowest dose in the second cycle if first-cycle dose is tied, or the lowest dose in the third cycle if first- and second-cycle doses are tied.

Table A1. Percentage of Dose Allocation Decisions for BaSyc Design Across 10,000 Simulated Trials

Decision	$y/n \in (0, p_T - \varepsilon_1)$	$y/n \in (p_T - \varepsilon_1, p_T + \varepsilon_2)$	$y/n \in (p_T + \varepsilon_2, 1)$
De-escalate	0	0	75.72
Stay	15.41	1	24.28
Escalate	84.59	0	0

NOTE. The true toxicity probabilities of three prespecified dose levels in the first cycle are 0.15, 0.3, and 0.45 (scenario 1). The target toxicity probability $p_T = 0.3$, and the choices of ε_1 and ε_2 are both 0.05. Symbol y denotes the number of dose-limiting toxicities, and n denotes the number of patients treated.
Abbreviation: BaSyc, Bayesian adaptive dose-cycle finding.

Simulation Results

Main results.

In scenario 1, sequence 1,1,1 is selected with the highest frequency (24.2%), and sequences 2,1,1 and 2,2,1 come in second and third, with a frequency of 19.7% and 19.0%, respectively. Sequence 2,2,1 is the true MTS. Sequence 2,2,1 is not selected with the highest frequency because of a small sample size (n = 15) and because the first dose has a nontrivial toxicity probability (0.15). This makes it difficult to escalate to dose 2 quickly. As summarized in Table 4, with a larger sample size (n = 30), sequence 2,2,1 gets selected the most frequently. Most patients are treated at doses with toxicity probabilities smaller than p_T. Because dose level 3 is overly toxic, few patients are allocated to it at any cycle. These results show that BaSyc is strong in protecting patient safety.

Scenario 2 presents a similar pattern of dose-toxicity relationship to scenario 1, except that the toxicity probability of dose level 2 is now below $p_T = 30\%$. The true MTS is 2,2,2, which is only selected 14.8% of the time with the fourth frequency. The same reason as scenario 1 is held.

Scenario 3 assumes that all the dose levels are safe with toxicity probabilities below p_T. With merely 15 patients, the BaSyc design is able to escalate quickly and select the sequence 3,3,2 and true MTS 3,3,3 with the highest and second highest probability of 20.9% and 20.4%, respectively. This suggests that the BaSyc design is nimble and can reach high dose-cycle combinations even with small sample sizes.

Scenario 4 presents an opposite case to scenario 3, because all the doses are overly toxic with toxicity probabilities above p_T. We can see that 99.21% of the simulated trials are terminated early without selecting any dose sequences as the MTS. On average, approximately eight patients are treated in the trial at low doses. This scenario demonstrates the ability of the BaSyc design to stop the trial quickly in case of excessive toxicity in all dose sequences.

Scenarios 5 and 6 reflect situations where there is a leap in toxicity probabilities across the three doses. In scenario 5, the leap lies between doses 2 and 3, with the first two doses below p_T. Therefore, the true MTS is 2,2,2. In scenario 6, the jump locates between doses 1 and 2, with 1,1,1 being the true MTS. The selection frequency of the true MTS is 39.7% for scenario 5 and 45.0% for scenario 6. In both scenarios, the MTS receives the most patients on average, whereas overly toxic doses receive much fewer patients. In scenario 6, 21.4% of patients (approximately three patients) are treated above the MTS because of the high probability of toxicity of the second dose level 0.46.

Sensitivity to sample size.

To evaluate the effect of sample size on the performance of BaSyc, we implement the BaSyc design for the same scenarios with a larger sample size of 30. Results are listed in Table 4. Across the scenarios, we see an increase of power in selecting the true MTS. The safety of the BaSyc design is again highlighted. For example, in scenario 4, toxic doses are assigned with a small number of patients, and trials are terminated early.