

Collection of Biostatistics Research Archive
COBRA Preprint Series

Year 2016

Paper 116

Semi-Parametrics Dose Finding Methods

Matthieu Clertant*

John O'Quigley†

*Cedar Sinai Medical Center, Cancer Center, Los Angeles, matthieu.clertant@cshs.org

†Universit'e Pierre et Marie Curie-Paris VI, Paris, France

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://biostats.bepress.com/cobra/art116>

Copyright ©2016 by the authors.

Semi-Parametrics Dose Finding Methods

Matthieu Clertant and John O'Quigley

Abstract

We describe a new class of dose finding methods to be used in early phase clinical trials. Under some added parametric conditions the class reduces to the family of continual reassessment method (CRM) designs. Under some relaxation of the underlying structure the method is equivalent to the CCD, mTPI or BOIN classes of designs. These latter designs are non-parametric in nature whereas the CRM class can be viewed as being strongly parametric. The proposed class is characterized as being semi-parametric since it corresponds to CRM with a nuisance parameter. Performance is good, matching that of the CRM class and improving on it in some cases. The structure allows theoretical questions to be more easily investigated and to better understand how different classes of methods relate to one another.

Semi-Parametric Dose Finding Methods

Clertant, M. and O'Quigley, J. ¹

Université Pierre et Marie Curie, Paris, France

Abstract: We describe a new class of dose finding methods to be used in early phase clinical trials. Under some added parametric conditions the class reduces to the family of continual reassessment method (CRM) designs. Under some relaxation of the underlying structure the method is equivalent to the CCD, mTPI or BOIN classes of designs. These latter designs are non-parametric in nature whereas the CRM class can be viewed as being strongly parametric. The proposed class is characterized as being semi-parametric since it corresponds to CRM with a nuisance parameter. Performance is good, matching that of the CRM class and improving on it in some cases. The structure allows theoretical questions to be more easily investigated and to better understand how different classes of methods relate to one another.

Keywords: adaptive designs; clinical trials; Continual Reassessment Method; dose finding methodology; hierarchical models; sequential designs.

Introduction

The importance of early phase dose finding studies - so called Phase I and Phase I/II clinical trials - is difficult to overstate. This is particularly so in oncology where it is believed that a significant number of the more than ninety percent of failed large scale randomized clinical trials can, to a more or lesser degree, be explained by an inefficient or an inaccurate early phase study. The recommended dose would have been either too high, and poorly tolerated, or too low and, in consequence, unable to provide an adequate anti-tumour response. It is widely recognized by statisticians and clinicians alike that the standard 3+3 dose finding design (Storer, 1989) widely employed in Phase I trials is fatally flawed and, in some sense, not fit for purpose. As a result the last twenty five years has seen considerable statistical research into early phase designs that are more efficient while simultaneously paying attention to the ethical constraints required in the running of any such trial.

Different approaches divide themselves into two classes; the first - examples include the 3+3 and the Rolling Six (Skolnik et al., 2008) - are called algorithmic designs since no modeling takes place and the escalation, de-escalation rules are determined solely as a function of some set of the most recent observations. They have a Markov property, sometimes referred to in this context as a lack-of-memory property. The second class of designs are called model-based designs. Their motivation is to impose greater structure on the observations in order to increase

¹Address for correspondence: Matthieu Clertant or John O'Quigley, LSTA, Université Pierre et Marie Curie Paris 6, 15-25, 4, place Jussieu 75005 Paris France
email: pr.clertant@upmc.fr ; john.oquigley@upmc.fr

the information obtained through sampling as well as to satisfy large sample convergence of the recommended dose to the true MTD. Statistical properties such as almost sure convergence are important in as much as, without such properties, it is difficult to feel confident in the solidity of any approach. At the same time, in real studies, sample sizes are often no more than 20 to 30 and so it is also crucial to have desirable finite sample properties such as coherence (Cheung, 2005). Simulations, across broadly varying situations, have been a useful help in developing methods. One clear advantage of the model-based designs is their ability to be generalized to deal with more complex situations such as group heterogeneity, combination therapies and toxicity attribution error.

The most well known of the model-based designs is the continual reassessment method (CRM) introduced by O’Quigley, Pepe and Fisher (1990). The method has been very successful but despite its now well known superior performance over the standard design, its use still lags behind that of the standard design. One explanation for this is that many clinicians are not at ease in using a method whose in-trial operating behaviour can not be immediately anticipated as well as the fact that help in the form of an able bodied statistician is recommended. But there are other statistical concerns (Azriel et al., 2011) that have led to the development of many competing model-based approaches. Among these are EWOC (Babb et al., 1998), mTPI (Ji et al., 2010) and BOIN (Liu and Yuan, 2015). Some authors have pointed out that the conditions for almost sure convergence in Shen and O’Quigley (1996) are very restrictive, and therefore not realistic. Cheung and Chappell (2002) and Azriel (2012) described ways to relax these assumptions but the concerns still remain.

The CRM is based on a strong parametrisation of the regression function, so much so that it is often described as an under parametrized model. Taking our cue from Cheung and Chappell’s work on the CRM, we introduce a semi-parametric characterization of the method (Section 2). This characterization can also be viewed as a hierarchical Bayesian model having as a first level the main parameter of interest, the MTD itself, and, on a secondary level, families of dose-toxicity curves constrained only by the location of the MTD. Within this framework we can characterize several other current methods. We immediately gain some theoretical advantages such as almost sure convergence to the Maximum Tolerated Dose (MTD) under weaker conditions than those currently admitted. It is also much easier to anticipate large sample behaviour in more general situations. The general structure allowed for by semi-parametric models enables deeper study of the various methods currently available. Perhaps no less importantly, we propose methods benefiting from improved asymptotics properties, computationally very fast, which, for small samples, obtain as good and sometimes better results than the CRM. Cheung and Chappell’s characterization for the MTD as an interval rather than a point paved the way to more realistic and achievable goals for dose finding studies. These ideas are fundamental to our development here and, within this context, we describe two central features of asymptotic behaviour. The first is that of being “sensitive” and the second that of being “balanced.” Under very wide and realistic dose-toxicity curves

we show that the semi-parametric method exhibits desirable large sample behaviour regarding these properties, and that this behaviour will indeed be reflected in commonly observed finite samples. For any given situation, the large sample behaviour of the semi-parametric method can be characterized by either almost sure convergence to the MTD or by infinite oscillation over two adjacent levels, the rates of visitation to either one of these levels being quantified according to Kullback-Leibler divergence.

A user friendly program in R is available at address https://github.com/MatthieuMC/SPM_project_01.

1 Context of model-based designs

1.1 Basic set-up and notation

The statistical purpose is to estimate the root of an unknown dose-toxicity regression function as observations are accumulated sequentially. The observations are the sequences: $(X_n, Y_n)_{n \in \mathbb{N}}$. At step n , the variable X_n is the dose selected from a range of available doses; $D = \{1, \dots, m\}$ and 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. The conditional distribution of Y_n given $X_n = d$ is Bernoulli with parameter β_d , which implies that at each dose is associated a probability of toxicity independent of the way in which patients are selected into the study. The range D has been chosen by clinical expertise so that the doses are ordered in terms of the probability of toxic response.

Assumption 1.1. $\forall n \in \mathbb{N}, \forall d \in D, \beta_d = \mathbb{P}(Y_n = 1 | X_n = d)$ with $\beta_1 < \dots < \beta_m$.

Estimating the root of the regression function enables us to determine which dose among those available in the range D suggests itself as having a probability of toxicity the closest to some maximum amount α chosen by the investigators. This dose, noted d^* , is called the MTD (maximum tolerated dose): $d^* = \arg \inf_{d \in D} |\beta_d - \alpha|$. As patients are included sequentially into the study, we suppose that all of the information contained in the sample, $(X_1^n, Y_1^n) = ((X_1, \dots, X_n), (Y_1, \dots, Y_n))$ is available to guide the selection of the dose X_{n+1} . The ethical constraints of the study imposed by the clinical team encourages us to use all the available information at each step in order to choose our best current estimate of the MTD as the dose to be given to the following patient. The following definition (Cheung, 2005) describes a property that any sensible design should have.

Definition 1.1. A method, \mathcal{M} , is said to be coherent if the selection of the next dose given the observed sample satisfies, for all $d \in D$ and $n \in \mathbb{N}$:

$$(X_n, Y_n) = (d, 0) \Rightarrow \mathcal{M}(X_1^n, Y_1^n) \geq d \quad \text{and} \quad (X_n, Y_n) = (d, 1) \Rightarrow \mathcal{M}(X_1^n, Y_1^n) \leq d,$$

where $\mathcal{M}(X_1^n, Y_1^n)$ denotes dose X_{n+1} given (X_1^n, Y_1^n) .

The following assumption restricts attention to design based estimates that are adapted to the accumulating observations.

Assumption 1.2. *The current estimator of the method satisfies: $\mathcal{M}(X_1^n, Y_1^n) \in \sigma(X_1^n, Y_1^n)$, where $\sigma(X_1^n, Y_1^n)$ is the sigma-algebra generated by the sample.*

Under this condition we are able to obtain classical asymptotic properties for frequentist estimators $\widehat{\beta}_{d,n}$ defined by:

$$\widehat{\beta}_{d,n} = \frac{n_d^1}{n_d^0 + n_d^1} = \frac{n_d^1}{n_d}, \text{ where } n_d^i = \sum_{j=1}^n \mathbb{1}_{\{X_n=d, Y_n=i\}}, i \in \{0, 1\}. \quad (1)$$

Lemma 1.1. *For all methods satisfying Assumption 1.2, we have:*

(i) *Law of large numbers: $\widehat{\beta}_{d,n} \xrightarrow[n_d \rightarrow +\infty]{} \beta_d$, a.s.*

(ii) *Law of the iterated logarithm: for all $\beta_d \in]0, 1[$, with $\sigma_d = \sqrt{\beta_d(1 - \beta_d)}$, we have*

$$\limsup_{n_d \rightarrow +\infty} \frac{\sqrt{n_d} (\widehat{\beta}_{d,n} - \beta_d)}{\sigma_d \sqrt{2 \log(\log(n_d))}} = 1 \quad \text{and} \quad \liminf_{n_d \rightarrow +\infty} \frac{\sqrt{n_d} (\widehat{\beta}_{d,n} - \beta_d)}{\sigma_d \sqrt{2 \log(\log(n_d))}} = -1, \text{ a.s.}$$

Proof. Part (i) is shown in Azriel, Mandel, and Rinott (2011, lemma 3) and Part (ii) in the supplementary material. □

This very general result provides no useful method in itself. Indeed, the event $\{n_d \rightarrow +\infty\}$ is random depending on the vector $\beta = (\beta_d)_{d \in D}$ and the chosen method. It is not immediately clear how to obtain a consistent estimator of the MTD because we do not wish for each dose to be observed infinitely often. A good dose finding method will be all the more desirable as it fulfills two criteria:

- (1) (TR, treatment): we would like the greatest possible number of patients to be treated at and close to the MTD during the study.
- (2) (PCS, percentage of correct selection): the method should lead us with high probability to a correct determination of the MTD.

Reconciling and jointly optimising these two criteria creates specific difficulties for dose finding studies. In this context it is worth recalling an impossibility theorem of Azriel et al. (2011) that throws a useful light on the asymptotic results obtained here. These authors have shown that no method exists that would, for all situations, allow the current estimator to be strongly consistent. Only particular configurations with respect to the employed method result in strong consistency.

Theorem 1.1. *Let \mathcal{M} be a method satisfying Assumption 1.2. A scenario β satisfying Assumption 1.1 exists such that:*

$$\mathbb{P}_\beta(\exists N : \forall n > N, \mathcal{M}(X_1^n, Y_1^n) = d^*) < 1.$$

Indeed, if the method recommends a single dose for n large enough, observations on competing doses will cease. The information we have at these doses is finished and, of course, we may then make an incorrect recommendation however large n .

Example 1.1. *Let $\beta = (0.05, 0.10, 0.20, 0.35, 0.55, 0.7)$ and $\alpha = 0.2$. Suppose that, for n large enough, the design selects only dose 2 and the results for dose 3 are 2 DLT among 5 observations: $\hat{\beta}_{3,n} = 0.4$. In such circumstances, for all dose $d \neq 2$, the law of large numbers will not apply to $\hat{\beta}_{d,n}$ because the dose d is not infinitely tested. The inaccuracy in the estimate of toxicity related to dose 3 will not be overcome by increasing sample size. However, note that the event $\{n_2 \rightarrow +\infty\}$ together with the assumption of monotonicity 1.1 allow us to eliminate almost surely dose 1 from the candidate doses for being the MTD.*

According to the Assumption 1.1, the two doses associated with toxic probabilities either side of the target dose α are consecutive. It would then appear desirable as a large sample property to concentrate experimentation on these doses. The class of methods introduced in this article arise in a natural way from a critical analysis of the asymptotic properties of the CRM (O’Quigley et al., 1990). This very general construct would allow us to include a wide range of, at first glance, diverse methods under a single general heading. This generalization opens the way to make progress on two fronts; that of critical evaluation of the overall strategy and that of more efficient parameterization of particular existing methods alongside their extensions when dealing with more complex clinical situations. In the rest of the article, we study this in relation to the CRM. The proposed parametrisation enables us to reproduce the global behavior of this method while obtaining better theoretical properties and allows us to escape those difficulties consequent on poor model specification (skeleton). In further unpublished work we study more deeply this generalization as it relates to the CCD (Ivanova et al., 2007), mTPI (Ji et al., 2010) and BOIN methods (Liu and Yuan, 2015), since the semi-parametric structure leads to immediate improvements in all of these designs.

1.2 Parametric methods: Continual Reassessment Method

In this paragraph, we recall the principle features of the continual reassessment method. We do this in a particular way which helps us to see how the new developments presented here sit quite naturally within the basic framework of the CRM. The CRM method works by approximating the dose-toxicity relationship $d \mapsto \beta_d$ by a family of continuous functions of a parameter a .

$$\begin{aligned} f : \mathcal{X} \times [A, B] &\rightarrow [0, 1] \\ (x, a) &\mapsto f(x|a), \end{aligned}$$

with $[A, B]$ a finite interval and \mathcal{X} a continuous set containing the range of doses D . We call the family of functions $(f(\cdot|a))_{a \in [A, B]}$ the model and the vector β the scenario (or reality). Note that $f(x|a)$ is the toxicity associated with the dose x , that is to say a model for the probability of toxicity at dose x under parameter a . The following algorithm describes the general working of CRM. In the Bayesian setting, G is the prior distribution of a and G_n its posterior given the observations $(X_1^n, Y_1^n) : G_n(da) \propto L_n(a) \times G(da)$, where $L_n(a) = \prod_{i=1}^n f(X_i|a)^{Y_i} (1 - f(X_i|a))^{1-Y_i}$.

Step 1. *Through the likelihood L_n , the current amount of information (X_1^n, Y_1^n) is used to update our estimate a_n of the parameter a or its posterior G_n .*

Step 2. *The estimator of the next dose, X_{n+1} is obtained as a function of a_n or G_n .*

Many possibilities are available for the second step. One approach is to calculate the estimators of toxicity at each dose $d \in D : \tilde{\beta}_{d,n} = f(d|a_n)$ or $\tilde{\beta}_{d,n} = \mathbb{E}_{G_n}[f(d, a)]$. The next dose is then the one whose estimated probability of toxicity is the closest to the desired target: $X_{n+1} = \arg \min_{d \in D} |\tilde{\beta}_{d,n} - \alpha|$. In order to set the context for the general semi-parametric model, we propose a new estimator for step 2 which is based on the analysis of Cheung and Chappell (2002). These authors provided an interpretation and insight into poor model specification by breaking down the parameter space: $[A, B] = \cup_{d \in D} H_d$, with $H_d = \{a \in [A, B] : |f(d|a) - \alpha| < |f(d'|a) - \alpha|, \forall d' \neq d\}$. The set H_d is the parametric space on which the model recommends dose d as the MTD. Asymptotic concerns together with the sequential nature of CRM and the partition of the parameter space leads to the following assumption. Let a_d be such that: $f(d|a_d) = \beta_d$.

Assumption 1.3. $a_{d^*} \in H_{d^*}$ and $\forall d \in D \setminus \{d^*\}, a_d \notin H_d$.

Assuming that the functions $f(\cdot|a)$ are increasing for all a , Assumption 1.3 is equivalent to the one under which Azriel (2012) shows the strong consistency of CRM. Theorem 1.1 rules out the existence of a method providing almost sure convergence to the MTD regardless of the circumstances. Indeed, Assumption 1.3 can not be checked because it requires control over the reality expressing itself via the parametric elements a_d (Cheung and Chappell, 2002, Figure 1) (Section 2.2, Figure 1). However, it does throw light on how the method works: the goal is to identify the MTD among a small range of doses D , at the same time the CRM leans on the estimation of a parameter in an infinite set $[A, B]$. The method tries to ascertain the belonging of this parameter to one of the sets of the family $(H_d)_{d \in D}$. On the basis of this analysis, we propose the following Bayesian estimator for the next dose (step 2) of CRM:

$$X_{n+1} = \arg \max_{d \in D} G_n(H_d). \quad (2)$$

The parametrization of the CRM can be expressed in terms of the MTD, θ . The prior Π is a distribution on the range of doses such that for all $\theta \in D$, we have: $\Pi(\theta) = G(H_\theta)$. The

family of priors $\Lambda = (\Lambda_\theta)_{\theta \in D}$ describe the dose-toxicity curve and we set $\Lambda_\theta(\cdot) = G(\cdot | H_\theta)$. This means that Λ_θ is the distribution G on the set H_θ . The prior G is then equal to the probability measure $\Lambda \otimes \Pi$ and the posterior Π_n , following the observations (X_1^n, Y_1^n) can be easily obtained by integrating the parameter of the dose-response curve for each $\theta : \Pi_n(\theta) \propto [\int L_n(a) \Lambda(da | \theta)] \Pi(\theta)$.

This hierarchical model being strictly equivalent to the Bayesian CRM, has no particular value in this form. It does though allow a greater conceptual understanding of the distributions $\Lambda(\cdot | \theta)$ and their topological supports. Relaxing the structure of these distributions and adding some flexibility amounts to a model involving a nuisance parameter. This semi-parametric setting provides methods that benefit from improved asymptotic properties while still conserving operating characteristics for small sample size that are similar and in some cases better to those of the CRM.

2 Semi-Parametric Models

2.1 General semi-parametric structure

Semi-parametric models (SPM) take as their starting point the direct modelling of the MTD itself. This is formalized within the framework of Bayesian hierarchical models. It can also be viewed in terms of model selection based on bayes factors. The hierarchical posterior allows us to compare and evaluate m classes indexed by the main parameter of interest, the MTD. These classes are structured by a prior referred to as a prior model (see Section 2.2). The initial topological support of the prior is a broad one, corresponding to a non-informative situation.

We model the accumulating information via an m -tuple of Bernoulli laws. To this end we introduce F an m -dimensional vector space of Bernoulli parameters covering a very wide range of situations. Let $q = (q_1, \dots, q_m) \in F$, and q_j the specific parameter corresponding to dose j . The set F is partitioned in terms of the main parameter of interest, θ , which provides us m distinct classes, each individual class containing an infinite set of members sharing the same MTD: $F = \bigcup_{\theta \in D} F_\theta$, with

$$F_\theta = \{q \in F : \forall j \in D, |q_\theta - \alpha| \leq |q_j - \alpha|, j < \theta \Rightarrow q_j < \alpha, j > \theta \Rightarrow q_j > \alpha\}. \quad (3)$$

Given that θ takes on the value of some particular dose-level, then F_θ is the collection of dose-toxicity curves having θ as the MTD. Lower levels in F_θ will necessarily have probabilities of toxicity less than α , and conversely for higher levels. For all β in F_θ , the MTD is θ . The vector $(f(d, a))_{d \in D}$, with $a \in H_\theta$ (see the preceding section), is then included in F_θ . Indeed, F_θ can be seen as a general extension of H_θ , such that F contains all the probability measure in compliance with Assumption 1.1. From our point of view, determining the MTD θ can be summarized by the following question. Which class is the most plausible one to have generated

the data? The likelihood is:

$$L_n(q) = \prod_{1 \leq i \leq n} (q_{X_i})^{Y_i} (1 - q_{X_i})^{1 - Y_i} = \prod_{1 \leq j \leq m} q_j^{n_j^1} (1 - q_j)^{n_j^0},$$

where n_j^1 is the number of toxicities at dose j and $n_j = n_j^1 + n_j^0$ is the number of patients treated at dose j . The set F is endowed with a probability measure $\Lambda \otimes \Pi$ such that Π is a measure on D and the support of the measure $\Lambda(\cdot|\theta) = \Lambda_\theta(\cdot)$ is included in the class F_θ . This means that Π is a vector of m non-negative numbers summing up to 1 and Λ is a family of distribution $(\Lambda_\theta)_{\theta \in D}$ indexed by the potential MTD, θ . The posterior distribution of θ given (X_1^n, Y_1^n) is:

$$\Pi_n(\theta) = \Pi(\theta|X_1^n, Y_1^n) \propto \left[\int \prod_{j=1}^m q_j^{n_j^1} (1 - q_j)^{n_j^0} \Lambda_\theta(dq) \right] \Pi(\theta). \quad (4)$$

When we focus on the class F_θ , the posterior distribution of Λ_θ given (X_1^n, Y_1^n) is: $\Lambda_{\theta,n}(dq) \propto \prod_{1 \leq j \leq m} q_j^{n_j^1} (1 - q_j)^{n_j^0} \Lambda_\theta(dq)$. By replacing this result in (4), we have:

$$\Pi_n(\theta) \propto \left[\int q_{X_n}^{Y_n} (1 - q_{X_n})^{1 - Y_n} \Lambda_{\theta,n-1}(dq) \right] \Pi_{n-1}(\theta). \quad (5)$$

Thus, each new observation leads first to an update concerning the distribution Π by weighting according to the expected value of the likelihood with respect to q conditioned by θ . In a second step, this observation is used to update the probability measures Λ_θ on classes using Bayes formula. In the following section, the family $(\Lambda_\theta)_{\theta \in D}$ will be called the prior model because of the predictive model-like role it plays in sequential decision making. Fitting the model is carried out by updating the prior Λ_θ . Finally, estimators of the MTD, $\hat{\theta}$ and of the toxicities at dose j , $\tilde{\beta}_j$, arise naturally and are the same we have already proposed and described for the CRM:

$$\hat{\theta}_n = \arg \max_{\theta \in D} \Pi_n(\theta), \quad \tilde{\beta}_j = \mathbb{E}_{(\Lambda \otimes \Pi)_n} [q_j] = \sum_{\theta=1}^m \left[\int q_j \Lambda_{\theta,n}(dq_j) \right] \Pi_n(\theta). \quad (6)$$

The general method might be summarized by the following points:

1. The current sample of observations (X_1^n, Y_1^n) is used to update the posterior Π_n .
2. The estimator of the next dose is $X_{n+1} = \hat{\theta}_n$.

We introduce this semi-parametric class of models in the most usual situation of a Phase study I summarized by Assumption 1.1. Under this assumption, a phase I study only deals with a finite dimensional parameter. The term 'semi-parametric' refers to the methodology used to build this structure. With the goal of determining the parameter of interest, the MTD, a parametric model is extended by using nuisance parameters: when the MTD is dose 3,

we do not have to estimate precisely q_1 or q_5 . The proposed model covers the whole range of scenarios, "which consists of all probability measures on the sample space for the observations" (Bickel et al., 2005), as a non-parametric model does. The gain of flexibility does not result in a loss of simplicity. As all the parameters are readily available, the semi-parametric class of models can be easily calibrated to reproduce the behavior of almost all of the designs currently in use, including the algorithmic designs. This broad generalization allows us to investigate theoretical questions, to more readily allow comparison between competing designs and to look for ways to improve on any given design. Finally, the simplicity of calibration (see below) together with the great range of possibilities suggest that the semi-parametric class of models might be extended to more complicated situations in phase I involving infinite dimensional parameters: known examples being continuous grade of toxicities, heterogeneous populations, combination studies and time-to-event toxicities.

2.2 A simple prior model specification

A general calibration of SPM can be obtained as follows. We focus on the prior distributions inside the classes, $(\Lambda_\theta)_{\theta \in D}$. The support S_θ of the distribution Λ_θ will be included in F_θ and will reflect locally the ordering of the parameters β_j (Assumption 1.1). This local property is sufficient to ensure that the design behaves in a sensible way. Indeed, after each new observation or set of observations, the practitioner would like the method to indicate if the dose appears too high, too low or acceptable. For this purpose, a natural partition of the interval $[0, 1]$ into 3 sets is introduced: $I_\epsilon = [\alpha - \epsilon, \alpha + \epsilon]$, $A = [\alpha + \epsilon, 1]$ and $B = [0, \alpha - \epsilon]$. The support of distributions that we choose are in m dimension as the whole space of probability measures for the observations. They are defined according to the constraints on the MTD. The Bernoulli parameters at each dose are considered as independent from the point of view of a single class θ .

Assumption 2.1. (i) The support S_θ^j of the marginal Λ_θ^j satisfies: $S_\theta^j = B$ when $j < \theta$, $S_\theta^j = I_\epsilon$ when $j = \theta$ and $S_\theta^j = A$ when $j > \theta$. We then have: $S_\theta = S_\theta^1 \times \dots \times S_\theta^m \subset F_\theta$. (ii) Λ_θ is a product of unidimensional distributions at each doses: $\Lambda_\theta(dq) = \Lambda_\theta^1(dq_1) \times \dots \times \Lambda_\theta^m(dq_m)$.

The point (i) of the above assumption means that the support of q_j depends on whether j is above, below, or at the level of the MTD. In each circumstance the support reflects the constraints on q_j . The point (ii) is an independence assumption given θ . This independence assumption could be relaxed although we have not studied this. Choosing the width of the central interval is an important step in the parametrisation as it determines consistency (see section 3). We propose here to use an interval centered in α . This leads the method to mimic the CRM under the most common scenarios. Non-centered intervals can also be used which provide more or less conservative results.

The direct calibration of the supports S_θ allows us to avoid those cases described in (Shen and O'Quigley, 1996) that can result in non-convergence (see Figure 1 and scenario 8 in

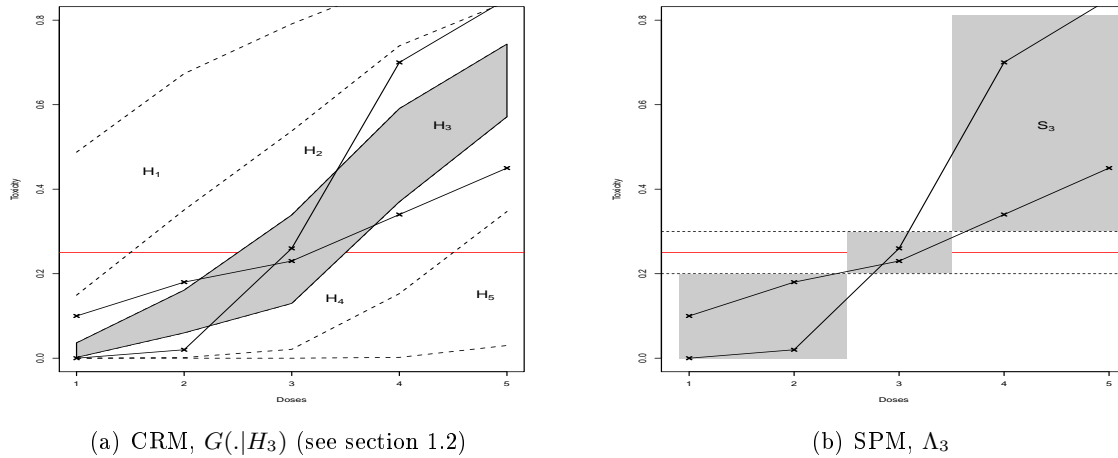


Figure 1: Support of the dose-response curves for distributions $G(\cdot|H_3)$ and Λ_3 . The broken lines represent two kinds of poor specification for CRM, the flatter curve fails Assumption 1.3.

Table 3) as well as poor performance consequent upon poor model specification (a slope in the neighborhood of the MTD, Figure 1 and scenario 5 and 6 in Table 3). Note that the independance of parameters q_j 's in Assumption 2.1 is only valid conditionally given θ . This means that all the information about dependence is captured by θ and is used to accomplish the main goal of the study: determining the MTD, θ . Thus, the independance in the class implies that the supports S_θ include locally decreasing scenarios of the dose-response phenomom, but all the scenarios in that class respect the monotonicity between θ and the other doses. Figure 2 (b) illustrates the dependence between the parameters q_j 's and how we can describe sub-models that include monotonicity restrictions. In this way, the family $(\Lambda_\theta)_{\theta \in D}$ can be built from the model of CRM. As it is used to make inference in place of the usual parametric model, the family $(\Lambda_\theta)_{\theta \in D}$ is called the prior model.




Figure 2: Prior model built on the CRM model

The Beta distribution linked to the likelihood has a key role as it allows us to use conjugate

priors. $\mathcal{B}(a + 1, b + 1)$ and $\mathcal{B}_I(a + 1, b + 1)$ denote the Beta distribution and its restriction to interval I with shape parameters $a + 1$ and $b + 1$. Let g be the following function: $g : [0, 1] \times \mathbb{R}_+^2 \rightarrow [0, 1]$, with $g(x, a, b) = x^a(1 - x)^b$. The density function of $\mathcal{B}(a + 1, b + 1)$ is $g(\cdot, a, b)/B(a + 1, b + 1)$, with $B(a + 1, b + 1) = \int_0^1 g(x, a, b)dx$. When conjugate priors are used, the whole prior model can be summarized by a triplet $[\epsilon, (q^\theta)_{\theta \in D}, c]$. The number ϵ belongs to $[0, 1]$ and determines the centered interval. The vectors $q^\theta \in [0, 1]^m$ are the modes of distributions Λ_θ . The positive real value c is the dispersion parameter of the distributions. Uniform priors on the topological supports are updated so that,

$$\Lambda_\theta \approx \prod_{j=1}^{\theta-1} \mathcal{B}_B(cq_j^\theta + 1, c(1 - q_j^\theta) + 1) \times \mathcal{B}_{I_\epsilon}(cq_\theta^\theta + 1, c(1 - q_\theta^\theta) + 1) \times \prod_{j=\theta+1}^m \mathcal{B}_A(cq_j^\theta + 1, c(1 - q_j^\theta) + 1) \quad (7)$$

In Figure 2 (b), the prior model is built by using a model f of CRM: $q^\theta = [f(j, |\alpha_\theta)]_{j \in D}$, with $f(d, \alpha_\theta) = \alpha$ (see Figure 2 (a)). All the parametrisations proposed in this article fulfill an assumption about stochastic ordering on the prior model. At each moment of the trial, for any dose j_0 , the marginal posterior $\Lambda_{\theta, n}^{j_0}(\cdot)$ should be stochastically greater than $\Lambda_{\theta', n}^{j_0}(\cdot)$ when θ is smaller than θ' (see supplementary material, Section 2). In the setting of SPM, it is the main argument needed to obtain the coherence property of Cheung (2005). The prior model should respect a local order between the classes corresponding to the different MTD, and not inside one class.

Theorem 2.1. *If the prior model satisfies 2.1 and the assumption about stochastic ordering (supplementary material), then SPM is coherent (Definition 1.1).*

This analysis throws a light on the natural similarities between the prior model and parametric models which satisfy the monotonicity assumption.

2.3 Summary and illustration

We have $q = (q_1, \dots, q_m)$ a possible dose-toxicity scenario, i.e., q_j is the probability of toxicity at dose j . $\Lambda = (\Lambda_\theta)_{\theta \in D}$ is a family of priors for q , indexed by θ , the parameter identifying the MTD. Given the MTD, θ , Λ_θ is made up of only those vectors q for which we know the MTD to be at dose θ . Π is the prior on the parameter θ , often a discrete uniform. The couple (Π, Λ) then describes a hierarchical model where the first level deals with the goal of the study, the MTD itself, and the second level concerns the dose-response curve. Bayes formula is used for obtaining the posterior Π_n according to the data already observed (see Section 2.1, Equation 4). The next selected dose is chosen to be the most probable: $X_{n+1} = \arg \max_{\theta \in D} \Pi_n(\theta)$. The prior model $(\Lambda_\theta)_{\theta \in D}$ can be calibrated in such a way to avoid those difficulties that arise for the parametric model as a result of misspecification. Given θ , we assume independence between the different probabilities of toxicity, which means that each prior Λ_θ is the simple product of his marginals Λ_θ^j on each dose j (Section 2.2). Under this assumption, the next dose selected

corresponds to the parameter θ that maximizes the product of the expected likelihood at each doses weighted by the prior $\Pi(\theta)$:

$$X_{n+1} = \arg \max_{\theta \in D} \Pi_n(\theta) = \arg \max_{\theta \in D} \left[\prod_{j=1}^m \int L_n(q_j) \Lambda_{\theta}^j(dq_j) \right] \Pi(\theta) \quad (8)$$

Our approach is based on the following two steps:

- **Calibration of Λ :** (a) A partition of the probability space of toxicities is used for calibrating the marginal laws: $B = [0, \alpha - \epsilon]$, $I_{\epsilon} = [\alpha - \epsilon, \alpha + \epsilon]$ and $A = [\alpha + \epsilon, 1]$. (b) The marginals are chosen from the Beta family (conjugacy). (c) The support S_{θ}^j of the marginal Λ_{θ}^j is restricted to one of the three preceding intervals according to the following rules: if $j < \theta$, the marginal has support limited to B , if $j = \theta$, the marginal has support limited to I_{ϵ} and if $j > \theta$, the marginal has support limited to A .
- **Calibration of Π :** This prior is a vector of probabilities on the range of doses. It can be easily calibrated. When the clinician has no extra-information to provide on the location of the MTD, a non-informative prior might be chosen, a discrete uniform being a natural candidate. However, given that this prior will drive the early escalation behavior until we encounter a DLT, it is appealing to note that, not only will this prior impact early behavior but, we can calibrate the prior in such a way as to obtain the very behavior we would like to see. Early escalation can be slowed down or speeded up by simple calibration of this prior.

The following examples are based on a very simple parametrisation. The target is fixed at $\alpha = 0.2$.

SPM (0.05,0,0,0): *The value ϵ is 0.05 such that $I_{\epsilon} = [0.15, 0.25]$. All the marginals are uniform on their respective interval. The distribution Π is uniform on the range of doses.*

SPM(0, 1/10, 1/3, 40) : *The size of the centered interval is null: $\Lambda_{\theta}^{\theta}$ is a constant random variable in α , for all $\theta \in D$. We set: $j < \theta \Rightarrow q_{\theta}^j = 1/10$, $j > \theta \Rightarrow q_{\theta}^j = 1/3$ and c is equal to 40 (See Equation (7)). The distribution Π is uniform.*

In Table 4, the overall performances are summarized. When escalating, skipping a dose is not allowed as this is now a requirement in these kinds of designs. These designs are computationally very fast. R codes are available from the authors on request or at address: https://github.com/MatthieuMC/SPM_project_01.

3 Large sample theory

The interval I_{ϵ} is centered on α : $I_{\epsilon} = [\alpha - \epsilon; \alpha + \epsilon]$. All of the results presented here remain valid for a non symmetric interval. Theorem 1.1 states that if treatment in a sequential experiment is determined by the current estimator of the MTD, then this estimator cannot be strongly

consistent. However, setting $\epsilon \geq 0$, if we assume that one or more doses are close enough to α with accuracy ϵ , then we shall obtain almost sure convergence of the design to this set of doses. Conversely, if we assume that neither dose is close enough to the threshold, the current estimator shall recommend alternatively the two doses with toxicities directly located on both sides of α . We introduce the following technical assumption which leans on the regularity of the prior model.

Assumption 3.1. *Let S_θ and S_θ^j be the topological support of Λ_θ and Λ_θ^j . The following conditions are valid except when Λ_θ^θ is a Dirac measure.*

(a) *For all $j \in D$, the marginal distribution Λ_θ^j is absolutely continuous with respect to Lebesgue measure and λ_θ^j denotes its density function.*

(b) *There exist two numbers s and S in $(0, \infty)$, such that, for all $(j, \theta) \in D^2$, we have:*

$$\forall q_j \in S_\theta^j, s < \lambda_\theta^j(q_j) < S.$$

When the density function can not be bound into the neighborhoods of 0 or 1, we can obtain compliance with the preceding assumption by using uniform priors on the small intervals $[0, \delta[$ and $]1 - \delta, 1]$:

$$\lambda_\theta^j(q_j) \propto \int_0^\delta g(q, cq_j^\theta, c(1 - q_j^\theta)) dq \mathbf{1}_{[0, \delta[}(q_j) + g(q_j, cq_j^\theta, c(1 - q_j^\theta)) \mathbf{1}_{]1 - \delta, 1]}(q_j).$$

Moreover, from the point of view of proving consistency, this assumption deals only with theoretical scenarios where there exist some doses which are never toxic and others which are always toxic. The following two definitions characterize the asymptotic behavior of SPM: ϵ -sensitivity is a property connected to indifference intervals. (Cheung and Chappell, 2002).

Definition 3.1. *Let $\epsilon \geq 0$ et $I_\epsilon = [\alpha - \epsilon; \alpha + \epsilon]$. We consider the collection of doses associated with a toxicity belonging to I_ϵ : $\mathcal{E}(I_\epsilon, \beta) = \{j \in D : \beta_j \in I_\epsilon\}$. A method, \mathcal{M} , is called ϵ -sensitive, if for all β such that $\mathcal{E}(I_\epsilon, \beta) \neq \emptyset$, we have:*

$$\mathbb{P}_\beta [\exists N, \forall n > N : \mathcal{M}(X_1^n, Y_1^n) \in \mathcal{E}(I_\epsilon, \beta)] = 1.$$

If the true situation is such that a unique dose is associated with a target in the interval I_ϵ , then a method that is ϵ -sensitive converges almost surely to the MTD. When no dose has a target located within I_ϵ , the SPM will assume an oscillating behavior between two doses with toxicities either side of the target α .

Definition 3.2. *The letter \tilde{D} denotes the set of doses infinitely observed:*

$$\tilde{D} = \{j \in D : n_j \xrightarrow{n \rightarrow +\infty} +\infty\}.$$

Let b (below) and a (above) be the two consecutive doses associated to toxicities either side of

the target α . A method, \mathcal{M} , is called ϵ -balanced, if for all β such that $\mathcal{E}(I_\epsilon, \beta) = \emptyset$, we have: $\tilde{D} = \{a, b\}$, a.s.

We might view oscillation as a desirable property for designs whose aim is to locate some dose, since, if it is not possible to obtain a method that converges almost surely in all circumstances (Theorem 1.1), it is nonetheless natural to want to construct an estimator, on the basis of observations, that is strongly consistent. As soon as a dose belongs to \tilde{D} , it becomes possible to reliably estimate its associated toxicity and the MTD belongs to the set $\{a, b\}$, which is the minimal set on which we need to have observation when the goal is that of determining almost surely the MTD.

Theorem 3.1. *Under the Assumptions 2.1 and 3.1, SPM is ϵ -sensitive and ϵ -balanced (see remark 1).*

Proof. The proof is given in the supplementary material. □

Remark 1. *In this theorem, and its proof, we consider that there exists no dose j_0 such that β_{j_0} equals $\alpha \pm \epsilon$, with $\epsilon > 0$. This assumption is made for the purpose of clarity in presenting the results.*

Large sample behavior of SPM is established by Theorem 3.1. In the case where $\mathcal{E}(I_\epsilon, \beta)$ is non empty, the sequence of doses selected by SPM converges almost surely to one or more elements belonging to $\mathcal{E}(I_\epsilon, \beta)$. In the case where $\mathcal{E}(I_\epsilon, \beta)$ is empty, the running estimate of SPM oscillates between those doses either side of the indifference interval. The two asymptotic properties of SPM are simultaneously complementary and antagonistic, since, whenever we diminish the size of the interval I_ϵ , we increase the set of circumstances where the method is ϵ -balanced and we diminish the ones where it is ϵ -sensitive. Furthermore, the interval I_ϵ can be chosen as small as we wish without having an effect on overall performance of the method. In the case of the simple parametrisation SPM(0, 1/10, 1/3, 40) presented in section 2.3, ϵ is zero, which aside from the case where a toxicity would equal α leads necessarily to an oscillation. This oscillation, giving an approximation of the two toxicities, allows for us to construct convergent estimators in all of the scenarios β :

$$\tilde{\theta}_n = \arg \min_{j \in \hat{D}_n} |\alpha - \hat{\beta}_{j,n}|, \tag{9}$$

where \hat{D}_n is the set of the last two selected doses: $\hat{D}_n = \{\hat{\theta}_n, \hat{\theta}_{n'}\}$, with $n' = \max\{j < n : \hat{\theta}_j \neq \hat{\theta}_n\}$.

Corollary 3.1. *Under the Assumptions 2.1 and 3.1, with $\epsilon = 0$, the estimator $\tilde{\theta}_n$ converges almost surely to the MTD.*

Proof. In the case where $\epsilon = 0$, the SPM is ϵ -balanced, which amounts to saying that: $\hat{D}_n \rightarrow \tilde{D} = \{a, b\}$. The law of large numbers (Lemma 1.1) leads to an immediate result. □

Strong consistency of the estimator based on isotonic regression of the observations could be obtained in the same way. According to the impossibility theorem of Azriel et al. (2011), the estimators which possess the property of strong consistency can not be the current estimator given by the method. The consistency of these estimators, regardless of the scenario, is possible because the adjacent doses a and b will be chosen infinitely often by the running estimator: $n_a \rightarrow \infty$ and $n_b \rightarrow \infty$. In the case of a finite sample size, experimentation will be spread over two doses and may give the impression of convergence if the observations on one dose are close enough to the expected rate. The following corollary considers an asymptotic characterization of the number of observations allocated to the dose a relative to the number allocated to b . For this it is helpful to recall entropy and divergence. Two Bernoulli laws P and Q , are denoted by their parameters p and q . The entropy of Q relative to P is: $H(q|p) = -p \log(q) - (1-p) \log(1-q)$, with $\log 0 = -\infty$ et $0 \times (-\infty) = 0$; we denote the entropy of P : $H(p) = H(p|p)$. The divergence of Kullbac-Leibler of P relative to Q is:

$$D_{KL}(p||q) = H(q|p) - H(p) = p \log\left(\frac{p}{q}\right) + (1-p) \log\left(\frac{1-p}{1-q}\right). \quad (10)$$

For $p \in [0, 1]$, the function $D_{KL}(p||\cdot)$ is strictly decreasing on $[0, p]$ and strictly increasing on $[p, 1]$ and its minimum in p is equal to 0.

Corollary 3.2. *Under the Assumptions 2.1 and 3.1, when $\epsilon = 0$ and when at least one of the toxicities β_a and β_b is different from α , we have:*

$$\frac{n_a}{n_b} \xrightarrow{n \rightarrow +\infty} \frac{D_{KL}(\beta_b||\alpha)}{D_{KL}(\beta_a||\alpha)}, \text{ a.s.}$$

Proof. The proof is given in the supplementary material. □

In situations more general than those for well specified models, the Kullback-Leibler divergence is often used as an appropriate distance measure between two probability laws. This pseudo-distance is the natural tool to use when showing consistency for Bayesian or maximum likelihood estimators. In this way, the running estimate for SPM oscillates between doses a and b according to an asymptotic ratio that is inversely proportional to the pseudo-distance of Kullback-Leibler between β_a and β_b at the chosen target: the greater the pseudo-distance between β_a and α relative to that between β_b and α the more SPM will recommend the dose b (and vice versa). The purpose of the following section is to highlight the practical performance of the usual SPM through a comparison with the CRM.

4 Simulations

The CRM demonstrates good performance with respect to the following criteria: (PCS), the percentage of correct selection at the final recommendation and (TR), the percentage of patients treated at the MTD. Here, we show that, if we so wish, the prior model of SPM can be

calibrated in such a way as to reproduce this same performance across many scenarios. This particular parametrisation is called SP-CRM. The simulations are carried out under a common situation ($\alpha = 0.2$, power model) where the CRM is considered as close to optimal under an adaptative minimum-variance criterion (Tian, 2016, Theorem 1 and its interpretation).

Calibration

The target rate is fixed at 0.20. The goal is to locate the MTD as one of 6 available doses. There are 25 patients in each study. We make use of two stage CRM (O’Quigley and Shen, 1996) based on some lead-in rule until we observe the first toxicity and then we use maximum likelihood. As proposed by Cheung (dfcrm documentation in the CRAN package), the chosen skeleton is $u = (0.05, 0.10, 0.20, 0.35, 0.50, 0.70)$ with Normal law for the prior $\mathcal{N}(0, 1.34^2)$ together with the power model. Each cohort is of size one meaning that we estimate the dose after each patient. We include the classical restriction that no skipping is allowed. For SP-CRM, the prior model verifies 2.1 and Equation (7). It is summarized by $(\epsilon = 0.015, (q^\theta)_{\theta \in D}, c = 48)$. The modes $(q^\theta)_{\theta \in D}$ are chosen close to the model of the CRM, as in Figure 2 a). They are given by Table 1. The law Π can be used as an alternative

Table 1: The modes of the prior model

q_j^θ	q_1	q_2	q_3	q_4	q_5	q_6
q_1	0.20	0.12	0.02	0.01	0.00	0.00
q_2	0.29	0.20	0.07	0.05	0.00	0.00
q_3	0.42	0.36	0.20	0.08	0.02	0.00
q_4	0.57	0.48	0.35	0.20	0.09	0.01
q_5	0.69	0.62	0.50	0.34	0.20	0.04
q_6	0.82	0.78	0.70	0.58	0.44	0.20

way to reproduce any initial dose escalation and we are able to choose it so that the method follows naturally a given increasing sequence of doses until we observe the first toxicity. An increasing sequence is denoted $s = (s_1, \dots, s_k)$ with $s_k \in D$. The modes and the dispersion being fixed, we define $\mathcal{B}(s)$ the set of distributions Π that produce the sequence s until the first observed toxicity: $\mathcal{B}(s) = \{\Pi : Y_j = 0, 1 \leq j \leq k \Rightarrow X_j = s_k, 1 \leq j \leq k\}$. The law \mathbb{P}_s is the one minimizing the distance in the sense \mathcal{L}^2 between the uniform distribution \mathcal{U} and the closure of the convex set $\mathcal{B}(s)$: $\mathbb{P}_s = \arg \min \|\Pi - \mathcal{U}\|_2$, for $\Pi \in \overline{\mathcal{B}(s)}$. This law does not belong to $\mathcal{B}(s)$, but there exist distributions in $\mathcal{B}(s)$ arbitrarily close to \mathbb{P}_s . It is then possible to find an approximation as accurate as we wish of the least informative distribution belonging to the closure of the set of measure providing the sequence (R codes available on request). Table 2 provides such distributions accurate to 10^{-3} for different sequences. These laws are not normalized but this does not impact the posterior Π_n . In the following simulations, the distribution corresponding to the sequence d is used. The CRM also produces this sequence while awaiting to observe the first toxicity.

Table 2: Prior Π , walk through the levels awaiting the first observed toxicity.

Sequences without toxicity	$\Pi(1)$	$\Pi(2)$	$\Pi(3)$	$\Pi(4)$	$\Pi(5)$	$\Pi(6)$
a: 111222333444555666	1	0.832	0.482	0.346	0.194	0.103
b: 112233445566*****	1	0.913	0.663	0.554	0.392	0.272
c: 123456*****	1	0.999	0.910	0.883	0.787	0.709
d: 123456*****	1	0.999	0.910	0.883	0.787	0.604




Figure 3: An example of a sequence for SP-CRM, $\beta = (0.01, 0.07, 0.10, 0.20, 0.40, 0.70)$; Δ : toxicity, \times : non-toxicity.

Model and prior model

Table 3 shows performance of SP-CRM when compared to the CRM according to the criteria (PCS) and (TR) for 10 000 replications. The findings show very similar behavior for the first 4 scenarios. When the data are generated exactly by the model being used for CRM (scenario 3), rather surprisingly, that does not appear to grant any advantage to the method and the SP-CRM appears to suffer no handicap as a price to pay for the extra-flexibility and adaptability of its prior model. On the other hand, scenario 4 presents an interesting illustration in which the CRM fails to satisfy the Assumption 1.3 and, as a result, does not possess the property of convergence to the MTD. Despite this, for a trial of 25 patients, it is difficult to observe any theoretical advantage of SP-CRM over CRM. However, it is enough to slowly increase sample size to observe convergence difficulty manifesting itself in practice (see figure 4). Increasing the number of patients included in the study fails to lead to improvement for CRM. In contrast, the SP-CRM is ϵ -sensitive and the portion of the curve that is traced out corresponds to almost sure convergence. In a real practical sense, as sample size increases, SP-CRM does better and better. In some ways, for CRM we were fortunate in that the best performance was already obtained around 25 subjects and increasing this number was not rewarded by increased accuracy. Beyond that sample size the handicap begins to show itself. In other scenarios, for instance 5 and 6, where the model specification is yet more severely

Table 3: Some varied scenarios.

Doses		1	2	3	4	5	6
Scenario 1		0.20	0.26	0.28	0.3	0.35	0.50
PCS	SP-CRM	49.4	21.5	13.2	9.6	5.4	0.6
	CRM	48.1	19.5	14.3	11.2	6.0	0.6
TR	SP-CRM	47.4	20.6	13.5	9.1	7.0	2.1
	CRM	47.6	17.6	14.1	10.7	7.6	2.2
Scenario 2		0.05	0.10	0.20	0.35	0.50	0.70
PCS	SP-CRM	2.3	22.7	54.0	19.7	01.2	0.0
	CRM	02.4	22.2	53.9	20.2	01.3	0.0
TR	SP-CRM	10.8	24.3	39.0	19.0	05.9	00.7
	CRM	12.3	22.1	37.7	20.4	06.4	00.8
Scenario 3		0.01	0.02	0.05	0.09	0.18	0.40
PCS	SP-CRM	0.0	0.2	2.8	20.3	59.2	17.3
	CRM	0.0	0.1	3.4	21.8	58.4	16.1
TR	SP-CRM	4.6	6.0	10.5	19.9	40.7	17.9
	CRM	4.9	5.3	9.7	20.7	40.1	19.0
Scenario 4		0.01	0.02	0.05	0.11	0.14	0.21
PCS	SP-CRM	0.0	0.1	3.2	15.7	31.0	49.8
	CRM	0.0	0.1	3.4	15.5	31.2	49.6
TR	SP-CRM	4.6	5.8	10.8	16.7	26.7	35.1
	CRM	4.9	5.3	10.2	16.7	25.9	36.0
Scenario 5		0.0	0.0	0.16	0.3	0.35	0.4
PCS	SP-CRM	0.0	2.3	51.7	31.5	11.1	3.2
	CRM	0.0	3.5	46.7	33.6	12.6	3.6
TR	SP-CRM	4.0	11.8	40.3	24.3	13.7	5.8
	CRM	4.7	11.0	36.3	26.7	14.5	6.5
Scenario 6		0.0	0.0	0.0	0.23	0.3	0.35
PCS	SP-CRM	0.0	0.0	10.2	56.8	23.6	9.2
	CRM	0.0	0.0	10.5	52.3	26.9	10.2
TR	SP-CRM	4.0	4.0	19	38.8	22.8	11.1
	CRM	4.0	4.0	16.9	37.8	24.4	12.7

tested and struggles to accommodate a slope in the neighbourhood of the MTD that is a strain to fit, SP-CRM shows clearly superior performance, both as measured by PCS and as measured by TR. The SP-CRM gains its advantage from the flexibility of the prior model that can readjust to each observation. The same argument underlies its asymptotic performance and its adaptability to those situations that appear far removed from the model. In order to confirm this impression, we randomly generated scenarios by making use of order statistics of “quasi-uniform” variates (see figure 5). For this purpose the following algorithm, called *the pseudo-uniform scenario* is used.

- The MTD is selected uniformly from a range of doses D : $MTD \sim \mathcal{U}_D$; resulting in the value k .






Figure 4: For scenario 6, (PCS) as a function of the number of included patients in the study. \square : SP-CRM ; \times : CRM.

- We randomly select an upper bound $B_s = \alpha + (1 - \alpha) \times M$; M is a random variable having a Beta law depending on the MTD and the number of doses m : $M \sim B(\max\{m - k; 0.5\}, 1)$.
- The random scenario β has the law of an ordered sample of m uniform laws on $[0, B_s]$ conditioned by the event $\{MTD = k\}$.



(a) Sample under the pseudo-uniform algorithm



(b) Sample under Λ

Figure 5: Scenarios tested and scenarios generated by the prior.

The second point downweights the importance of the more extreme scenarios in which the toxic probabilities following the MTD rise very sharply. Such scenarios can still be sampled but less frequently. Sampling of the law M is natural; indeed when we have 6 doses and the MTD is located at level 2, B_s is the maximum of 4 uniform laws on $[\alpha, 1]$. Table 4 compares the performance of the CRM and its semi-parametric version over the set of 100 000 randomly generated scenarios. Three additional criteria enabling comparison are introduced. $(TR(a, b))$

is the percentage of patients treated at doses a and b (see the definition 3.2); (Δ) is the mean of the difference between the toxic rates at which patients are treated and the toxicity at the MTD. The fifth criteria (R- Δ) is an index based on the statistic (Δ) relative to that value obtained by the optimal design (O’Quigley et al., 2002). The optimal design is based on the idea of complete and incomplete information. We can use it sequentially, in a theoretical setting, to provide a running best estimate of the MTD, the level at which we would like to treat the next included patient in the study. In order to maintain comparability, at least early on, we constrain the optimal design to similar behaviour such as that imposed on the CRM, i.e., only increases in level by one level at a time. This helps provide a reference for the criterion (Δ) : $\Delta(\text{Opt}) = 9.75$. The base reference is calculated for the CRM.

$$\text{R-}\Delta(\mathcal{M}) = \frac{\Delta(\mathcal{M}) - \Delta(\text{OPT})}{\Delta(\text{CRM}) - \Delta(\text{OPT})} \quad (11)$$

As R- Δ gets closer to 0, all the more the considered method gets close to the optimal design. In all categories, the SP-CRM obtains the best results (Table 4). Regarding the criteria (Δ) ,

Table 4: Comparison from a sample scenario (size=100 000).

Criterion	PCS	TR	TR(a,b)	Δ	R- Δ
CRM	50.43	39.23	59.68	10.05	1.0
SP-CRM	51.45	39.56	60.22	9.93	0.6
SPM(0, 1/10, 1/3, 40)	51.16	39.19	59.80	10.12	1.23

the difference between the SP-CRM and the optimal method is 40% smaller than that which obtains when comparing the CRM to the optimal. This significant gain can be explained in part by the fact that we are very close to the performance of the optimal method. The parametrisation SPM(0, 1/10, 1/3, 40) from Section 2.2 obtains good results, even if its distance to the optimal method looks slightly greater than that for the CRM. This leads us to conclude that very simple parametrisations of SP-CRM can attain comparable performance to those of the CRM, even in cases where the CRM is considered to be near optimal (Tian, 2016).

In these simulations, it is important to keep in mind that our goal was to emulate as best we could the behaviour of the CRM. Since this can be accomplished we can conclude that we do no worse than the CRM. However, the greater flexibility allows us to do better in those particular cases that prove thorny for the CRM since the explanation for the awkward behaviour here is the strong parametrisation of the CRM, a feature that is greatly relaxed in SP-CRM. It remain to study the great range of SPM parametrisation under different circumstances. The posterior Π_n on the doses suggests some avenues of exploration for estimating the MTD by groups of doses which may show itself to be of value when we move beyond Phase I to the Phase I/II setting. In Figure 6, the target is $\alpha = 0.2$ and the MTD is located between 3 and 4. At the end of the trial the posterior puts 81% of its mass over these two doses. In

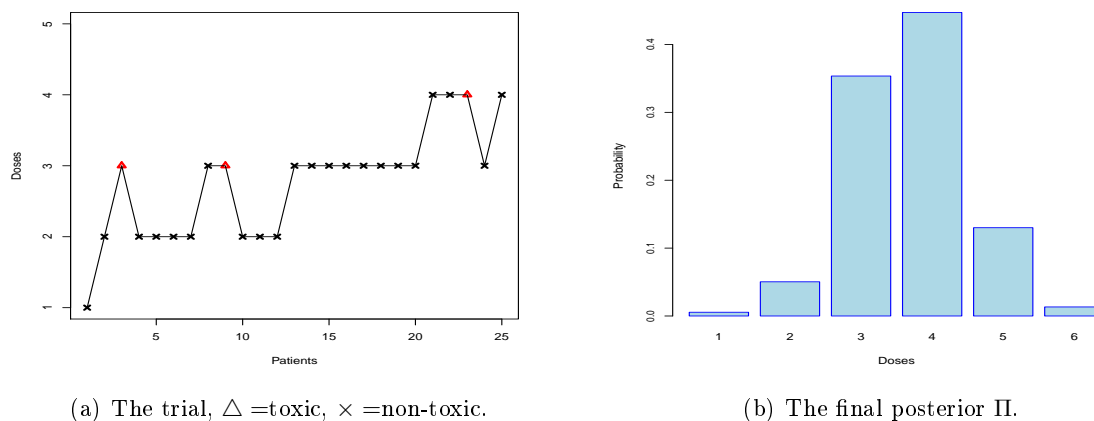


Figure 6: A balanced scenario: $\beta = (0.04, 0.08, 0.16, 0.24, 0.35, 0.45)$.

this article we have not presented any results concerning SPM and methods other than the CRM. Such comparisons can be readily carried out both in theory and in practice leading to improvements on methods currently in use. We study this more deeply in a separate paper.

Conclusion and perspectives

The central feature of the SPM is the direct modelling of the key parameter of interest, the dose itself, structured around a regression function that is not fully specified. The approach is expressed via a hierarchical Bayesian structure. So far, we have not tried to increase the flexibility of the usual prior as a way to better deal with poor model specification. Instead we replace the model by a prior that we call the prior model. The topological support of this is indexed by the parameter of main interest, the MTD. From the asymptotic standpoint, we no longer seek to obtain convergence on the set of posterior laws, but only on the particular law surrounding the parameter of interest, the MTD. This, albeit small, change in emphasis leads to improved asymptotic behaviour. In particular we obtain the almost sure convergence of the estimator of the MTD built on the observations obtained by the method.

As a by-product we obtain much more and we note that the generalization is sufficiently flexible to allow to include almost all of the currently used model-based designs as special cases. Here we have applied the generalization to the specific case of the CRM which is a strongly parametrized method. We are currently working on doing the same thing for the CCD, the mTPI and the BOIN methods, all of which can be seen to be less parametric since they do not explicitly model the relation $Y \approx f(X)$. All of these methods achieve good performance both in the treatment of patients and in the accurate locating of the MTD. Being able to put them under a single umbrella - semi-parametric dose finding methods - will enable us to better study the differences and the similarities between them and, ultimately, to construct improvements. The SPM framework allows for theoretical study, the results of

which will then apply to all of these special cases. Furthermore, the SPM can be used in its own right as a method, as it stands, and our theoretical and simulation based investigations suggest that it is at least as good, and in most cases better, than all the methods we have already tested.

ACKNOWLEDGEMENT:

The authors would like to acknowledge the many insightful comments of the reviewers and the editors. They have helped us to write a broader and more accessible text.

References

- David Azriel. A note on the robustness of the continual reassessment method. *Statistics & Probability Letters*, 82(5):902–906, 2012.
- David Azriel, Micha Mandel, and Yosef Rinott. The treatment versus experimentation dilemma in dose finding studies. *Journal of Statistical Planning and Inference*, 141(8): 2759–2768, 2011.
- James Babb, Andre Rogatko, and Shelemyahu Zacks. Cancer phase i clinical trials: efficient dose escalation with overdose control. *Statistics in medicine*, 17(10):1103–1120, 1998.
- Peter J Bickel, CAJ Klaassen, Ya’acov Ritov, and Jon A Wellner. Semiparametric inference and models. Technical report, Mimeo, University of California, Berkeley, 2005.
- K. Cheung. Coherence principles in dose-finding studies. *Biometrika*, 92(4):863–873, 2005.
- Ying Kuen Cheung and Rick Chappell. A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics*, 58(3):671–674, 2002.
- Anastasia Ivanova, Nancy Flournoy, and Yeonseung Chung. Cumulative cohort design for dose-finding. *Journal of Statistical Planning and Inference*, 137(7):2316–2327, 2007.
- Yuan Ji, Ping Liu, Yisheng Li, and B Nebiyu Bekele. A modified toxicity probability interval method for dose-finding trials. *Clinical Trials*, page 1740774510382799, 2010.
- Suyu Liu and Ying Yuan. Bayesian optimal interval designs for phase i clinical trials. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 64(3):507–523, 2015.
- John O’Quigley and Larry Z Shen. Continual reassessment method: a likelihood approach. *Biometrics*, pages 673–684, 1996.
- John O’Quigley, Margaret Pepe, and Lloyd Fisher. Continual reassessment method: a practical design for phase 1 clinical trials in cancer. *Biometrics*, pages 33–48, 1990.

John O'Quigley, Xavier Paoletti, and Jean Maccario. Non-parametric optimal design in dose finding studies. *Biostatistics*, 3(1):51–56, 2002.

Larry Z Shen and John O'Quigley. Consistency of continual reassessment method under model misspecification. *Biometrika*, 83(2):395–405, 1996.

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

Barry E Storer. Design and analysis of phase i clinical trials. *Biometrics*, pages 925–937, 1989.

Tian Tian. A note on continual reassessment method. *Statistics & Probability Letters*, 2016.

