Statistica Sinica **32** (2022), 1983-2005
doi:https://doi.org/10.5705/ss.202020.0248

SEMIPARAMETRIC DOSE FINDING METHODS FOR PARTIALLY ORDERED DRUG COMBINATIONS

Matthieu Clertant, Nolan A. Wages and John O'Quigley

Université Sorbonne Paris Nord, University of Virginia and University College London

Abstract: We investigate a statistical framework for Phase I clinical trials that test the safety of two or more agents in combination. For such studies, the traditional assumption of a simple monotonic relation between the dose and the probability of an adverse event no longer holds. Nonetheless, the dose toxicity (adverse event) relationship does obey an assumption of partial ordering in that there will be pairs of combinations for which the ordering of the toxicity probabilities is known. Some authors have considered how to best estimate the maximum tolerated dose (a dose providing a rate of toxicity as close as possible to some target rate) in this setting. A related and equally interesting problem is to partition the two-dimensional dose space into two sub-regions: doses with probabilities of toxicity lower and greater than the target. We carry out a detailed investigation of this problem, using the recently presented semiparametric dose finding method as the theoretical framework. This results in a number of proposals, one of which can be viewed as an extension of the product of independent beta probabilities escalation (PIPE) method. We derive useful asymptotic properties, which also apply to the PIPE method when it is seen as a special case of the more general method given here. Simulation studies provide added confidence concerning the good behavior of the operating characteristics.

Key words and phrases: Bayesian method, dose-finding design, partial ordering, phase I clinical trials, semiparametric method.

1. Introduction

The importance of multi-agent Phase I trials in drug development has grown in recent years. The practical benefits of drug combinations are numerous: several modes of action can be combined, or the negative side effects of one drug can potentially be attenuated by the presence of a second compound. The aim of Phase I oncology trials is to find one or more maximum tolerated dose (MTD) combinations that have a probability of toxicity as close as possible to some threshold α, specified in advance by clinicians (common values are 20%, 25%, and 33%). Algorithmic designs remain a popular approach to identifying the

Corresponding author: Matthieu Clertant, LAGA, LabEx Inamex, Universite Paris Sorbonne Nord, 93430 Villetaneuse, France. E-mail: clertant@math.univ-paris13.fr.

MTD within a discrete set of levels, for both single-agent trials (Storer (1989)) and dual-agent trials (Huang et al. (2007)). These designs make no appeal to any model, and the escalation/de-escalation rules are determined as a function of the most recent set of observations. The simplicity of these methods leads to their frequent implementation in practice. Algorithmic designs have a Markov property, sometimes referred to in this context as a lack-of-memory property. These methods are not efficient and lack desirable statistical properties, such as almost sure convergence. Many model-based designs have been suggested, but their widespread adoption in practice remains an unaccomplished goal. This is particularly true in dual-agent trials because of the complexity and lack of interpretability of some of the proposed models, as well as operational difficulties that can discourage clinical investigators. Model-based designs come under two headings: parametrics and nonparametric.

The design and analysis of dual-agent trials can be carried out based on the work of Wages, Conaway and O'Quigley (2011), the partial-ordering continual reassessment method (poCRM) and an extension of the CRM by Wang and Ivanova (2005); the work of Thall et al. (2003) with a six-parameter logistic-type model and the work of Braun and Wang (2010) who use beta distributions and log-linear models. The second one is the class of nonparametric models. For dual-agent trials, Mander and Sweeting (2015) proposed the product of independent beta probabilities escalation design (PIPE), and Lin and Yin (2017) propposed the Bayesian Optimal Interval (2d-BOIN). The primary focus of many methods for dual-agent trials has been to find a single MTD for recommendation in Phase II studies. However, an arguably more relevant task is centered around the idea that, for combination studies, multiple, essentially equivalent MTDs may exist. The collection of these MTDs form a maximum tolerated contour (MTC) in two-dimensional space. In this case, the primary objective of the study may be to identify multiple MTDs for further testing, and investigators may desire a trial designed to meet this objective. This is particularly so as it is becoming common to study more than a single dose in subsequent dose expansion cohorts (Iasonos and O'Quigley (2016, DEC)).

Our motivation here is to describe a journey that starts with estimation of the MTD and ends with estimation of the MTC. Our signposts are built from illustrations and our ultimate purpose is to allow deeper study of the problem of early-phase dose finding for combinations of treatments. We set out on our journey (Section 2) via a look at the landscape, a look at the context in which these problems arise, including an actual trial carried out at the University of Virginia School of Medicine, and a particular hypothesized set of rates of DLT.

Table 1. Scenario T corresponds to a hypothetical pattern of toxicity probability in a study of Vinblastine and C6 Ceramide NanoLiposome (CNL). The 5×4 grid of dose combinations is fixed and chosen by the investigator. Dose combinations from the minimal set are shown in bold (see Definition 2).

Toxicity probabilities (%) of dose combinations		Scenario T Drug A1: CNL				
	Level (mg/m^2)	*54*	*81*	*122*	*183*	*215*
	2.25	**24**	29	34	59	70
Drug A2:	*1.5*	13	**18**	**22**	49	53
Vinblastine	*0.75*	7	9	**16**	35	42
	0.375	2	5	9	**21**	30

Next stop (Section 3) outlines the tools that we use in the more usual setting of an MTD target, and Section 4 describes how to solve MTC estimation. Dose finding in two dimensions is necessarily more complex, and Section 4.2 examines how actual experimentation can be implemented. Although sample sizes will not typically be large, large sample theory is helpful in providing confidence in how well the designs should work in practice (Section 5). In particular, Theorem 1 shows that, for all the methods, the almost sure convergence to the MTD is equivalent to the almost sure convergence to the MTC. Section 6 discusses a simulated comparison with an available method (PIPE), and Section 7 re-examines our original illustration (Section 2) and studies it in greater detail.

The online Supplementary Material contains a table of notation, a discussion of the dose allocation strategies, how to calibrate the methods, numerical experiments in the MTC setting, and proofs of the theorems and other properties.

2. Context and Motivation

We describe a Phase 1 clinical trial on drug combinations that was carried out at the University of Virginia School of Medicine. The trial involved C6 Ceramide NanoLiposome (CNL) and Vinblastine in patients with relapsed/refractory acute myeloid leukemia and patients with untreated acute myeloid leukemia. A grid of five by four dose-combinations is considered: five levels for CNL (from 54 to 215 mg/m^2), and four levels for Vinblastine (from 0.375 to 2.25 mg/m^2). The target rate α was taken to be 20%, and the maximum sample size was fixed at 60 patients. This particular application allows us to demonstrate the workings of such a Phase I drug combinations design, and to introduce some of the essential ideas developed in this paper. Table 1 describes a hypothesized dose toxicity function, called Scenario T.

Let dose $d = (i, j)$ represent the combination of the ith dose of an agent A1 and the jth dose of an agent A2, with $i \in \{1, \ldots, I\}$ and $j \in \{1, \ldots, J\}$. In Scenario T, I and J are respectively equal to five and four and dose $(2, 3)$ corresponds to 81 mg/m^2 of CNL and 1.5 mg/m^2 of Vinblastine. Let D be the set of all dose combinations: $D = \{1, \ldots, I\} \times \{1, \ldots, J\}$. The sequence $(X_n, Y_n)_{n\in\mathbb{N}}$ is a sample of observations. At step n, corresponding to the nth patient enrolled in the trial, the variable X_n is the dose selected among the $I \times J$ combinations. The variable Y_n is the observed binary response at this dose taking values $\{0, 1\}$, 1 for a dose limiting toxicity (DLT), and 0 otherwise. Each combination of drug d is associated with a toxicity probability P_d.

Assumption 1. $\forall n \in \mathbb{N}, \quad \mathbb{P}(Y_n = 1 | X_n = d) = P_d$.

In Scenario T, $P_{(2,3)}$ is equal to 0.18. Note that the ranges $\{1, \ldots, I\}$ and $\{1, \ldots, J\}$ are ordered in terms of the probability of toxic response. Such an assumption is common in Phase I trials for cytotoxic agents. The ordering of the marginal doses induces a partial ordering on the full range of doses D. The sign $<$ or $\leq$ will be used for the total ordering on $\mathbb{R}$ and the partial ordering on the set of doses D: $(i, j) \leq (r, s)$ if and only if $i \leq r$ and $j \leq s$.

Assumption 2. $(i, j) < (r, s) \Rightarrow P_{(i,j)} < P_{(r,s)}$.

In Scenario T, dose $(2, 3)$ is ordered with dose $(3, 3)$, as are their respective toxicity probabilities, 18% and 22%. The ordering assumption is fundamental to extracting information on a dose d from observations on neighboring doses. In particular, given any dose d, the partial ordering allows a decomposition of the set D into four subsets: $D = \{d\} \cup \mathcal{A}_d \cup \mathcal{B}_d \cup \mathcal{C}_d$, where $\mathcal{A}_d = \{d' \in D : d' > d\}$, $\mathcal{B}_d = \{d' \in D : d' < d\}$, and $\mathcal{C}_d = D \setminus (\mathcal{A}_d \cup \mathcal{B}_d)$. Those combinations where d belongs to the set $\mathcal{A}_d$ are associated with toxicity probabilities higher than P_d. The set $\mathcal{B}_d$ contains those combinations of doses below d associated with toxicity probabilities lower than P_d. The dose combinations in $\mathcal{C}_d$ are not ordered with dose d, which means that no prior assumption of order exists between their respective probabilities of toxicity. In Scenario T, the set $\mathcal{C}_{(2,3)}$ contains doses associated with toxicity probabilities below or above $P_{(2,3)} = 0.18$, for example, $P_{(3,1)} = 0.09$ and $P_{(1,4)} = 0.24$. Not unlike the single-agent setting, the primary objective of a Phase I combination study is often to identify a single MTD for further testing in the Phase II setting.

Numerous existing designs try to determine which combination d^* among those available in D has a probability of toxicity closest to the target toxicity rate α so that $d^* = \arg\inf_{d\in D} |P_d - \alpha|$. For Scenario T, the MTD, dose $(4, 1)$,

is associated with a toxicity probability 0.21, just above the threshold, $\alpha = 0.2$, fixed by the investigators. However, in the context of dual-agent dose finding, the notion of a single MTD is less clear cut. Indeed, aiming for the MTD is justified by an assumption of an ordering on the probabilities of efficacy, corresponding to an analogous ordering of the doses. For a single agent (full ordering), choosing the MTD amounts to maximizing the efficacy probability (probability of observing a sufficient therapeutic effect) under the constraint of being close to some chosen threshold of toxicity. This is no longer true in the dual-agent setting. Indeed, there exist situations in which a particular combination, less toxic than the MTD, may be more efficacious overall (Karapanagiotou et al. (2012)). In Scenario T, let the efficacy probabilities of doses $(4, 1)$ and $(2, 3)$ be equal to 0.10 and 0.90, respectively. Such a situation is coherent with the partial ordering on D, because dose $(2, 3)$ belongs to $\mathcal{C}_{(4,1)}$. In this case, the MTD $(4, 1)$ is not particularly advantageous. As a result, a dual-agent trial may aim to locate more than one MTD, forming an MTC (a cut or a contour) across the dose surface comprised of multiple combinations with similar acceptable toxicity profiles. The objective of the trial then becomes one of determining an MTD curve c^* consisting of a set of combinations with toxicity probabilities close to α. In Scenario T, the MTD contour is represented by the line splitting the set of doses in two parts: the dose combinations associated with toxicity probabilities below the target α and those associated with toxicity probabilities above α. The minimal set related to this specific contour contains the closest doses from the line according to the partial ordering (see doses with bold toxicity probabilities in Table 1). This set is defined more formally in Section 4.

However, a design aiming for a single MTD will be of interest in some situations. This is the case if there are too few patients available to determine a complete contour of the MTD (20 or 30 patients in a grid of five by four), or if the investigators have some strong prior assumptions or requirements on the set of potential outcomes. We present an example of such an assumption. In the context of CNL and Vinblastine combinations, the investigators may want to reach a dose combination with a level of Vinblastine at least equal to 1.5mg/m^2, because this agent has previously shown positive results in term of efficacy at this level. The general structure of the semiparametric class of dose finding methods is described by a hierarchical model, where the first level deals with the goal of the study - the doses themselves - and the second level is devoted to the dose-response curve. Here, the goal of the study is to find a single MTD. The first level of our model is then a distribution Π on the MTD itself. This distribution is updated in the Bayesian setting after each patient (see Section 3 for further

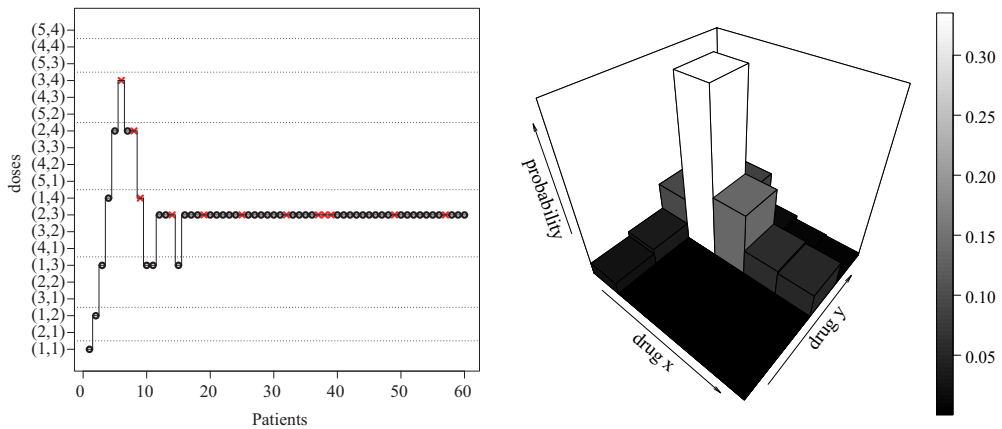

Figure 1. A simulated trial under scenario T using poSPM and the final posterior Π_{60}, which means the probabilties of each dose combination for being the MTD conditional on observations on 60 patients. Red cross: dose limiting toxicity (DLT); black circle: no toxicity response.

details).

In Figure 1, the posterior distribution on the MTD at the end of the trial is summarized by a histogram. Throughout the trial, the next dose combination is selected as the one maximizing the posterior Π according to the observations. The prior distribution on the range of doses D is a matrix of positive weights summing to one. It can easily be tuned to allow a progressive exploration in the range of doses or to correspond to expert knowledge: here, if possible, the investigators prefer to increase the level of Vinblastine until 1.5 mg/m^2 (see Section S5.2 in Supplementary Material for the exact prior distribution). This is what is done in the trial. At the beginning, the method recommends increasing the dose levels of Vinblastine while keeping the CNL level at 54 mg/m^2. After 15 patients, the method stays at dose $(2, 3)$. This dose is the one recommended at the end of the trial: nine patients among 48 treated at this dose experienced a DLT (rate $= 18.75\%$) accounting for 33.5% of the posterior mass. According to the posterior distribution, the two other most probable dose combinations are $(1, 4)$ and $(3, 3)$ which, respectively, account for 13.4% and 13.7% of the posterior mass. Note that semiparametric methods can be easily adjusted either to meet a specific criteria or to work under general conditions with uninformative priors. The analysis based on the semiparametric methods for scenario T continues in Section 7.

The trial in Figure 1 shows the classical behavior of a model-based design that tends to converge to a single MTD (Bayesian model in Section 3). However,

a more useful goal may be to explore the MTD curve in order to determine multiple safe dose combinations associated with potentially different probabilities of efficacy (Section 4).

3. Semiparametric Methods and a Single MTD Target

We briefly recall the semiparametric dose finding method introduced in (Clertant and O'Quigley (2017, SPM)), and then exploit these same ideas to tackle the problem of partial ordering. We refer to this as poSPM. The way in which we extend the simpler to the more complex situation is to build a model around the non-ordered dose combinations. This extension corresponds to a non-informative prior, structured in such a way that no distinction can be made between two doses d_1 and d_2 on the basis of observations on a dose d_3 that is not ordered with respect to d_1 and $d_2 : d_3 \in \mathcal{C}_{d_1}$ and $d_3 \in \mathcal{C}_{d_2}$.

Consider the set $S \subset [0,1]^D$, having elements $P = (P_d)_{d\in D}$. Then, P_d provides the probability of toxicity for scenario P at dose d. The set S is partitioned according to the different values (dose levels) that θ can assume.

$$S = \bigcup_{\theta\in D} S_\theta, \quad \text{where} \quad S_\theta \subset \{P \in S \mid \forall d \in D,\ |P_\theta - \alpha| \leq |P_d - \alpha|\}\,.$$

S breaks down into $I \times J$ classes indexed by θ, where every scenario from the same class has the same MTD. Any P belonging to S_θ is associated with the MTD at dose θ. For more discussion on this topic, refer to (Clertant and O'Quigley (2017, Sec. 4)). The definition of S_θ used in practice is expressed in Assumption 3 below. For a fixed $P \in S$, the likelihood that we can associate with the history $(X_1^n, Y_1^n) = (X_k, Y_k)_{1\leq k\leq n}$ is:

$$\mathcal{L}(P, X_1^n, Y_1^n) = \prod_{k=1}^{n} (P_{X_k})^{Y_k} (1 - P_{X_k})^{1-Y_k} = \prod_{d\in D} P_d^{n_d^1} (1 - P_d)^{n_d^0}\,, \tag{3.1}$$

where n_d^1 represents the number of DLTs occurring at dose d, and $n_d = n_d^1 + n_d^0$ is the number of patients that are treated at dose d. Let Π be the prior for the parameter θ. Then, Π is a probability measure on the discrete space D, that is a matrix of positive numbers with I rows and J columns summing to one. Each class S_θ is endowed with a prior denoted as $\Lambda_\theta = \Lambda(.|\theta)$. The family of distributions $(\Lambda_\theta)_{\theta\in D}$ and Π defines a prior on the whole set of scenarios. This prior can be updated sequentially as observations are made. The posterior Π_n on the parameter of the MTD, θ, is obtained by integration over the support S_θ

of the scenarios associated with this dose:

$$\Pi_n(\theta) = \Pi(\theta|X_1^n, Y_1^n) \propto \int_{S_\theta} \mathcal{L}(P, X_1^n, Y_1^n)\Lambda_\theta(dP)\,\Pi(\theta). \tag{3.2}$$

The family of distributions $(\Lambda_\theta)_{\theta\in D}$ plays a predictive role, and provides the basis that underlies the practical implementation of the adaptive process. We call this the prior model. In the sequential study, the next patient or cohort of patients is allocated to the dose having the highest probability of being the MTD, according to the posterior Π_n. The MTD and the probability of toxicity at any dose d, P_d, are estimated sequentially by

$$\hat{\theta}_n = \underset{\theta\in D}{\operatorname{argmax}}\ \Pi_n(\theta) \quad \text{and} \quad \hat{P}_d^{(n)} = \mathbb{E}_{(\Lambda\otimes\Pi)_n}[P_d] = \sum_{\theta\in D}\left[\int P_d\,\Lambda_{\theta,n}(dP_d)\right]\Pi_n(\theta). \tag{3.3}$$

The prior model is a family of $I\times J$ distributions on the set of scenarios. Moreover, each scenario can be viewed as a matrix with I rows and J columns having elements P_d, the toxicity probability at dose d. The following assumption expresses a simple structure for the prior Λ_θ. For this purpose, the Bernoulli parameter space is broken down into three sets. The high and low toxicity probabilities belong to $A = [\alpha+\epsilon, 1]$ and $B = [0, \alpha-\epsilon]$, respectively. The probabilities of toxicity in the neighborhood of the target α lie in the centered interval $I = [\alpha-\epsilon, \alpha+\epsilon]$. This interval results from our choice of ϵ, and can be made arbitrarily small. Indeed, it can be reduced to a single point, α, by choosing ϵ equal to zero. For the MTD parameter θ, the marginal Λ_θ^d distributes its probability mass according to the relative position of the doses d. The prior Λ_θ weights only those scenarios in which the MTD is θ. Given θ, Λ_θ^d is the prior of the toxicity probability at dose d itself. The position of d relative to θ, summarized in the following assumption, indicates that: (i) if $d=\theta$, as θ is the parameter indicating the MTD, the marginal weights are on the centered interval I; (ii) if $d>\theta$ (i.e., $d\in\mathcal{A}_\theta$), then Λ_θ^d puts its mass on the highest probability interval A; (iii) if $d<\theta$ (i.e., $d\in\mathcal{B}_\theta$), then Λ_θ^d puts its mass on the lowest probability interval B; (iv) if the dose d is non-ordered with θ (i.e., $d\in\mathcal{C}_\theta$), the marginal Λ_θ^d is uninformative, and the weights are on the whole interval $[0,1]$.

Assumption 3. *Λ_θ is a product of unidimensional distribution, i.e.*

$$\Lambda_\theta(dP) = \prod_{d\in D}\Lambda_\theta^d(dP_d),$$

and the support S_θ^d *of the distribution* $\Lambda_\theta^d(dP_d)$ *satisfies:* $d \in \mathcal{A}_\theta \Rightarrow S_\theta^d = A$; $d \in \mathcal{B}_\theta \Rightarrow S_\theta^d = B$; $d \in \mathcal{C}_\theta \Rightarrow S_\theta^d = [0,1]$; $d = \theta \Rightarrow S_\theta^d = I$.

The probabilities of toxicity in each set $\mathcal{A}_\theta$, $\mathcal{B}_\theta$, and $\mathcal{C}_\theta$ are conditionally independent once we have fixed S_θ. Marginally, of course, they are not independent. Thus any independency assumption fails to be valid when considering the whole model $\{(\Lambda_\theta)_{\theta \in D}, \Pi\}$. This is intuitively clear when a stochastic order is chosen for the marginal $(\Lambda_\theta^d)_{\theta \in D}$, at a certain dose $r : \forall r\,,\ d < d' \Rightarrow \Lambda_{d',n}^r \preceq \Lambda_{d',n}^r$. The coherence principle introduced by Cheung (2005) can be extended to the case of partial ordering.

Definition 1 (Partial ordering coherence)**.** A method $\mathcal{M}$ is coherent if the sequence of selected doses satisfies: $(X_n, Y_n) = (d, 0) \Rightarrow X_{n+1} \in D \backslash \mathcal{B}_d$ and $(X_n, Y_n) = (d, 1) \Rightarrow X_{n+1} \in D \setminus \mathcal{A}_d$.

This definition matches that in Cheung (2005) in the case of a single-agent trial, and can be seen as a logical extension of this criterion in the case of partial ordering. Note that the stochastic ordering implies that the poSPM is coherent.

Proposition 1. *Under Assumption* 3 *(above) and Assumption* 1 *in Supplementary Material, the poSPM is coherent.*

Working with a simple product of the marginals allows us to choose conjugate priors for the likelihood, in particular, beta distributions, of which the uniform distribution is a special case.

4. Identifying the Maximum Tolerated Contour (MTC)

4.1. Semiparametric model

In the setting of partially ordered doses, the semiparametric class of methods on the contour (poSPMc) takes its inspiration from the PIPE model (Mander and Sweeting (2015)) and a method for single agents (Clertant and O'Quigley (2018a,b)).

The set of dose combinations D can be naturally partitioned into two subsets, defined by the toxicity target α. The partition includes one set of dose combinations associated with toxicity probabilities below $\alpha : D^- = \{d \in D : P_d < \alpha\}$, and a second set of dose combinations associated with toxicity probabilities above $\alpha : D^+ = \{d \in D : P_d > \alpha\}$. We have: $D = D^- \cup D^+$. The maximum tolerated contour, also called the MTC curve, is a line separating the set D^- and D^+ (see Figure 2). The goal of poSPMc is to allocate patients at acceptable doses around the contour in order to better study dose combinations close to the

toxicity threshold. If these dose combinations are not ordered, then they do not share information in terms of the probability of efficacy. In general, a contour c is a line that separates the set of dose combinations into two ordered sets of doses: one set of dose combinations above the contour, $\mathcal{A}_c$, and another set of dose combinations below the contour, $\mathcal{B}_c$. For any given contour c, we note that

$$d \in \mathcal{A}_c \text{ and } d' \in \mathcal{B}_c \Rightarrow d > d' \text{ or } d \text{ and } d' \text{ non-ordered.}$$

In Figure 2, we can describe any contour as a polygonal chain tracing out a path from $(0.5, J + 0.5)$ to $(I + 0.5, 0.5)$, with steps of size 1 along the abscissa, and of size -1 on the ordinate axis. Only rightward and downward steps are permitted. As a result, for the range D containing $I \times J$ dose combinations, there exist $\binom{I+J}{I}$ possible contours c. Let C be the set of these contours, and $c^\star$ be the maximum tolerated contour (MTC) defined by the true unknown toxicity probabilities at each dose combination. Estimating the MTC may appear to be a considerable challenge, because the number of contours increases exponentially with the number of doses: 70 possible contours for a 4×4 grid and 962 for a 6×6 grid. Fortunately, each contour c is associated with a minimal set $\mathcal{M}_c$ of dose combinations on which we need to have observations in order, with probability one, to confirm or disprove that this contour can be identified as the MTC. The dose combinations belonging to $\mathcal{M}_c$ are the closest doses of the contour in our partially ordered set of doses D (see Figure 2). In the following definition, the minimum and maximum operators are applied to partially ordered sets.

Definition 2. The minimal set of dose combinations, $\mathcal{M}_c$, associated with the contour c is: $\mathcal{M}_c = \min \mathcal{A}_c \cup \max \mathcal{B}_c$.

For a grid with I doses for agent A_1 and J doses for agent A_2, the number of dose combinations belonging to a minimal set varies between 1 and $2 \times \min(I, J)$, except in the case $I = J$, where the maximum is equal to $2I - 1$. The average number of combinations inside a minimal set is equal to $2IJ/(I+J)$. If a minimal set contains only one dose combination, all of the toxicity probabilities $(P_d)_{d \in D}$ are on the same side of the target, either below or above. If a minimal set contains only two doses, it corresponds to a vertical or a horizontal contour.

Let $\mathcal{M}_{c^\star}$ be the minimal set associated with the MTC. Dose combinations in $\mathcal{M}_{c^\star}$ are not necessarily those closest to α in terms of toxicity probabilities. However, $\mathcal{M}_{c^\star}$ satisfies the following properties, which justify that a model for a dose combination study focuses on the contour or the associated minimal set.

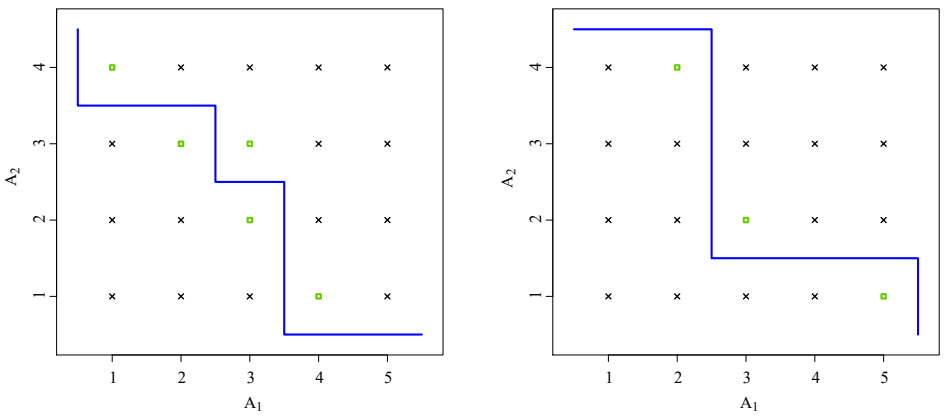

Figure 2. **Two contours called (a) and (b):** the green squares represent the dose combinations belonging to the minimal set. Contour (a) corresponds to that of Scenario T (Table 1).

Proposition 2.

1) *For all dose combinations* $d \notin \mathcal{M}_{c^\star}$, *there exists a dose* $d' \in \mathcal{M}_{c^\star}$ *ordered with* d, *such that:* $|P_d - \alpha| > |P_{d'} - \alpha|$.

2) *The MTD belongs to the minimal set associated with the MTC:* $d^\star \in \mathcal{M}_{c^\star}$.

The proof is immediate from Definition 2. The proposed model takes place in the setting of SPM; γ is the parameter of interest, the MTC itself. As in Section 3, S is a broad set of scenarios, and $P = (P_d)_{d \in D}$ is an element of S, with P_d the toxicity probability of the scenario P at dose d.

The set S is partitioned according to $\gamma : S = \bigcup_{\gamma \in C} S_\gamma$, where $S_\gamma = \{P \in [0,1]^D : d \in \mathcal{A}_\gamma \Rightarrow P_d > \alpha \text{ and } d \in \mathcal{B}_\gamma \Rightarrow P_d < \alpha\}$. Furthermore, S is partitioned into $\binom{I+J}{I}$ classes indexed by γ, where every scenario of the same class has the same MTC.

Let Π be the prior for the parameter γ, which means that Π is a probability measure on the discrete space C. Each class S_γ is endowed with a prior denoted as $\Lambda_\gamma = \Lambda(.|\gamma)$, and the family $(\Lambda_\gamma)_{\gamma \in C}$ is called the prior model. This family can be easily shaped by the following assumption and the use of truncated beta distributions for each marginal.

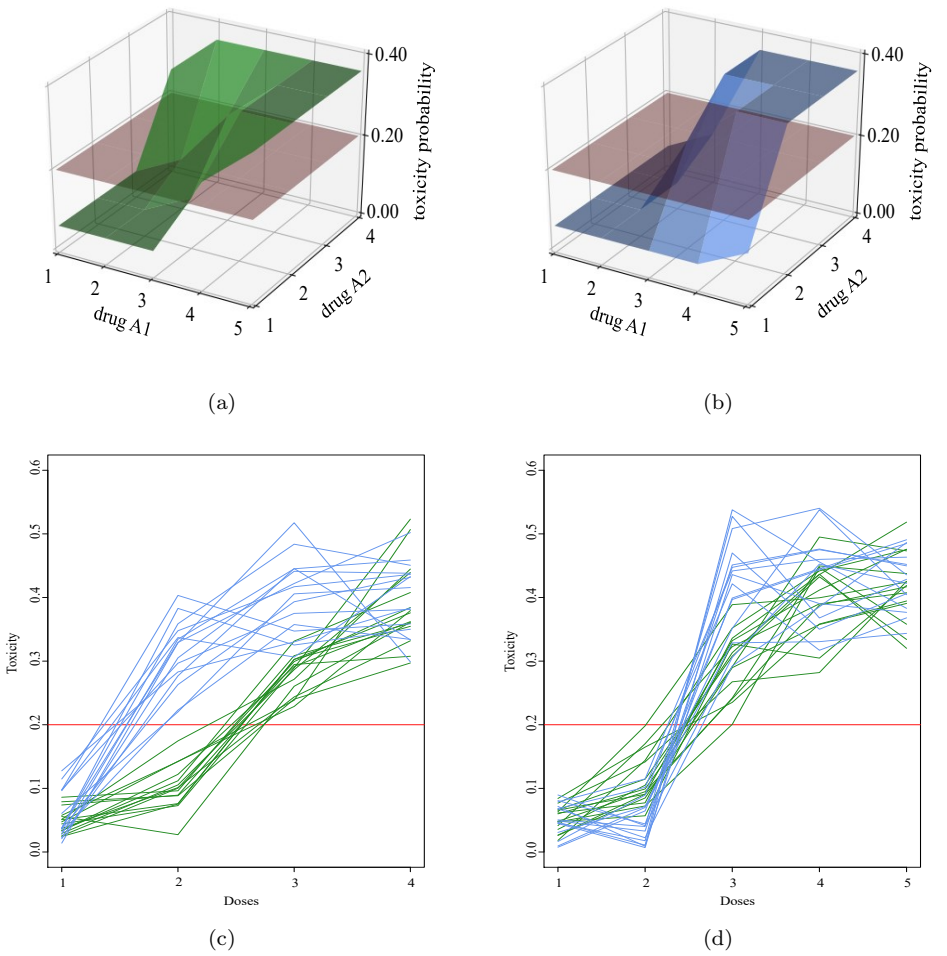

Figure 3. **Model of two competitive contours:** The blue and green models correspond to contours (a) and (b) (Figure 2). Figures (a) and (b) show the modes of $\Lambda^d_{(a)}$ and $\Lambda^d_{(b)}$, for each dose d in comparison to the target $\alpha = 0.20$. Figures (c) and (d) represent the sampling of marginal scenarios under these models when agents A1 and A2, respectively, are fixed at dose level 3.

Assumption 4. *Λ_γ is a product of one-dimensional distributions, i.e.*

$$\Lambda_\gamma = \prod_{d \in D} \Lambda^d_\gamma,$$

where Λ^d_γ is a prior for P_d, the toxicity probability at dose d. The support S^d_γ of the distribution Λ^d_γ satisfies: $d \in \mathcal{A}_\gamma \Rightarrow S^d_\gamma = [\alpha, 1]$; $d \in \mathcal{B}_\gamma \Rightarrow S^d_\gamma = [0, \alpha]$.

In Figure 3, the distribution Λ_γ is illustrated where γ represents the two contours in Figure 2. The parameters of these two distributions are those used in the numerical experiments. Note that, in Figure 3 (d) the marginal prior model $\Lambda_{(a)}$ is steeper than $\Lambda_{(b)}$; the doses (2,3) and (3,3) are in the minimal set for contour (a), but not for contour (b).

The family of distributions $(\Lambda_\gamma)_{\gamma\in C}$ and Π define a prior on the whole set of scenarios S. The posterior Π_n on the parameter of the MTC, given the data (X_1^n, Y_1^n), is:

$$\Pi_n(\gamma) \propto \int_{S_\gamma} \mathcal{L}(P, X_1^n, Y_1^n)\Lambda_\gamma(dP)\,\Pi(\gamma). \tag{4.1}$$

The current estimate of the contour is

$$\hat{\gamma}_n = \underset{\gamma\in D}{\operatorname{argmax}}\ \Pi_n(\gamma). \tag{4.2}$$

The algorithm used to allocate a patient (or the next cohort of patients) during a trial is based on the following two steps: (i) the prior Π_n is updated and the most probable contour $\hat{\gamma}_n$ is selected; and (ii) the dose for the next patient is selected using an allocation strategy in the most probable minimal set, according to our posterior: $X_{n+1} \in \mathcal{M}_{c^\star}$. Such an allocation strategy that spreads the observations inside the estimated minimal set is developed in Section 4.2.

The PIPE method (Mander and Sweeting (2015)) might be seen as the first example of using a semiparametric method on the contour. However, it has some strong limitations in terms of calibration, because it does not clearly specify priors on the MTC and the space family $(S_\gamma)_{\gamma\in C}$. For this method, the toxicity probability at each dose d is modeled by a weak beta prior with parameters a_d and b_d. For each contour, a posterior probability is obtained by multiplying the posterior of the interval $[0, \alpha]$ for every dose below the contour and the posterior of the interval $[\alpha, 1]$ for those doses above the contour. Let $B_I(a, b)$ be the truncated beta distribution with parameters a and b on the interval I, and let $B(x, a, b)$ be the incomplete beta function. In the SPM setting, the PIPE model is defined by:

$$\begin{cases} \Lambda_c \ \sim \prod\limits_{d\in\mathcal{B}_c} B_{[0,\alpha]}(a_d, b_d) \times \prod\limits_{d\in\mathcal{A}_c} B_{[\alpha,1]}(a_d, b_d) \\ \Pi(c) \propto \prod\limits_{d\in\mathcal{B}_c} B(\alpha; a_d, b_d) \times \prod\limits_{d\in\mathcal{A}_c} (1 - B(\alpha; a_d, b_d)). \end{cases}$$

Note that, for the PIPE method, the prior model $(\Lambda_c)_{c\in C}$ and the prior on the contour Π are defined by the $I \times J$ prior of the toxicity probabilities at each dose combination. The indirect calibration of the prior model and the prior on the set

of contours can undermine the performance of this method. In particular, the prior Π defined using the PIPE model is not chosen to allow a progressive exploration of the range of doses, but is dependent on the target toxicity probability α.

4.2. Allocation strategy

In a minimal set, many of the doses are not ordered and the question of allocation within this set needs to be addressed. This is not straightforward, and we do not exhaust all possibilities here. Several allocation strategies can be made available and any particular choice will result in some particular trial behaviour. The allocation strategy introduced here is a compromise between one that spreads experimentation "equally" and one that tests doses that will, in expectation, bring more information into the study. We also need to restrict experimentation to an area of the minimal set considered to be safe. After n observations, the next dose X_{n+1} is

$$X_{n+1} = \min_{d \in \mathcal{M}_{\hat{\gamma}_n} \setminus T_n} \frac{1}{k_d} \left\{ \left(\sum_{j \in \mathcal{B}_d \cup \{d\}} n_j^0 \right) \times H(\alpha|0) + \left(\sum_{j \in \mathcal{A}_d \cup \{d\}} n_j^1 \right) \times H(\alpha|1) \right\}, \tag{4.3}$$

where $H(p|q)$ is the entropy of p relative to q : $H(p|q) = -q \times \log(p) - (1 - q)\log(1-p)$, with the convention $\log(0) = -\infty$ and $0 \times (-\infty) = 0$, and $k_d = \#\mathcal{A}_d + \#\mathcal{B}_d + 1$. In Equation 4.3, the relative entropy, which corresponds to a quantity of information, is counter-balanced by k_d, the number of dose combinations ordered with d, that is, the number of dose combinations on which we aim to learn something subsequent to an observation on d.

The next dose X_{n+1} is selected from among the safe dose combinations in the estimated minimal set: $\mathcal{M}_{\hat{\gamma}_n} \setminus T_n$. The set of dose combinations on which we already have some evidence of overly toxicity is denoted T_n. It is estimated by using local Bayesian tests at each of the doses and the assumption of partial ordering. A dose d_0 is excluded from the study if

$$d_0 \in T_n = \{ d \in D : \exists\, d' \in H_n,\, d \geq d' \},$$

with $H_n = \{ d \in D : \mathbb{P}_{\mathcal{U}} \left[P_d > \theta_T | (n_d, n_d^1) \right] > \delta_T \}$, and $\mathcal{U}$ is the uniform prior on the parameter space $[0, 1]$.

This allocation strategy spreads the observations along the estimated contour, and tests slightly more often the doses in the middle of the grid. Because a DLT corresponds to a greater quantity of information, those doses with a

smaller number of toxicities are tried more often (for $\alpha = 0.25 : H(\alpha|1)/H(\alpha|0) = \log(\alpha)/\log(1-\alpha) \approx 4.82$). A longer description of this allocation strategy can be found in Supplementary Material (S3).

5. Large-Sample Theory

Two distinct and complementary asymptotic behaviors are described here. The first kind of large sample behavior is that of ϵ-sensitivity, introduced by Cheung (2011). This corresponds to the almost sure convergence to a dose combination for which the true probability of toxicity lies inside the interval $[\alpha-\epsilon;\alpha+\epsilon]$.

Definition 3. Let $\varepsilon \geq 0$ and $I = [\alpha-\varepsilon;\alpha+\varepsilon]$. We consider the set $\mathcal{E}(I,P)$ of the collection of dose combinations associated with a toxicity probability belonging to I, that is, $\mathcal{E}(I,P) = \{d \in D : P_d \in I\}$. A method is ε-sensitive if for all scenarios $P = (P_d)_{d\in D}$, such that $\mathcal{E}(I,P) \neq \emptyset$, we have

$$\mathbb{P}\left[\exists N\,,\, \forall n > N : X_n \in \mathcal{E}(I,P)\right] = 1.$$

The ε-sensitivity corresponds to strong consistency for a single dose associated with a toxicity rate close to the threshold α. This dose is not necessarily the MTD if there exist two or more doses in the interval I. The almost sure convergence to the MTD is obtained if the MTD is a unique dose in $\mathcal{E}(I,P)$. This assumption about the shape of the scenario is necessary in the light of the impossibility theorem of Azriel, Mandel and Rinott (2011). In our semiparametric method for a single MTD (poSPM), the interval I is a parameter that can be directly calibrated. This requires that we choose a small interval I, which could potentially fail to contain any of the toxicity probabilities associated with the given doses. For this reason, the complementary behavior of the ϵ-sensitivity introduced by Clertant and O'Quigley (2017) is extended to the case of partial ordering. We define as balanced behavior the almost sure convergence to a set of doses. More precisely, a sequence $(X_n)_{n\in\mathbb{N}}$ converges to a set B, denoted by $X_n \xrightarrow{S} B$, when

$$\max_{x\in B}\left(\liminf_{n\to+\infty}\delta(X_n,x)\right) = 0\,, \tag{5.1}$$

and where $\delta(.,.)$ is the Euclidean distance. In the case of a complete ordering, the sequence $(X_n)_{n\in\mathbb{N}}$ converges to the set consisting of the two consecutive doses on either side of the target α. This set of two doses corresponds to the definition of the minimal set for the maximum tolerated contour (see definition 2): $\mathcal{M}_{c^\star} = \max D^- \cup \min D^+$.

Definition 4. Let D be a range of doses. A method is balanced if, for all scenarios, we have

$$X_n \xrightarrow{S} M_{c^\star}\ ,\ a.s.$$

The doses recommended infinitely often by the current estimator of a balanced method are all the doses in $M_{c^\star}$, and only these doses. The following theorem completes in a natural way the impossibility theorem formulated by Azriel, Mandel and Rinott (2011).

Theorem 1. *For any allocation method, the three following statements are equivalent:*

(i) For all scenarios, there exists a statistic F on the sample (X_1^n, Y_1^n) such that $F(X_1^n, Y_1^n) \longrightarrow d^\star$, a.s.

(ii) For all scenarios, there exists a statistic F on the sample (X_1^n, Y_1^n) such that $F(X_1^n, Y_1^n) \longrightarrow c^\star$, a.s.

(iii) For all scenarios, there exists a set of dose combinations D' with $M_{c^\star} \subset D'$ such that the sequence of recommended doses satisfies

$$X_n \xrightarrow{S} D'\ ,\ a.s.$$

In particular, this basic theorem expresses a necessary condition for any method: the minimal set $M_{c^\star}$ is the smallest set of dose combinations on which we need to have observations in order to find almost surely the MTD and the MTC under general circumstances. Some assumptions are needed on our models in order to obtain the desired asymptotic behavior. The regularity assumption 2 (Section S8.1 in Supplementary Material) on the prior model of the poSPM and the poSPMc are readily met. The following assumption concerns the allocation strategy.

Assumption 5.

(a) The next dose combination X_{n+1} is selected from $\mathcal{M}_{\hat{\gamma}_n}$.

(b) All the doses in the limit supremum of $(\mathcal{M}_{\hat{\gamma}_n})_{n\in\mathbb{N}}$ are selected infinitely often:

$$\forall d \in \bigcap_{n\geq 1}\bigcup_{k\geq n} \mathcal{M}_{\hat{\gamma}_k},\ \left\{\sum_{i=1}^{n} \mathbb{1}_{\{X_n=d\}} \underset{n\to\infty}{\longrightarrow} \infty\right\},\ a.s.$$

Assumption 5 (b) means that if a contour is selected infinitely often by $\hat{\gamma}_n$, all the doses in the minimal set associated with this contour are explored infinitely

often (a.s). This is a necessary condition to obtain balanced behavior (a simple consequence of Theorem 1's proof). This constructive assumption is satisfied by the allocation strategy defined in Section 4.2.

Theorem 2.

(a) Suppose that P^T is the true scenario. Under assumptions 3 *and* 2 *(Section* S8.1 *in Supplementary Material), we have: if $\epsilon > 0$ and $\mathcal{E}(I, P^T)$ contains no couple of ordered doses then the poSPM is ε-sensitive; when $\epsilon = 0$, the poSPM is balanced.*

(b) Under assumptions 4, 5 *and* 2 *(Section* S8.1 *in Supplementary Material), the poSPMc is balanced.*

When $\epsilon = 0$, the poSPM is balanced, which means that the method explores all the dose combinations included in $\mathcal{M}_{c^\star}$ when the sample size becomes sufficiently large. However, this only describes the asymptotic behavior and, at finite sample sizes, the simulations will often converge to a single dose identified as the MTD (see Figure 1). Conversely, the balanced behavior of poSPMc at a finite sample size can be observed in Figure 4.

6. Numerical Experiments: MTC Setting

In the MTC setting, we compared the poSPM to the PIPE design in terms of each method's ability to identify and treat patients along a target MTC. Each trial targets a toxicity level of $\alpha = 0.20$, with a total sample size of 50 and one patient per cohort. The calibration of the two methods is fully described in the Supplementary Material S5.1 and S5.2.

Table 2 shows the operating characteristics of the 2 methods under the four scenarios shown in Table 1 (Section S1 in the Supplementary Material). As in the MTD setting, for each scenario, we report the percentage of patient allocation (experimentation %) and percentage of MTD selection (recommendation %) for doses contained within five different ranges of true toxicity probabilities. Because $\alpha = 0.20$, the target interval containing the true MTD is $[0.15, 0.25]$. In Scenarios 1 and 4, the poSPM and PIPE method have very similar operating characteristics with regard to recommending and experimenting at combinations in the target interval. The difference between the methods in Scenarios 1 and 4 can be observed in the intervals outside the target range. In both scenarios, when outside the target interval, the poSPM allocates to combinations below the target range, whereas the PIPE method tends to be more aggressive. This has an impact on the accuracy index (see Equation S4.1 in the Supplementary Material) for

Table 2. Experimentation and recommendation percentages for the poSPM and PIPE design. Column containing the MTD in bold. Scenarios 1 to 4 correspond to scenarios A, C, E, G in Mander and Sweeting (2015). They can be found in the Supplementary Material S.1.

		Experimentation %						
Scenario	Method	$[0, 0.10)$	$[0.10, 0.15)$	$[\mathbf{0.15}, \mathbf{0.25}]$	$(0.25, 0.30]$	$(0.30, 1]$	Acc	%DLTs
1	poSPM	9.0	25.8	**56.9**	7.7	0.4	0.35	16.9
	PIPE	2.8	7.5	**59.4**	26.1	4.2	0.39	22.0
2	poSPM	0.0	26.4	**45.0**	9.4	19.1	0.88	23.5
	PIPE	0.0	6.2	**21.2**	11.5	61.0	0.63	36.0
3	poSPM	12.2	27.0	**36.1**	22.8	1.9	0.38	18.0
	PIPE	0.0	6.6	**21.8**	11.5	60.2	0.33	35.9
4	poSPM	31.2	19.4	**32.4**	13.7	3.2	0.67	14.2
	PIPE	12.3	14.0	**34.0**	26.0	13.4	0.55	22.1
		Recommendation %						
		$[0, 0.10)$	$[0.10, 0.15)$	$[\mathbf{0.15}, \mathbf{0.25}]$	$(0.25, 0.30]$	$(0.30, 1]$	Acc	$\overline{N}_{MTD}$
1	poSPM	2.0	19.9	**71.6**	6.3	0.0	0.63	3.73
	PIPE	0.3	5.2	**70.3**	23.7	0.4	0.64	2.23
2	poSPM	0.0	17.8	**54.3**	12.8	15.0	0.92	2.60
	PIPE	0.0	8.8	**38.4**	17.4	35.3	0.84	1.44
3	poSPM	3.8	29.1	**36.6**	29.5	0.9	0.44	3.18
	PIPE	0.0	9.8	**37.7**	17.5	35.0	0.28	1.44
4	poSPM	20.1	22.6	**39.1**	16.8	1.2	0.77	4.2
	PIPE	17.4	23.5	**38.0**	19.1	2.0	0.77	2.77

experimentation in Scenario 4, with the poSPM yielding a larger value: 0.67 versus 0.55. For both scenarios, the PIPE induces a higher overall percentage of DLTs, on average.

In Scenario 2, the poSPM is the best performing method by all metrics. It demonstrates better performance from the viewpoints of experimentation (45.0% vs. 21.2%) and recommendation (54.3% vs. 38.4%) percentages in the target interval. The PIPE again tends to be more agressive, allocating 61% of patients to combinations with toxicity probabilities in the interval $(0.30, 1.00]$. The improvement of the poSPM over the PIPE is also reflected in the accuracy indices in Scenario 2. In Scenario 3, the methods have similar operating characteristics in terms of the recommendation percentage in the target interval (36.6% vs. 37.7%), although the poSPM maintains an advantage when considering experimentation and accuracy. Similarly to Scenario 2, the PIPE allocates 60.2% of patients to combinations with toxicity probabilities in the interval $(0.30, 1.00]$. The more aggressive behavior of the PIPE is further indicated in the overall percentage of DLTs observed, as this value is both higher than the corresponding value for poSPM, as well as larger than the target $\alpha = 0.20$ in all scenarios. Considering the average performance over the four scenarios, the poSPMc is the best method

Table 3. Experimentation and recommendation percentages for poSPMc, poSPM1, and poSPM2 on Scenario T. The column containing the MTD is shown in bold.

	Experimentation %						
Method	$[0, 0.10)$	$[0.10, 0.15)$	$[\mathbf{0.15}, \mathbf{0.25}]$	$(0.25, 0.30]$	$(0.30, 1]$	Acc	%DLTs
poSPMc	27.5	8.8	**47.3**	8.9	7.4	0.52	18.1
poSPM1	14.7	19.4	**55.7**	8.7	18.7	0.48	22.2
poSPM2	10.7	20.1	**42.9**	14.1	11.9	0.48	21.6
	Recommendation %						
	$[0, 0.10)$	$[0.10, 0.15)$	$[\mathbf{0.15}, \mathbf{0.25}]$	$(0.25, 0.30]$	$(0.30, 1]$	Acc	$\overline{N}_{MTD}$
poSPMc	13.9	8.4	**60.7**	11.5	5.4	0.63	4.31
poSPM1	5.4	2.2	**73.0**	9.1	10.1	0.67	1
poSPM2	1.7	18.2	**62.9**	11.5	5.6	0.65	1

by all the metrics. In particular, the poSPMc recommends a larger number of doses (3.43 doses vs. 1.97), while maintaining greater accuracy than the PIPE method (0.69 of accuracy recommendation vs. 0.63).

7. CNL and Vinblastine Combination: A Study of Scenario T

Three versions of our semiparametric class of methods are compared on Scenario T introduced in Table 1 in the context of the CNL and Vinblastine combination:

- **poSPMc** is the semiparametric method targeting the MTD contour (or the associated minimal set); the calibration is the same as that in Section 6,
- **poSPM1** is the semiparametric method targeting a single MTD, without including expert knowledge (uninformative prior); the calibration is the same as that in the Supplementary Material S.4,
- **poSPM2** is the semiparametric method targeting a single MTD that is calibrated in line with expert knowledge, as described in Section 2 ('if possible, the investigators prefer to increase the level of Vinblastine until 1.5 mg/m^2'), the calibration is the same as that in the Supplementary Material S4 except that the mass of the prior Π on doses $(i, j) \in \{2, \ldots, 5\} \times \{1, 2\}$ is divided by two.

In Table 3, the three methods are compared on Scenario T. The target α is 20% and 60 patients are enrolled in the study. We simulate 10,000 trials for each method. Cheung's accuracy metric is defined in Section S4. The metric $\overline{N}_{MTD}$ is the average number of dose combinations recommended at the end of the trial: 4.31 for the poSPMc (less than the number of doses in the minimal set

Table 4. Experimentation and percentage of trials recommending at least a dose in sets $H_1 = \{(2,3),(3,3)\}$, $H_2 = \{(3,2),(3,3)\}$, $H_3 = \{(4,1)\}$, and $H_4 = \{(1,4)\}$ for poSPMc, poSPM1, and poSPM2. Each set may be the only one containing dose combitions sufficiently efficacious and not associated with overly high toxicity probabilities.

		Experimentation %			
Scenario T	Method	H_1	H_2	H_3	H_4
	poSPMc	18.8	16.7	10.6	8.0
	poSPM1	32.6	37.2	4.3	2.9
	poSPM2	12.6	24.5	0.1	17.8
		% of trials Recommendation			
Scenario T	Method	H_1	H_2	H_3	H_4
	poSPMc	65.2	74.3	64.8	62.8
	poSPM1	44.9	43.6	6.5	3.75
	poSPM2	21.3	41.3	0.2	20.2

of Scenario T), and one for the other methods. For the final recommendation percentages in interval $[0.15, 0.25]$, the poSPM1 outperforms the poSPM2 and poSPMc. For the poSPM2, this is due to the steep toxicity profile around the MTDs in the middle of the grid, rather than around the dose (1,4), which is then more difficult to locate. The difference between the poSPMc and poSPM1 in terms of allocation and recommendation percentage in interval $[0.15, 0.25]$ is due both to the exploration behavior of poSPMc and the multiple recommendation at the end of the trial (it is more difficult to recommend 4 dose combinations in the interval $[0.15, 0.25]$ than it is to recommend one). However, the percentage of dose combinations recommended by the poSPMc that are inside the interval $[0.15, 0.25]$ is still satisfactory (60.7%).

In order to show the benefits of targeting the MTD curve (poSPMc), four hypotheses concerning the efficacy profile of CNL and Vinblastine are considered. The minimal set of Scenario T contains dose combinations (1,4), (2,3), (3,2), (3,3) and, (4,1). Note that the true MTD always belongs to the minimal set; for Scenario T, this is $(4, 1)$. The hypotheses rely on the set of doses with sufficient efficacy probabilities in the minimal set. In the four hypotheses, the only dose combinations that are sufficiently efficacious and not associated with overly toxicity probabilities belong to set $H_1 = \{(2,3),(3,3)\}$, $H_2 = \{(3,2),(3,3)\}$, $H_3 = \{(4,1)\}$, and $H_4 = \{(1,4)\}$, respectively. These four hypotheses are coherent with the partial order applied to the efficacy probabilities. They represent four possible profiles of efficacy. Note that a dose combination of CNL and Vinblastine is considered to have a sufficient efficacy profile if its efficacy rate exceeds that associated with competitive therapies. The aim of the Phase I design is to

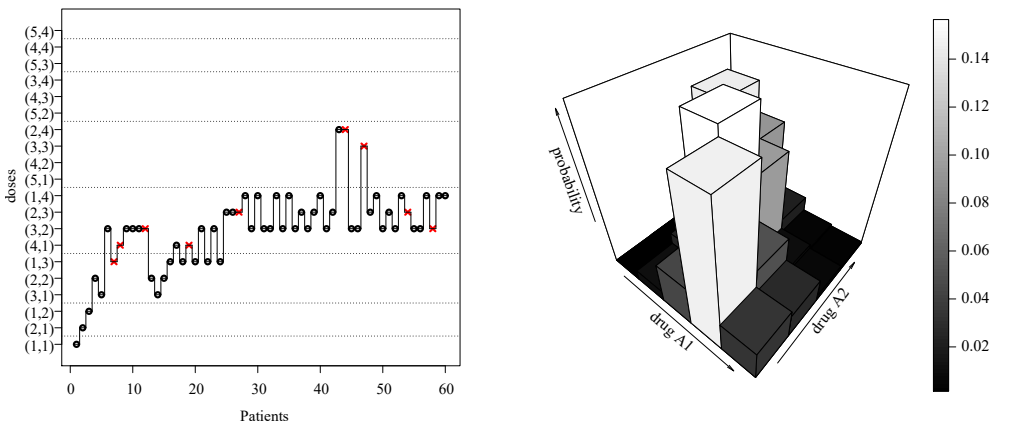

Figure 4. A simulated trial under scenario T using the poSPMc and the final posterior $\tilde{\Pi}_{60}$ on the dose combinations. Red cross: dose limiting toxicity; black circle: no toxicity response.

then recommend such a dose combination for further exploration in Phase II. In Table 4, percentages of trials (among the 10,000) which recommend a dose in the sets H1, H2, H3 and H4 are detailed for each of the three methods. The advantage of the method targeting the MTD curve over methods aiming for a single MTD is significant. For the poSPMc, the percentage of trials in which at least one dose with a sufficient efficacy profile is recommended is between 60% and 75% for all the hypotheses. The same metric goes from 2% to 45% for the poSPM1 and poSPM2. This does not mean that the two methods are not accurate, but rather that they are not equipped to deal with this situation.

The balanced behavior of the poSPMc is outlined using a trial presented in Figure 4. At the beginning of the trial, the method explores each of the dose combinations of the progressive inverted diagonals (e.g. (1,2) and (2,1)) until a first DLT is observed. Following the 20th patient, almost all the recommended dose combinations are in the minimal set of Scenario T.

Note that the poSPMc updates a distribution Π on the possible maximum tolerated contour. The method does not immediately lead to a posterior distribution on the MTD. However, because the model (Λ, Π) is a distribution on the range of scenarios, one can still calculate a distribution $\tilde{\Pi}$ on the MTD parameter by integrating out the toxicity probabilities conditional upon each possible value that could be assumed by the MTD. Instead of doing this complex calculation, we use the following formula which holds for all n:

$$\forall c \in C\,,\ \forall d \in D\,,\ \tilde{\Pi}_n \propto \frac{1}{\overline{\mathcal{V}_d}} \sum_{c \in \mathcal{V}_d} \Pi_n(c)\,, \tag{7.1}$$

where $\mathcal{V}_d$ is a neighborhood in terms of the contours for dose combination d : $\mathcal{V}_d = \{c \in C : d \in \mathcal{M}_c\}$, and $\overline{\mathcal{V}_d}$ is its cardinality. The following proposition shows that $\tilde{\Pi}_n$ is the desired result by considering a poSPM model $(\tilde{\Lambda}, \tilde{\Pi})$ for which the updating process is conjugate with the one of the poSPMc model.

Proposition 3. *For all $n \in \mathbb{N}$, the probability $\tilde{\Pi}_n$, defined by Equation 7.1, is the posterior on the set of dose combinations from the poSPM model $(\tilde{\Lambda}, \tilde{\Pi})$, where*

$$\tilde{\Pi}(d) \propto \frac{1}{\overline{\mathcal{V}_d}} \sum_{c \in \mathcal{V}_d} \Pi(c) \ \textit{and}\ \tilde{\Lambda}_d = \sum_{c \in \mathcal{V}_d} r_c^d \Lambda_c \ \textit{with}\ r_c^d = \frac{\Pi(c)}{\sum\limits_{c \in \mathcal{V}_d} \Pi(c)}\,. \tag{7.2}$$

Thus, there always exists a model $(\tilde{\Lambda}, \tilde{\Pi})$ of the poSPM setting such that, for all $n \in \mathbb{N}$, the distributions $\tilde{\Pi}_n$ correspond to the posterior of $\tilde{\Pi}$. This conjugacy property associates, with every poSPMc model on the MTD curve, a poSPM model for a single MTD. However, this association is not trivial, and appealing to an allocation strategy for the poSPMc allows the design to explore the full MTD curve. At the end of a trial with the poSPMc, the final recommendation relies on the posterior $\tilde{\Pi}_n$. In Figure 4, more than 60% of the posterior mass of $\tilde{\Pi}$ is on the minimal set of Scenario T.

Supplementary Material

The eight sections contain scenarios used in the simulation, a notation table, a numerical experiment in the MTD setting, the calibration methods, and the proofs of all the properties and theorems.

Ackowledgments

Dr. Clertant is supported by Investissements d'Avenir programme ANR-11-IDEX-0005-02 and 10-LABX-0017, Laboratoire d'excellence INFLAMEX.Dr. Wages is supported by National Institutes of Health Award R01CA247932. Additional philanthropic support was provided by Elizabeth Cronly to the Patients & Friends Research Fund of University of Virginia Cancer Center. Thanks to Dr. Roxane Duroux for her participation in the early phase of this project.

References

Azriel, D., Mandel, M. and Rinott, Y.(2011). The treatment versus experimentation dilemma in dose finding studies. *Journal of Statistical Planning and Inference* **141**, 2759–2768.

Braun, T. M. and Wang, S. (2010). A hierarchical Bayesian design for phase I trials of novel combinations of cancer therapeutic agents. *Biometrics* **66**, 805–812.

Cheung, K. (2005). Coherence principles in dose-finding studies. *Biometrika* **92**, 863–873.

Cheung, Y. K. (2011). *Dose Finding by the Continual Reassessment Method*. Chapman and Hall/CRC, Boca Raton.

Clertant, M. and O'Quigley, J. (2017). Semiparametric dose finding methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **79**, 1487–1508.

Clertant, M. and O'Quigley, J. (2018a). The role of minimal sets in dose finding studies. *Journal of Nonparametric Statistics* **30**, 1016–1031.

Clertant, M. and O'Quigley, J. (2018b). Semiparametric dose finding methods: Special cases. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **68**, 271–288.

Huang, X., Biswas, S., Oki, Y., Issa, J. and Berry, D. A. (2007). A parallel phase I/II clinical trial design for combination therapies. *Biometrics* **63**, 429–436.

Iasonos, A. and O'Quigley, J. (2016). Dose expansion cohorts in phase I trials. *Statistics in Biopharmaceutical Research* **8**, 161–170.

Karapanagiotou, E. M., Roulstone, V., Twigger, K., Ball, M., Tanay, M., Nutting, C. et al. (2012). Phase I/II trial of carboplatin and paclitaxel chemotherapy in combination with intravenous oncolytic reovirus in patients with advanced malignancies. *Clinical Cancer Research* **18**, 2080–2089.

Lin, R. and Yin, G. (2017). Bayesian optimal interval design for dose finding in drug-combination trials. *Statistical Methods in Medical Research* **26**, 2155–2167.

Mander, A. P. and Sweeting, M. J. (2015). A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Statistics in Medicine* **34**, 1261–1276.

Storer, B. E. (1989). Design and analysis of phase I clinical trials. *Biometrics* **45**, 925–937.

Thall, P. F., Millikan, R. E., Mueller, P. and Lee, S. (2003). Dose-finding with two agents in phase I oncology trials. *Biometrics* **59**, 487–496.

Wages, N. A., Conaway, M. R. and O'Quigley, J. (2011). Continual reassessment method for partial ordering. *Biometrics* **67**, 1555–1563.

Wang, K. and Ivanova, A. (2005). Two-dimensional dose finding in discrete dose space. *Biometrics* **61**, 217–222.

Matthieu Clertant

LAGA, LabEx Inflamex, Université Paris Sorbonne Nord, 93430 Villetaneuse, France.

E-mail: clertant@math.univ-paris13.fr

Nolan A. Wages

Translational Research and Applied Statistics, University of Virginia, Charlottesvile, USA.

E-mail: nwages@virginia.edu

John O'Quigley

Department of Statistical Science, University College London, London, WC1E 6BT, UK.

E-mail: john.oquigley@upmc.fr

(Received June 2020; accepted March 2021)