



HAL
open science

Semi-parametric bayesian model, applications in dose finding studies

Matthieu Clertant

► **To cite this version:**

Matthieu Clertant. Semi-parametric bayesian model, applications in dose finding studies. Statistics [math.ST]. Université Pierre et Marie Curie - Paris VI, 2016. English. NNT : 2016PA066230 . tel-01822162

HAL Id: tel-01822162

<https://theses.hal.science/tel-01822162v1>

Submitted on 24 Jun 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Université Pierre et Marie Curie

Ecole doctorale 386

Sciences mathématiques de Paris Centre

Laboratoire de Statistiques Théoriques et Appliquées

Semi-parametric bayesian model, applications in dose finding studies.

Par Matthieu Clertant

Thèse de Doctorat de Mathématiques

Spécialité : Statistiques

Dirigée par Michel Broniatowski et John O'Quigley

Présentée et soutenue publiquement le 22/06/16

Devant un jury composé de :

M. David AZRIEL, Lecturer, Technion, Rapporteur

M. Gérard BIAU, Professeur, Université Paris 6, membre du jury

M. Michel BRONIATOWSKI, Professeur, Université Paris 6, Directeur

M. John O'QUIGLEY, Professeur, Université Paris 6, co-Directeur

M. Christian ROBERT, Professeur, Université Paris Dauphine, membre du jury

M. Nicolas SAVY, Maître de conférences, Université Paul Sabatier, Rapporteur

Remerciements

Durant ce travail de recherche, j'ai pu observer que la productivité du chercheur solitaire est engendrée par la bonne humeur, l'écoute, la collaboration et le plaisir au travail de la tribu à laquelle il appartient. Pour avoir eu l'occasion d'appréhender cet important résultat, je remercie chaleureusement tous les membres du LSTA.

Et plus particulièrement, mes directeurs de thèse, Michel Broniatowski et John O'Quigley, qui m'ont dès le début gratifié de leur confiance. Je retiens de leur enseignement, leur vision de la recherche, les intuitions qu'ils partagent amicalement et l'art mystérieux du café mathématique et de ses digressions. Je les remercie très sincèrement pour tous les projets qu'ils m'ont aidé à initier et pour le plaisir que j'ai à travailler avec eux.

Je souhaite exprimer ma gratitude aux directeurs du LSTA, M. Paul Deheuvels et M. Gérard BIAU pour leur accueils bienveillants, leur disponibilité et les conseils amicaux qu'ils m'ont prodigué.

Je suis très honoré que M. Nicolas Savy et M. David Azriel se soient prêtés à la relecture attentive de ma thèse. Je leur suis gré de tous les commentaires de qualité qu'ils m'ont faits et je leur exprime ma gratitude ainsi qu'à M. Christian Robert et M. Gérard Biau pour avoir accepté d'être membre du jury.

Je remercie vivement M. André Rogatko et M. Mourad Tighiouart pour leur accueil si amical dans leur laboratoire à Los Angeles.

Un grand merci aussi aux membres de l'administration du LSTA, Mme Louise Lamart pour ses grandes qualités professionnelles et humaines, Mme Corinne Van Vlieberghe pour sa gentillesse et sa précieuse collaboration durant l'organisation du symposium.

Enfin, un salut ému et reconnaissant à mes coreligionnaires-thésards, compagnons d'aventures et amis: Assia, Bruno, Cécile, Diaa, Dimibihery, Erwan, Félix, José, Kévin, Le Quyen, Lise, Lucie, Lucas, Manu, Marcio, Meggy, Mokhtar, Nazih, Pierre, Roxane, Sophie, Svetlana, Tarn et le dernier mais non le moindre à remercier, Xavier. Une pensée particulière pour ma famille, mes parents et pour toi, Layla, dont j'ai l'honneur de partager la vie.

A mes proches,

Résumé

Les Phases I sont un domaine des essais cliniques dans lequel les statisticiens ont encore beaucoup à apporter. Depuis trente ans, ce secteur bénéficie d'un intérêt croissant et de nombreuses méthodes ont été proposées pour gérer l'allocation séquentielle des doses aux patients intégrés à l'étude. Durant cette Phase, il s'agit d'évaluer la toxicité, et s'adressant à des patients gravement atteints, il s'agit de maximiser les effets curatifs du traitement dont les retours toxiques sont une conséquence. Parmi une gamme de doses, on cherche à déterminer celle dont la probabilité de toxicité est la plus proche d'un seuil souhaité et fixé par les praticiens cliniques. Cette dose est appelée la MTD (maximum tolerated dose). La situation canonique dans laquelle sont introduites la plupart des méthodes consiste en une gamme de doses finie et ordonnée par probabilité de toxicité croissante. Dans cette thèse, on introduit une modélisation très générale du problème, la SPM (semi-parametric methods), qui recouvre une large classe de méthodes. Cela permet d'aborder des questions transversales aux Phases I. Quels sont les différents comportements asymptotiques souhaitables? La MTD peut-elle être localisée? Comment et dans quelles circonstances? Différentes paramétrisations de la SPM sont proposées et testées par simulations. Les performances obtenues sont comparables, voir supérieures à celles des méthodes les plus éprouvées. Les résultats théoriques sont étendus au cas spécifique de l'ordre partiel.

La modélisation de la SPM repose sur un traitement hiérarchique inférentiel de modèles satisfaisant des contraintes linéaires de paramètres inconnus. Les aspects théoriques de cette structure sont décrits dans le cas de lois à supports discrets. Dans cette circonstance, de vastes ensembles de lois peuvent aisément être considérés, cela permettant d'éviter les cas de mauvaises spécifications.

Summary

Phase I clinical trials is an area in which statisticians have much to contribute. For over 30 years, this field has benefited from increasing interest on the part of statisticians and clinicians alike and several methods have been proposed to manage the sequential inclusion of patients to a study. The main purpose is to evaluate the occurrence of dose limiting toxicities for a selected group of patients with, typically, life threatening disease. The goal is to maximize the potential for therapeutic success in a situation where toxic side effects are inevitable and increase with increasing dose. From a range of given doses, we aim to determine the dose with a rate of toxicity as close as possible to some threshold chosen by the investigators. This dose is called the MTD (maximum tolerated dose). The standard situation is where we have a finite range of doses ordered with respect to the probability of toxicity at each dose. In this thesis we introduce a very general approach to modeling the problem - SPM (semi-parametric methods) - and these include a large class of methods. The viewpoint of SPM allows us to see things in, arguably, more relevant terms and to provide answers to questions such as asymptotic behavior. What kind of behavior should we be aiming for? For instance, can we consistently estimate the MTD? How, and under which conditions? Different parametrizations of SPM are considered and studied theoretically and via simulations. The obtained performances are comparable, and often better, to those of currently established methods. We extend the findings to the case of partial ordering in which more than one drug is under study and we do not necessarily know how all drug pairs are ordered.

The SPM model structure leans on a hierarchical set-up whereby certain parameters are linearly constrained. The theoretical aspects of this structure are outlined for the case of distributions with discrete support. In this setting the great majority of laws can be easily considered and this enables us to avoid over restrictive specifications than can results in poor behavior.

Contents

Remerciements	2
Résumé	4
Summary	5
1 Introduction	11
1.1 Le statut de la MTD	13
1.2 Une courte revue des méthodes proposées par les statisticiens	17
1.2.1 Méthodes aléatoires	17
1.2.2 Méthodes basées sur un modèle	19
1.2.3 Méthodes locales	23
1.3 Problèmes issus de la pratique des phases I	25
1.3.1 Les grades de toxicité	25
1.3.2 Informations pharmacocinétiques et pharmacodynamiques	25
1.3.3 Groupes hétérogènes	26
1.3.4 Étude conjointe de la toxicité et de l'efficacité	26
1.3.5 Étude simultanée de la toxicité de deux substances	27
1.4 Nouvelles études et résultats proposés	28
2 Modèle hiérarchique subordonné à une équation estimatrice	35
2.1 Introduction	35
2.2 Contexte et notations	36
2.3 Heuristique du modèle hiérarchique	39
2.4 Convergence au sens de Doob et convergence faible	41
2.5 Conclusion et perspectives	44
2.6 Annexe: Construction formelle du modèle hiérarchique	46
3 Semi-Parametric Dose Finding Methods	48
3.1 Basic set-up and notation	50
3.2 Parametric methods: Continual Reassessment Method	52
3.3 Semi-Parametric Models	56

3.3.1	General setting	56
3.3.2	A simple prior model specification	58
3.4	Large sample theory	61
3.5	Simulations	64
3.6	Appendix 1: Material for the proof of asymptotic results	72
3.6.1	Progressive elimination of unsuitable candidate doses	72
3.6.2	Proof of theorem 3.4.1	74
3.6.3	Proof of corollary 3.4.2	76
3.7	Appendix 2: Consistency and oscillations in the case of variable support	78
3.7.1	Asymptotic results	78
3.7.2	Progressive elimination of unsuitable candidate doses in case of variable support	80
3.7.3	Proof of theorem 3.7.1	83
4	From CRM to Local Methods	86
4.1	Introduction	86
4.2	Semi-parametric dose finding methods	87
4.3	Global models of the regression function: SPM and CRM	88
4.4	From local methods to local models	90
4.5	Simulations	96
4.6	Conclusion and prospective work	100
5	The Role of Minimal Sets in Early Phase Dose Finding Studies	102
5.1	Introduction	102
5.2	Semi-Parametric Modelling of the Maximum Tolerated Dose	103
5.3	Minimal sets, cuts and large sample behavior	105
5.4	Modelling of the Maximum Tolerated Cut (SPMc)	107
5.5	A family of allocation strategies	111
5.6	Simulations	116
5.7	Conclusion	118
5.8	Appendix: Comparison of two discrete priors	119
6	The Case of Partial Ordering	121
6.1	Introduction	121
6.2	Context and Notations	122
6.3	Modelling from an MTD point of view	123
6.3.1	General setting	123
6.3.2	A simple and coherent prior model	125
6.3.3	Asymptotical results and perspective	128
6.4	Modelling from an MTC point of view	129

6.5	Appendix: Proof of asymptotical results for po-SPM	134
6.5.1	A lemma	134
6.5.2	Proof of theorem 6.3.2	136
	Bibliographie	138

Notations

Symbole	Description
a	Dose dont la probabilité de toxicité est la plus proche supérieurement de α ou, dans quelques paragraphes, paramètre de la CRM.
b	Dose dont la probabilité de toxicité est la plus proche inférieurement de α .
A	Espace paramétrique supérieur à α ($[\alpha, 1]$ ou $[\alpha + \epsilon, 1]$).
B	Espace paramétrique inférieur à α ($[0, \alpha]$ ou $[0, \alpha - \epsilon, 1]$).
c^*	"True cut" ou contour à déterminer, limite dans la gamme D partageant les doses associées à des probabilités de toxicité inférieures à α de celles associées à des probabilités supérieures.
d^*	dose maximum tolérée à déterminer.
D	Gamme de dose ($D = \{1, \dots, m\}$).
F	Ensemble possible des scénarios de toxicité (voir β).
I, I_ϵ	Intervalle paramétrique central ($[\alpha - \epsilon, \alpha + \epsilon]$ ou $\{\alpha\}$)
m	Dernière dose de la gamme, la plus toxique.
q	Élément générique de F ou S .
S	Espace des lois de probabilité, support topologique de l'a priori $\Lambda \otimes \Pi$.
S_θ	Sous-espace de l'ensemble S , support topologique de l'a priori Λ_θ .
$S_\theta^j, S_\theta^{(i,j)}$	Sous-espace de S_θ relatif à la dose j ou (i, j) .
α	Seuil de probabilité de toxicité fixé par les praticiens.
β	Vecteur inconnu des probabilités de toxicités supposées engendrer les observations aléatoires, parfois appelé le scénario.
β_j	Probabilité d'observer une DLT en la dose j .
$\hat{\beta}_n, \tilde{\beta}_n$	Estimateurs de β au rang n .
θ	paramètre d'intérêt issu de la contrainte linéaire
$\hat{\theta}_n, \tilde{\theta}_n$	Estimateurs de θ au rang n .
Λ	Noyau de transition ou plus simplement famille de loi indexée par θ , parfois appelé le "prior model".
$\Lambda_\theta(\cdot)$ ou $\Lambda(\cdot \theta)$	A priori relatif à la classe S_θ .
Λ_θ^j	Loi de probabilité marginale en la dose j de support S_θ^j .
Π	A priori sur le paramètre d'intérêt θ .

Sigle	En toute lettre
BOIN	bayesian optimal interval
CCD	cumulative cohort distribution
CRM	continual reassessment method
EWOC	escalation with overdose control
MTC	maximum tolerated cut (contour)
MTD	maximum tolerated dose
mTPI	modified toxicity probability interval
po	partial ordering
SPM	semi-parametric methods

Chapter 1

Introduction

Un essai clinique portant sur un médicament est une recherche biomédicale organisée et pratiquée sur l'Homme en vue de recueillir des connaissances biologiques ou médicales. De telles investigations précèdent toujours l'autorisation d'un nouveau traitement et ont pour objectif d'en évaluer la toxicité potentielle et l'efficacité. Ils sont traditionnellement répartis en quatre phases:

- **Phase I** : Il s'agit de l'évaluation de la toxicité et de la tolérance de l'homme à ce médicament. Celle-ci peut donner lieu à l'étude de la cinétique et du métabolisme de la nouvelle substance. Elle est en général menée sur de petits échantillons (de 20 à 80). La nature et la fréquence de la toxicité éventuelle sont appréciées au regard des bénéfices curatifs attendus.
- **Phase II** : Aussi appelée étude pilote, elle consiste en une première analyse de l'efficacité de la substance testée (phase IIa, de 100 à 200 sujets), puis en une recherche de la dose optimale ou thérapeutique (phase IIb, de 100 à 300 sujets ou plus).
- **Phase III** : L'étude pivot a pour objectif d'établir l'efficacité comparative de la nouvelle substance. Les résultats sont obtenus au moyen d'étude statistique sur de grands échantillons.
- **Phase IV** : Celle-ci consiste en un suivi à long-terme alors que le médicament est autorisé sur le marché. Elle a pour but de dépister des effets secondaires rares ou tardifs.

Les méthodes étudiées et proposées dans cette thèse s'inscrivent dans le cadre des études précoces de phases I pour des substances destinées à traiter des maladies graves (cancer, HIV, maladies rares ...). Le pronostic vital des patients étant en jeu, on cherche à maximiser les effets curatifs du traitement. Cela justifie que l'on ne recherche pas une totale innocuité des médicaments, en particulier, lorsque les effets bénéfiques sont intrinsèquement liés aux avatars toxiques de la substance. Il y a alors un risque élevé d'observer des toxicités, c'est pourquoi les essais ne sont pas conduits chez des sujets sains, mais parmi une population de

patients malades, parfois en impasse curative, qui peuvent tirer un bénéfice des effets positifs du traitement.

La menée d'une telle étude doit amener l'ensemble des praticiens et des scientifiques qui y participent à un questionnement éthique. Le "code de Nuremberg" est un texte fondateur pour l'encadrement des essais biomédicaux sur l'être humain. Il est extrait du jugement du procès des médecins nazis (1947). Il n'a pas, à l'origine, pour vocation d'énoncer des maximes éthiques, mais de définir des critères de licéité utile au juge pénal. Ces critères permettent cependant de définir un cadre et des conditions d'"expériences acceptables" sur l'être humain. Nous reproduisons dans leur version originale accompagnée de la traduction de Amiel and Vialla (2009) trois de ces critères. Les sept autres ne sont pas moins importants et nous invitons chaleureusement le lecteur à les consulter.

- 4. The experiment should be so conducted as to avoid all unnecessary physical and mental suffering and injury. *L'expérience doit être conduite de façon telle que soient évitées toute souffrance et toute atteinte, physiques et mentales, non nécessaires.*
- 6. The degree of risk to be taken should never exceed that determined by the humanitarian importance of the problem to be solved by the experiment. *Le niveau des risques devant être pris ne doit jamais excéder celui de l'importance humanitaire du problème que doit résoudre l'expérience.*
- 8. The experiment should be conducted only by scientifically qualified persons. The highest degree of skill and care should be required through all stages of the experiment of those who conduct or engage in the experiment. *Les expériences ne doivent être pratiquées que par des personnes scientifiquement qualifiées. Le plus haut degré de compétence professionnelle doit être exigé tout au long de l'expérience, de tous ceux qui la dirigent ou y participent.*

Aujourd'hui, l'éthique des essais biomédicaux se coordonne autour de la déclaration d'Helsinki, élaborée par l'Association médicale mondiale. La première version fut adoptée en 1964 et a été l'objet de révision au cours d'assemblée générale de l'Association. La dernière a eut lieu en 2013. Une réflexion sur ces problèmes existe aussi au niveau européen. On cite notamment la directive 2001/20/CE qui fait obligations aux états membres d'harmoniser leurs dispositions législatives et administratives relatives aux essais cliniques de médicaments à usage humain et introduit la notion d'estimation de la balance "bénéfice/risque". Enfin, les avis du Conseil consultatif français d'éthique (www.ccne-ethique.fr) sont un point de repère important pour l'ensemble des personnes impliquées dans des essais sur l'être humain (avis numéros 2, 7, 11, 34, 38, 41, 58, 73, 79). Il faut rappeler que la décision finale, en dernier ressort, revient aux investigateurs et en particulier aux médecins, sous la condition du "consentement volontaire du sujet humain" (critère 1).

Conjointement aux enjeux éthiques dans lesquels le statisticien doit se sentir impliqué, le rôle de celui-ci est de proposer un protocole qui tiennent compte des informations disponibles et

de la capacité d'un traitement raisonnable des données. Par un raisonnement mathématique, il est possible d'anticiper le déroulement d'un essai et d'optimiser à chaque étape de l'étude les attentes légitimes des praticiens qui découlent du critère 4. Cela est d'autant plus vrai que l'allocation d'un patient à une dose est effectuée selon un procédé séquentiel, l'ensemble des informations disponibles sur la substance testée augmentant naturellement au cours de l'essai. Le secteur des phases I a suscité un intérêt croissant des statisticiens au cours des 30 dernières années. Ruberg (1995) dit à ce sujet:

'statisticians have gained prominence in clinical research by contributing statistical and scientific principles to phase clinical trials. Over the past two decades, statisticians have become a well-accepted part of the pharmaceutical industry and the academic community engaged in large-scale, definitive clinical trials ... Now that the phase III barriers have been broken down for statisticians, the next great challenge for statisticians is to play a larger role in phase I... study design.... This is where decision making and strategy can be paramount, and this is where statisticians who want to contribute to the overall drug development process can make their greatest impact'.

Les phases I d'essai cliniques sont conduites dans de nombreux secteurs de recherches médicales. Elles sont particulièrement importantes et intéressantes du point de vue statistiques dans la recherche sur le cancer, où la toxicité des traitements doit être contrôlée dans le même temps qu'un seuil acceptable de toxicité est défini pour traiter les patients à des doses suffisamment élevées pour être efficace. Cela vaut notamment pour les substances cytotoxiques. Dans certain cas, l'enjeu vital et la rareté des traitements efficaces impliquent que les toxicités encourues par les patients soient lourdes: myélosuppression, suppression immunitaire, alopecie, vomissements, anorexie. Le but de l'essai de phase I est alors de déterminer une dose acceptable, la MTD (Maximum Tolerated Dose), associée à un risque de toxicité voisin ou ne dépassant pas le seuil préalablement fixé, α . Lors de la phase pré-clinique, la molécule, sa structure, son effet sur les cellules sont étudiées *in vitro* et *in vivo* sur l'animal (critère 3 du "code de Nuremberg"). Si les essais sont concluants, une gamme de dose $D = \{1, \dots, m\}$ est établie à partir des retours de toxicités sur les animaux. La MTD est à déterminer parmi cette gamme de doses. Pour chacun des patients, on observe ou non une toxicité, aussi appelée DLT (Dose Limiting Toxicity).

1.1 Le statut de la MTD

Pour répondre à cet objectif, de nombreuses méthodes, ainsi que leur variantes, ont été proposées dans les trois dernières décennies. Elles peuvent être classées en deux groupes au regard de la nature mathématique qu'elles attribuent à la MTD. Les premières, appelées méthodes standards ou algorithmiques (Korn et al., 1994), considèrent que la MTD est directement observable à partir des données. La MTD est alors vue comme une statistique de l'échantillon.

Elles organisent l'allocation des doses selon une méthode dites de 'up and down'. le but est de recueillir le maximum d'information autour de et à la MTD. La plus connue, appelée '3+3', alloue les patients par groupe de 3 à une dose d , la plus petite étant choisie pour débiter l'essai:

- Si aucune toxicité n'est observée, le prochain groupe de patients sera traité à la dose $d + 1$,
- Si 1 toxicité sur 3 est observée, le prochain groupe de patients sera traité à la dose d ,
- Si 2 ou 3 toxicités sont observées, le prochain groupe de patients sera traité à la dose $d - 1$;
- En fin d'essai, la MTD est déterminée pour correspondre à certains critères: une toxicité moyenne inférieur à un certains ratios ($1/6$, $1/5$, $1/4$, $1/3$) et au moins deux toxicités observées à la dose supérieures.

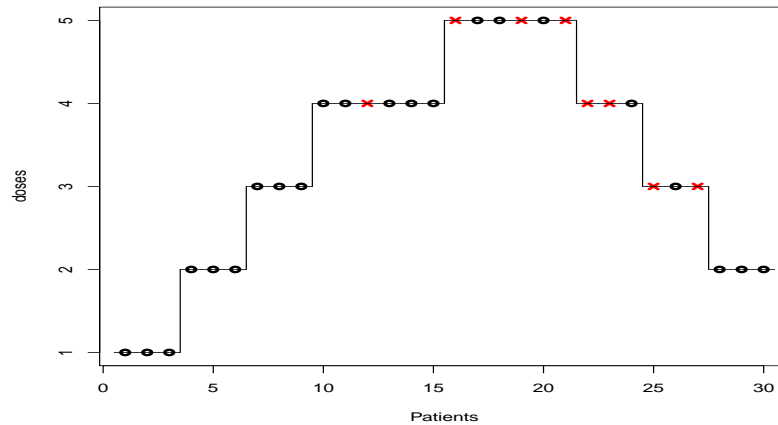


Figure 1.1: Un essai clinique simulée avec la méthode standard (toxicités: croix rouges).

Une telle méthode part du constat que le risque de toxicité augmente avec la quantité de produit et que les doses de la gamme D sont ordonnées en ce sens. Elle est issue des usages des praticiens d'essais cliniques. Le premier à avoir étudiées leurs propriétés statistiques est Storer (1989). Reiner, Paoletti, and O'Quigley (1999) disent, à propos de cette méthode, que "la probabilité de sélectionner un niveau incorrect est plus élevée qu'il n'est généralement cru". De plus, la variabilité de cette méthode conduit à traiter les patients à des doses sous-optimales ou trop toxiques. Cette variabilité dans l'allocation des doses est illustrée à la figure 1.1, sans présager de ce qu'est la MTD ou de ce que serait un bon essai clinique de phase I. Si le ratio utilisé pour la décision finale est $1/3$, ce qui semble être le plus cohérent avec l'allocation des doses en cours d'essai, les doses 3 et 4 correspondent à la définition de la MTD. Si l'on abaisse le ratio, la MTD est la dose 2. Les doses 5 et 1 qui sont rejetées quelque soit le ratio recueillent

près d'un tiers des observations. Pour cette méthode, le risque de toxicité associé à la MTD est observé directement à partir des données. En statistique, le risque est modélisé par une probabilité. C'est le sens des nombreuses méthodes proposées par les statisticiens. La MTD n'est plus alors une statistique de l'échantillon mais un paramètre à estimer. Ceci constitue le principal saut qualitatif dans les méthodes d'allocation de doses proposées en phase I. Cette modélisation repose sur l'idée que chaque patient inclus dans l'étude possède un seuil propre à partir duquel il déclenche une toxicité. Cela veut dire qu'un patient souffrant d'une DLT à une certaine dose en aurait souffert à n'importe quelle dose plus élevée. À l'inverse, si un patient tolère une certaine dose, il aurait supporté toutes les doses plus faibles. Ainsi, l'information relative à un patient se résume à un vecteur de toxicité associé à l'ensemble des doses: des "0" pour une réponse non-toxique jusqu'à une certaine dose d et des "1" pour une réponse toxique à partir de la dose $d + 1$. Dans la figure 1.2, le patient dont la fonction dose/réponse est en

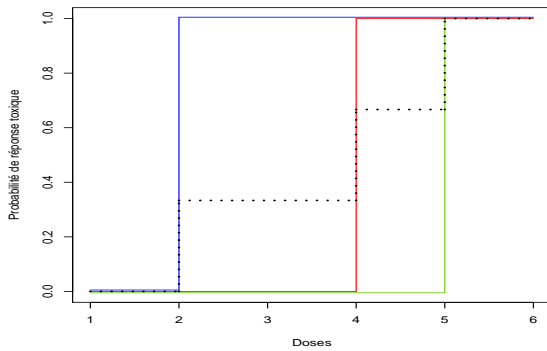


Figure 1.2: Fonction dose/réponse pour trois patients et, en pointillé, fonction dose réponse de la population des trois patients.

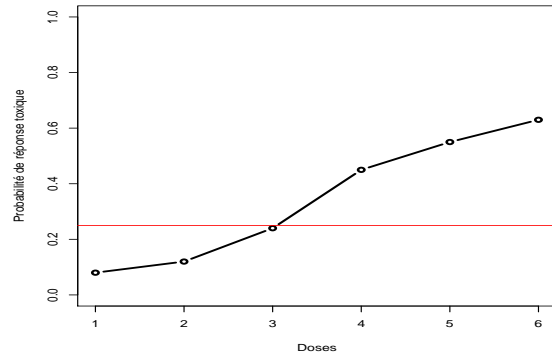


Figure 1.3: Fonction dose réponse d'une grande population et seuil $\alpha = 0.25$.

rouge déclenche une toxicité à partir de la quatrième dose; son vecteur de toxicité associé aux 6 doses est $(0, 0, 0, 1, 1, 1)$. Des fonctions dose/réponse particulières, on passe à une fonction dose/réponse pour une population dans la figure 1.3, sur le même principe par lequel on obtient la fonction de répartition d'un échantillon.

Il serait alors possible d'avoir un estimateur fiable de cette fonction de régression si un échantillon suffisamment grand de vecteurs de toxicité était connu. C'est ainsi que fonctionne la méthode optimale de O'Quigley et al. (2002), en prenant la moyenne des vecteurs de toxicité. La MTD est alors celle dont le risque de toxicité minimise la distance au seuil α . Cependant, on ne dispose jamais du vecteur de toxicité d'un patient, mais uniquement de sa réaction toxique à la dose à laquelle il a été traité. Ainsi, en pratique, la méthode optimale n'existe pas. Elle est néanmoins un outil théorique utile, permettant de borner les performances des autres méthodes. Si, lors des simulations, une méthode semble posséder de meilleurs résultats pour une certaine fonction dose/réponse, c'est que d'une certaine manière la méthode est biaisée et possède une connaissance a priori de certaines propriétés de la "vraie" fonction de régression

(aussi appelée scénario) ou de la MTD.

Finalement, pour le statisticien, le problème des phases I sous sa forme la plus simple s'exprime dans le contexte de l'analyse de régression. Un relevé séquentiel de données $(X_1^n, Y_1^n)_{n \in \mathbb{N}}$ est observé. X_n est le prédicteur, la dose sélectionnée pour le n -ième patient. Elle appartient à la gamme $D = \{1, \dots, m\}$. Y_n est la réponse en terme de toxicité de ce patient (1 pour une DLT, 0 sinon). La distribution conditionnelle de Y sachant $X = d$ est une Bernoulli de paramètre β_d . La gamme de doses D a été choisie de telle sorte que les lois de Bernoulli soient stochastiquement ordonnées.

Hypothèse 1.1.1. $\forall n \in \mathbb{N}, \beta_d = \mathbb{P}(Y_n = 1 | X_n = d)$, with: $\beta_1 < \beta_2 < \dots < \beta_m$.

Le but d'un essai de Phase I est de déterminer la MTD. Celle-ci est définie pour une certaine métrique Δ , usuellement la métrique euclidienne.

Définition 1.1.1. La MTD, notée d^* , est la dose dans la gamme D vérifiant: $\forall d \in D, \Delta(\beta_d, \alpha) \geq \Delta(\beta_{d^*}, \alpha)$.

Les méthodes d'allocation séquentielle de doses proposées par les statisticiens se basent sur l'ensemble de l'information disponible à chaque étape de l'essai. Elles ont pour objectif de satisfaire les exigences suivantes:

- Contraintes éthiques : Minimiser le nombre de patients traités à des doses trop toxiques et à des doses trop basses pour être bénéfique.
- Contraintes pragmatiques : Localiser le plus souvent possible la MTD avec le minimum de sujets dans l'essai (on note que minimiser le nombre de sujets dans un étude de Phase I est aussi une réponse aux exigences éthiques).

Cela se traduit par un ensemble de comportements observables durant un essai: la méthode doit s'adapter aux données observées, proposer rapidement une augmentation de dose en l'absence d'observation de toxicité et proposer rapidement une descente dans la gamme de dose en présence de nombreuses toxicités observées. Sous sa forme minimale, ce comportement pratique est résumé par Cheung (2005) qui introduit le concept de cohérence d'une méthode.

Définition 1.1.2. Soient \mathcal{M} une méthode et $\mathcal{M}(X_1^n, Y_1^n) = X_{n+1}$, la $n + 1$ -ième dose choisie à partir de l'échantillon au rang n . La méthode \mathcal{M} est cohérente si pour tout $d \in D$ et $n \in \mathbb{N}$, on a:

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

On observe ce comportement sur l'essai simulé présenté à la figure 1.4. La méthode utilisée est décrite dans la dernière section du chapitre 5. C'est un essai particulièrement réussi car la

fonction de régression utilisée pour le simuler est celle de la figure 1.3. La MTD est localisée à partir du huitième patient et la méthode ne recommande plus que cette dose; au total 26 des 30 patients inclus dans l'essai ont été traités à la MTD. On observe une forme de convergence de la méthode que nous détaillerons par la suite. Ce type de résultat ne survient pas tout le temps, car il est dépendant de la variabilité des observations. Parmi les méthodes décrites ci-après, celles qui sont basées sur un modèle sont celles s'approchant le plus souvent de ce type de comportement.

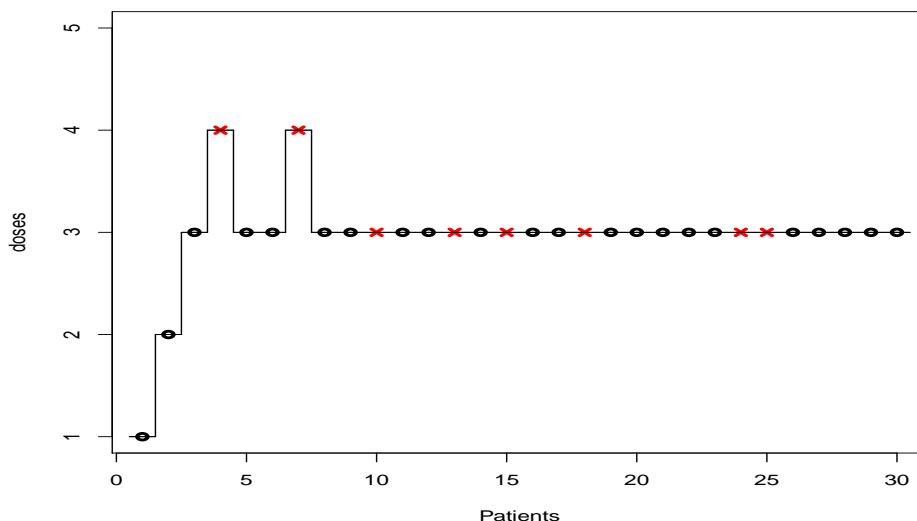


Figure 1.4: Un essai clinique simulé avec la méthode SPMc natural (voir 5).

1.2 Une courte revue des méthodes proposées par les statisticiens

Il existe de nombreuses méthodes dans la littérature statistique pour répondre aux problèmes des phases I. La plupart sont des variantes de méthodes déjà existantes. Elles sont, de notre point de vue, organisées en trois classes.

1.2.1 Méthodes aléatoires

Les règles des marches aléatoires, abrégées en RWR (pour Random Walk Rules) sont une théorie décrivant des extensions de la méthode standard du '3+3'. Ces méthodes ont d'abord été développées dans Storer (1989) pour le quantile $\alpha = 1/3$, puis pour tout les quantiles dans Durham and Flournoy (1994). La version qui est recommandée comme "la plus appropriée pour les phases I" (Rosenberger and Haines, 2002) est la RWR(A) étudiée dans (Durham et al., 1997). Comme le '3+3', elle fonctionne sur le principe du 'up and down'; un patient

étant traité à la dose d :

- si une toxicité a été observée, le prochain patient sera traité à la dose $d - 1$
- si on n'observe pas de toxicité, on choisit la dose $d + 1$ pour une certaine probabilité $p \in [0, 0.5]$ ou l'on reste à la dose en cours avec une probabilité $1 - p$.

Durham and Flournoy (1994) ont montré que si p était choisie égale à $\alpha/(1 - \alpha)$ la distribution asymptotique des doses sélectionnées était une loi unimodale centrée sur la MTD. Le but est d'explorer autour et à la MTD qui doit concentrer le plus grand nombre d'observations. Le mode ou la moyenne de la distribution des doses observées sont des estimateurs finaux de la MTD. La principale critique que l'on peut faire à ces méthodes est de faire un usage très partiel de l'information obtenue durant l'essai; en effet seul le résultat du dernier patient est pris en compte au moment de choisir la prochaine dose. Cependant, comme la distribution des doses observées se centre en la MTD, il est légitime de penser qu'asymptotiquement l'alternative d , $d - 1$ ou $d + 1$ devient plus raisonnable.

Une autre méthode aléatoire est introduite dans (Azriel et al., 2011). Elle utilise une régression isotonique, \hat{f} , pour approximer la fonction de régression générant les données. Il y a alors deux doses consécutives, d et $d + 1$, dont l'estimateur de toxicité se situe de part et d'autre du seuil α . Soit k le nombre d'observations effectuées sur l'ensemble $D' = \{d, d + 1\}$ et $\hat{d} \in D'$, la dose semblant la plus proche du seuil α au sens de la distance euclidienne pour la régression isotonique. Le choix de la prochaine dose est effectué aléatoirement dans l'ensemble D' : la dose \hat{d} est recommandée avec une probabilité $1 - 1/k$ et son alternative avec une probabilité $1/k$. Cette méthode, notée RAD (pour Randomized Allocation Design), bénéficie de très bonnes propriétés asymptotiques car elle converge en probabilité vers la MTD. La convergence est même presque sûre vers l'ensemble des deux doses consécutives associées à des probabilité de toxicité de part et d'autre du seuil α (e.g. les doses 3 et 4 dans la figure 1.3). Cela est dû à la bonne approximation locale de la fonction de régression.

Une critique globale peut être formulée à l'encontre des méthodes aléatoires sur la base de cas particuliers. On peut se demander s'il existe certaines circonstances dans lesquelles le choix de la prochaine dose paraisse évident. Supposons, comme lors d'un essai réel, que l'on ignore tout de la réalité générant les données. On cherche la MTD correspondant à un seuil $\alpha = 1/3$. 3 toxicités sur 3 patients ont été observées à la dose 4, 1 toxicité sur 4 patients à la dose 3. La dernière observation est effectuée sur la dose 3 et ne relève pas de toxicité. Sur la base de ces données, il est bien sûr impossible de savoir avec certitude laquelle des deux doses est associée au risque le plus proche du seuil α . On peut néanmoins penser qu'une attitude raisonnable serait de continuer à expérimenter la dose 3 afin, peut-être, d'obtenir un doute plus légitime pouvant nous conduire à expérimenter de nouveau la dose 4, perçue jusqu'alors comme très toxique. Cependant, dans cette circonstance, la dose 4 sera recommandée avec une probabilité $1/2$ par la RWR(A) et une probabilité $1/7$ par la RAD. Cela est dû au fait que ces méthodes ne tirent pas entièrement profit des observations précédentes pour leurs décisions locales. Les

probabilités d'allocation de la dose 4 seraient inchangées si l'on avait observées seulement 1 toxicité sur 3 patients à cette dose. Les méthodes présentées dans la partie suivante mettent en balance l'ensemble des observations pour recommander la prochaine dose.

1.2.2 Méthodes basées sur un modèle

Il existe plusieurs méthodes qui utilisent un modèle de la fonction de régression générant les données. La première introduite est la CRM (pour Continual Reassessment Method)(O'Quigley et al., 1990). C'est la méthode la plus citée et la plus utilisée lors de vrais essais cliniques parmi les méthodes basées sur un modèle. À notre connaissance, et malgré le grand nombre de méthodes proposées, la communauté scientifique ne s'est jamais accordée à dire qu'une autre méthode fonctionnait mieux au regard des exigences pratiques et éthiques tel qu'elles ont été définies (traitement global des patients et localisation finale de la MTD). Cette méthode repose sur un principe simple: utiliser une famille de fonction de la relation dose/réponse, paramétrée par une variable unidimensionnelle a , pour approximer localement, en la MTD, la fonction de régression générant les données. Cette fonction de régression est résumée par le vecteur β des toxicités à chaque dose: $\beta \in [0, 1]^{\#D}$. Initialement, ces fonctions sont définies sur un ensemble continu \mathcal{X} , contenant la gamme de dose D .

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

où l'intervalle $[A, B]$ est l'espace paramétrique. En pratique, la gamme D est souvent remplacée par une gamme D' d'éléments appartenant à l'ensemble $[0, 1]$. D' est appelée le "skeleton" (Cheung, 2011). Elle est obtenue à l'aide d'une transformation croissante t de la gamme D . Elle permet une construction plus intuitive du modèle. En effet, si on pose: $f(x|a) = t(x)^{\exp(a)}$, avec $t(D) = \{0.05, 0.10, 0.20, 0.35, 0.55, 0.70\}$ et que le scénario est $\beta = (0.05, 0.10, 0.20, 0.35, 0.55, 0.70)$, alors la fonction $f(\cdot|0)$ s'ajuste parfaitement au modèle. On ne peut pas espérer que l'ajustement ait lieu en même temps pour chacune des doses avec un unique paramètre. Cependant, on ne souhaite pas observer longtemps les doses autres que la MTD et, dans la pratique, pour que la méthode se comporte bien il suffit d'avoir un ajustement locale en chacune des doses.

Hypothèse 1.2.1. *le modèle est suffisamment riche tel que pour toute dose d , il existe une valeur du paramètre $\tilde{a}_d \in [A, B]$ tel que: $f(d|\tilde{a}_d) = \beta_d$.*

La méthode fonctionne en deux temps, la mise à jour du paramètre en fonction de l'échantillon (X_1^n, Y_1^n) , puis le choix de la dose en minimisant la distance des toxicités estimées au seuil. Le paramètre actualisé à l'étape n , noté a_n , peut être obtenu en maximisant la vraisemblance $L_n(a)$:

$$L_n(a) = \prod_{i=1}^n f(X_i|a)^{Y_i} (1 - f(X_i|a))^{1-Y_i} .$$

L'estimateur du vecteur des toxicités β est alors: $(\tilde{\beta}_{d,n})_{d \in D} = (f(d|a_n))_{d \in D}$. Il est aussi possible

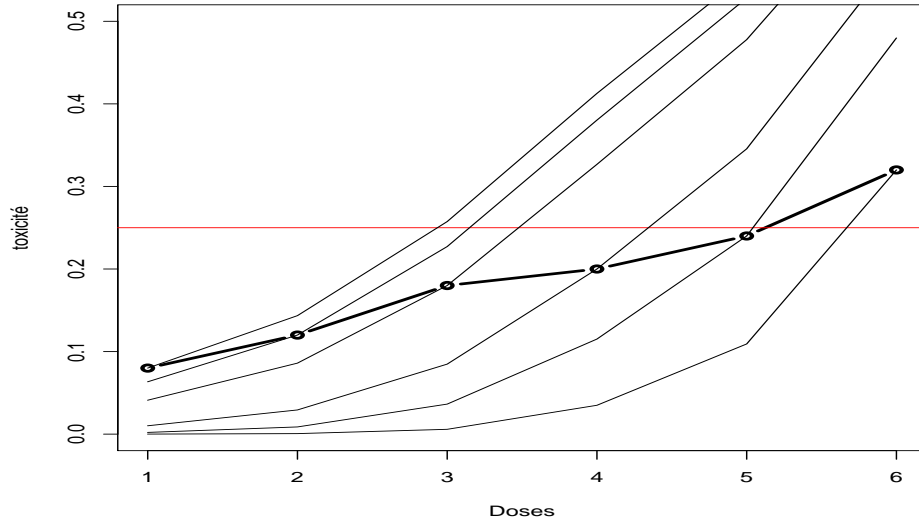


Figure 1.5: Approximation locale du modèle en chaque élément du scénario.

d'utiliser le principe bayésien. Soient G l'a priori et g sa densité. La densité de l'a posteriori G_n conditionnellement aux données est:

$$g_n(a) = \frac{L_n(a) g(a)}{\int_{a \in [A, B]} L_n(a) g(a) da}.$$

L'estimateur de β est alors l'espérance des toxicités sous l'a posteriori G_n : $\tilde{\beta}_{d,n} = \mathbb{E}_{G_n}[f(d, a)]$. Finalement, l'algorithme séquentiel est le suivant:

- Mettre à jour le paramètre a_n à l'aide des données au rang n .
- Sélectionner la dose à recommander pour le patient $n + 1$: $X_{n+1} = \mathcal{M}_1(X_1^n, Y_1^n) = \arg \min_{d \in D} |\tilde{\beta}_{d,n} - \alpha|$.

On note que le principe du maximum de vraisemblance ne peut pas s'appliquer tant qu'il n'y pas d'hétérogénéité dans les réponses des patients (au moins une toxicité et une tolérance observée). Cependant cette méthode peut fonctionner en deux étapes dissociées (Two stage CRM). On suppose que les premières observations sont des tolérances. Jusqu'à l'observation de la première DLT, on alloue des doses aux patients selon une séquence prédéterminée ou en utilisant une CRM bayésienne, ce qui revient exactement au même en calibrant correctement l'a priori. Puis, on utilise le principe du maximum de vraisemblance. Dans cette circonstance, la convergence presque sûre de la CRM est obtenue par Shen and O'Quigley (1996) sous l'hypothèse 1.2.1 et les conditions suivantes:

Hypothèse 1.2.2. Pour tout $t \in [0, 1]$ et tout x , la fonction $a \mapsto t \times (f'/f)(x|a) + (1 - t) \times (f'/(f - 1))(x|a)$ est continue et strictement monotone et on a :

- (a) For every value of the parameter a , the function $f(\cdot|a)$ is strictly increasing.
 (b) The function $f(x|\cdot)$ is continuous and strictly monotone in a , in the same direction for all x .

Hypothèse 1.2.3. Soient $(a_d)_{d \in D}$ les paramètres tels que définis dans l'hypothèse 1.2.1, on a: $\forall d \in D, \tilde{a}_d \in S = \{a : |f(d^*|a) - \alpha| < |f(d|a) - \alpha|, \forall d \neq d^*\}$.

La première de ces deux hypothèses ne pose pas de problème. Il est à la charge du statisticien de choisir des fonctions régulières et monotones. En revanche, la seconde spéculé sur la proximité du modèle avec le scénario inconnu β sur l'ensemble de la gamme de dose. Il est possible de relaxer cette hypothèse en adoptant un raisonnement plus séquentiel. Du point de vue théorique, la CRM ne converge pas si l'estimation d'une dose d , différant de la MTD, et l'ajustement locale du modèle en cette dose conduit à recommander la dose d . Ce comportement n'est souhaitable que pour la MTD. On trouve de tels exemples à la figure 1.5. La MTD est alors la dose 5 et, dans le cadre du principe du maximum de vraisemblance, les fonctions $f(\cdot|\tilde{a}_3)$, $f(\cdot|\tilde{a}_4)$ et $f(\cdot|\tilde{a}_6)$ conduise à recommander les doses 3, 4 et 6, respectivement. Ces doses constituent des points d'attraction de la méthode qui l'empêche de converger presque sûrement vers la MTD. Cheung and Chappell (2002) formalise cette analyse. Soient les intervalles

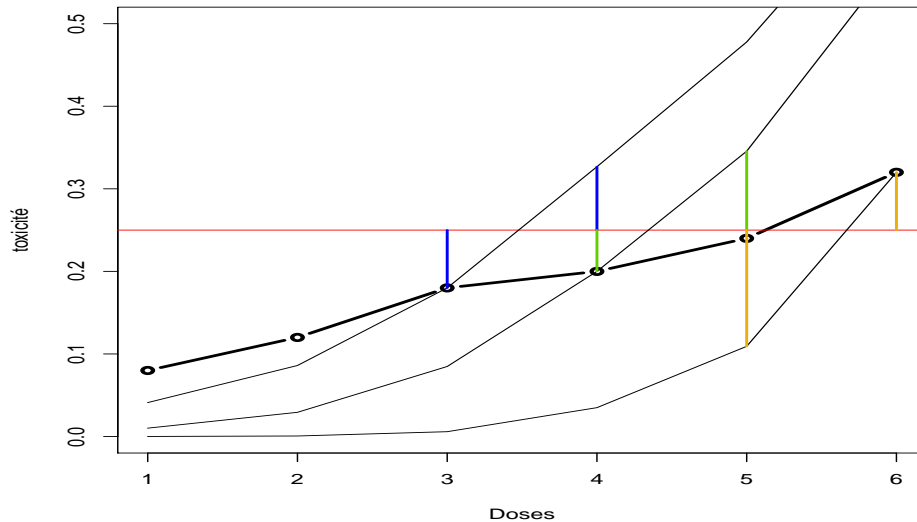


Figure 1.6: Problème de convergence en trois doses de la CRM.

$H_1 = [A, b_1)$, $H_m = (b_{m-1}, B]$ et $H_d = (b_{d-1}, b_d)$ pour $d = 1, \dots, m - 1$, où les réels b_d sont solutions de l'équation: $f(d|b_d) + f(d+1|b_d) = 2\alpha$. Cheung et Chappell montrent que la paramètre appartient à H_d si et seulement si le modèle recommande la dose d comme étant la MTD: $H_d = \{a : |f(d|a) - \alpha| < |f(d'|a) - \alpha|, \forall d' \neq d\}$. Cela leur permet de poser l'hypothèse suivante:

Hypothèse 1.2.4. $a_{d^*} \in H_{d^*}$ et $\forall d \in D \setminus \{d^*\}, a_d \notin H_d$.

Azriel (2012) démontre la convergence presque sûre de la CRM en remplaçant l'hypothèse 1.2.3 par celle de Cheung et Chappell.

Malgré ses très bonnes performances comparativement aux autres méthodes, la CRM n'est pas utilisée pleinement. Cela est dû au problème théorique posé par sa convergence et à la relative complexité de la paramétrisation de cette méthode. En effet, l'a priori et le skeleton sont souvent confondus de prime-abord et il est difficile de comprendre que ce dernier n'a pas pour but d'être proche du scénario β qui est inconnu. Sa paramétrisation dépend plutôt d'une bonne réactivité de la méthode à la nature des données binaires, et la méthode se comporte parfois mieux sous certains scénarios distants du modèle que lorsque les données sont générées sous le modèle utilisé. Dans la pratique de la paramétrisation de cette méthode, on se rend compte que les effets du skeleton et de l'a priori sont joints. Il existe de nombreuses publications sur ce sujet. Un point d'entrée est le livre de Cheung (2011).

Au vu des problèmes théoriques de convergence de la CRM, la flexibilité obtenue en prenant un paramètre de dimension plus élevée peut paraître une solution. C'est en ce sens que de très nombreuses méthodes comprenant deux paramètres ont été proposées ; en général un paramètre d'échelle et un paramètre de localisation. Cependant la nature séquentielle de l'expérience et le but qui est de recommander le plus possible de patients à la bonne dose implique que l'on ne soit pas en capacité d'estimer deux paramètres. Cela se rapproche d'une forme de non-identifiabilité. Une procédure utilisant le maximum de vraisemblance sera instable et peut mener à des sur-ajustements du modèle empêchant toute progression, tandis qu'une approche bayésienne peut marcher correctement, en pratique, en contrôlant la variance d'un des deux paramètres, ce qui signifie réduire l'impact du second paramètre de sorte que l'on soit ramené au comportement d'une méthode utilisant un unique paramètre (O'Quigley et al., 1990)(Gatsonis and Greenhouse, 1992)(Whitehead and Williamson, 1998). Il n'est pas démontré que la petite variabilité du second paramètre permette d'obtenir de meilleures performances. Cependant, certaines de ces méthodes obtiennent de bons résultats (comparables à celles de la CRM) et ont des paramétrisations intéressantes. L'EWOC (pour Escalation With Overdose Control) introduit par Babb et al. (1998) fonctionne avec une famille de fonctions à paramètres d'échelle et de localisation:

$$\forall d \in D, \beta_d = F[(d - b)/a],$$

où F est une fonction de répartition. Soient μ la MTD sur la gamme continue des doses \mathcal{X} contenant D et ρ la probabilité de toxicité à la dose d' , usuellement la première dose. Il existe un homéomorphisme entre (a, b) et (μ, ρ) :

$$a = \frac{\mu - d'}{F^{-1}(\alpha) - F^{-1}(\rho)} \text{ et } b = \frac{d'F^{-1}(\alpha) - \mu F^{-1}(\rho)}{F^{-1}(\alpha) - F^{-1}(\rho)}. \quad (1.1)$$

L'évaluation des paramètres (a, b) au rang n s'effectue à l'aide d'une approche bayésienne. La loi a posteriori sur ce couple peut être transformée en utilisant le changement de variable (1.1)

afin d'obtenir une loi a posteriori sur le couple (μ, ρ) . En intégrant selon ρ , on obtient une loi a posteriori sur μ , la MTD sur l'espace continue des doses. La dose $d \in D$ recommandée pour le prochain patient peut alors être une approximation de la médiane ou du quartile de la distribution finale sur μ . Les résultats sont similaires à ceux de la CRM pour la médiane et plus conservatifs pour le quartile.

Enfin, les méthodes proposées par Gasparini and Eisele (2000) et Whitehead et al. (2010) modélisent la fonction de régression à travers des lois a priori sur l'ensemble des probabilités de toxicités β_d pour $d \in D$. La première des méthodes est dite curve-free car elle ne nécessite pas de spécifier la famille de régression utilisée. La seconde utilise seulement l'hypothèse de monotonie de la relation dose/réponse pour évaluer sur une grille les probabilités de toxicité. Ces méthodes sont intéressantes car elles possèdent une volonté d'aller vers une solution non-paramétrique du problème.

1.2.3 Méthodes locales

Trois méthodes ont été récemment proposées: la CCD ou Cumulative Cohort Distribution (Ivanova et al., 2007), le mTPI ou modified Toxicity Probability Interval (Ji et al., 2010) et le BOIN ou Bayesian Optimal INterval (Liu and Yuan, 2015) qui comprend une version globale et locale. Ces méthodes ne modélisent pas la fonction de régression mais cherchent à déterminer la MTD par une exploration raisonnée de la gamme de doses, elles ont, elles aussi, de ce point de vue une nature non-paramétrique. Elles sont toutes trois basées sur les mêmes principes. Le premier est celui d'une décision locale de recommandation entre trois doses: la dose en cours et les deux doses de part et d'autre. Il s'agit d'augmenter d'un palier, de descendre d'un palier ou de rester à la dose précédemment observée. Cela est en adéquation avec la nature séquentielle des méthodes de phase I et la prise en compte d'une nécessaire prudence qui voudraient que la séquence de recommandation ne saute pas de dose. On note que cette contrainte paraît très légitime lorsqu'on augmente, mais peut susciter le débat en cas de diminution. Le second principe est de n'utiliser qu'une information très locale pour effectuer la recommandation: celle portée par l'ensemble des observations récoltées sur la dose en cours. Ces méthodes sont d'une certaine manière des extensions de la méthode '3+3' qui fournit une recommandation à partir des dernières observations sur la dose en cours. L'étape principale de ces méthodes peut être résumée ainsi:

- À la dose en cours $X_n = d$, mettre à jour une statistique exhaustive $S_{n,d}$ de la probabilité de toxicité à cette dose. Évaluer cette statistique pour recommander la dose X_{n+1} parmi $d - 1$, d et $d + 1$.

Pour la CCD et le BOIN local, la statistique $S_{n,d}$ est simplement le ratio de n_d^1 , le nombre de toxicités observées à la dose d , par n_d , le nombre de sujets observés à la dose d . Elle est évaluée en fonction d'un découpage de l'espace paramétrique $[0, 1]$ en trois intervalles, $[0, \eta_L]$, $[\eta_L, \eta_U]$ et $[\eta_U, 1]$ associés respectivement aux doses $d + 1$, d et $d - 1$. L'appartenance de $S_{n,d}$ à l'un des

trois intervalles conduit à recommander la dose correspondante. Le mTPI et le BOIN globale ont une approche bayésienne de l'espace paramétrique, mais cela peut être traduit, lorsqu'on utilise la statistique fréquentiste du risque de toxicité, par une variation des bornes η_L et η_U en fonction du nombre d'observations sur la dose en cours.

Ces méthodes font un usage partiel de l'information contenue dans l'échantillon. Elles ne pourraient pas fonctionner correctement sans les règles ad hoc introduites dans (Ji et al., 2010). Celles-ci ont pour but de contrôler et limiter les déplacements possibles et de fournir un estimateur de la MTD en fin d'essai. Les deux premières fonctionnent à l'aide d'une approche bayésienne posant une loi uniforme sur l'espace des paramètres $[0, 1]$.

- *Règle d'arrêt* : L'expérience s'achève si $\mathbb{P}(\beta_1 \in [\alpha, 1] | n_1^1, n_1^0) > \zeta$, avec ζ proche de 1 (par exemple $\zeta = 0.95$).
- *Exclusion de dose* : L'ensemble des doses susceptibles d'être explorées se restreint à $\{1, \dots, j - 1\}$ dès lors que $\mathbb{P}(\beta_j \in [\alpha, 1] | n_j^1, n_j^0) > \zeta$, avec ζ proche de 1 (par exemple $\zeta = 0.95$).
- *Recommandation finale* : Parmi les doses encore présentes dans l'étude, choisir celle dont l'estimateur isotonique, $\bar{\beta}_j$, est le plus proche de α . Si plusieurs doses sont dans cette circonstance: prendre la plus grande quand l'estimateur isotonic est plus petit que α et la plus petite dans le cas contraire.

Ces règles ont un caractère général et peuvent être utilisées avec toutes les méthodes de Phase I. Elles sont cependant essentielles à l'obtention de résultats satisfaisants pour les méthodes locales. Dans le cas des méthodes basées sur un modèle, leur utilisation doit être discutée. Il est notamment possible de choisir un seuil ζ plus proche de 1. L'utilisation d'une régression isotonique en fin d'essai offre a minima un point de vue différentiel sur l'estimation de la MTD.

Les méthodes introduites dans cette partie, possèdent chacune leurs avantages. La CRM obtient de très bonnes performances et utilise un paramètre unidimensionnel, ce qui est en relation avec la nature du problème posé : trouver la MTD et traiter le plus de patients possible à cette dose. L'EWOC obtient des performances comparables et donne accès à une loi a posteriori sur la MTD. Pour autant, ces méthodes ne sont pas utilisées pleinement par les investigateurs qui supervisent les phases I et il existe, encore à ce jour, de nombreux essais menés à l'aide de la méthode traditionnelle du '3+3'. Cela est peut être du à la difficulté apparente de la CRM et de l'EWOC. Les méthodes locales (CCD, mTPI, BOIN) répondent à cette demande de simplicité et obtiennent, lorsqu'elles sont utilisées avec leurs règles ad hoc, des résultats proches de ceux de la CRM et de l'EWOC. Ces méthodes ont toutes été introduites dans la situation canonique des phases I: une unique substance à tester, une réponse binaire à chaque dose et une relation dose/réponse croissante. La pratique des phases I offre des circonstances parfois plus complexes. Nous en décrivons quelques aspects dans la partie suivante.

1.3 Problèmes issus de la pratique des phases I

Les méthodes qui ont été présentées dans la partie précédente peuvent être adaptées aux différentes situations pratiques qui vont être décrites dans cette partie. Cela dit, les modélisations à adopter ne sont pas évidentes et font l'objet de publications contradictoires. Dans cette thèse, nous n'avons pas pu aborder toutes ces situations. Elles doivent alors être vues comme des perspectives de travail.

1.3.1 Les grades de toxicité

Dans les essais actuels, les grades de toxicité ne sont pas simplement la DLT (dose limiting toxicity) et l'absence de DLT. Il existe une vaste classe de toxicités potentielles, qui varient en fonction du degré, de la réversibilité, de la sévérité des effets secondaires subits. Leur spécification par les observateurs (médecins, investigateurs) dépend du type de maladies en étude, du pronostic clinique du patient traité aussi bien que des potentiels effets bénéfiques du traitement testé. Ces sévérités sont usuellement ordonné par gravité. Il s'agit alors de savoir

Table 1.1: Exemple de grades.

Label	Degré de toxicité
0	Non-toxique
1	Toxicité bénigne
2	Toxicité non-benigne
3	Toxicité sévère
4	Toxicité très sévère (DLT)

comment les praticiens veulent en tenir compte. Le plus souvent, ces grades variant entre 0 et 4, les toxicités rangées sous le label 4 sont celles dont on veut limiter la probabilité de survenue au seuil α . Dans ce cas, les toxicités rangées sous les autres labels servent à anticiper le degré 4 de toxicité. Il serait aussi possible de spécifier différents seuils que l'on ne souhaite pas dépasser pour différents types de toxicités. On note du reste que ces attentes peuvent varier en fonction des résultats observés par les praticiens, qui doivent tenir compte de la gravité effective des toxicités observées et de la mise en balance des bénéfices et des risques.

1.3.2 Informations pharmacocinétiques et pharmacodynamiques

Durant une phase I ont lieu des études pharmacocinétiques et pharmacodynamiques de la substance testée. Ces deux branches de la pharmacologie s'occupent respectivement du devenir physique de la substance testée (absorption, diffusion, transformation, évacuation) et des effets que le principe actifs produit sur l'organisme (interaction entre la substance active et les récepteurs). Évidemment de telles informations peuvent avoir un grand impact pour savoir si

un patient donné subira une DLT ou non, ou, rétrospectivement pour savoir pourquoi certains patients sont capables de tolérer certaines doses et pas d'autres. Il y a beaucoup de paramètres d'intérêts pour le pharmacologiste, e.g., l'aire sous la courbe de concentration en fonction du temps, le taux de clairance de la substance, la concentration maximale ...

Le plus souvent les informations issues de l'étude pharmacocinétique et pharmacodynamique suivent l'administration de la dose et proposent des explications rétrospectives des effets toxiques ou bénéfiques observés. Cependant il est possible de croiser ces informations avec les caractéristiques spécifiques des patients pour choisir une dose selon un procédé plus individualisé. On peut alors calibrer des modèles plus complexes, nous permettant peut-être d'obtenir de meilleurs résultats. Les travaux dans ce domaine ont été initiés par Piantadosi and Liu (1996) qui mettent à jour le potentiel de ces modélisations plus complexes.

1.3.3 Groupes hétérogènes

Dans les essais cliniques de phase I, les investigateurs recherchent un effet moyen. Dans le même état d'esprit d'individualisation que le paragraphe précédent, il est parfois possible de différencier le traitement par groupes de patients hétérogènes. Ainsi, pour certains types de traitements, il a été observé que les enfants pouvaient tolérer des doses plus élevées relativement à leur poids. De même, des patients pré-traités lourdement subiront plus facilement des effets secondaires toxiques que des patients plus légèrement pré-traités. Dans de telles situations, il semble préférable de séparer les groupes afin de localiser une MTD spécifique à chaque groupe. Usuellement, les cliniciens pratiquent deux essais séparés ou séparent l'essai en deux études différentes à partir de la première DLT observées. Cela présente l'inconvénient de ne pas utiliser l'information commune aux deux groupes.

Une CRM à double échantillon a été développée (O'Quigley et al., 1999) de telle sorte que chacun des essais tire partie des informations relatives aux deux groupes. En extension, il serait possible d'imaginer une CRM à multiples échantillons, bien qu'il faille rester réaliste vis à vis de la qualité des résultats obtenus à la lumière de la taille prévue des échantillons. Il a été montré que, sous certaines conditions usuelles pour la CRM, cette méthode converge vers les MTD de chacun des groupes. De plus, elle améliore les résultats que l'on pourrait obtenir en effectuant deux études séparées. Cette méthode se base sur l'hypothèse d'un ordre de probabilité de toxicité entre les deux groupes. On pourrait aussi imaginer une hypothèse spéculant sur la proximité des MTD relatives aux différents groupes.

1.3.4 Étude conjointe de la toxicité et de l'efficacité

L'une des particularités des phase I menée sur le VIH est que l'on dispose assez souvent d'information sur l'efficacité du traitement dans le même ordre de temps pour lequel a lieu l'observation des toxicités. Il est alors absolument nécessaire d'en tenir compte, en particulier, lorsque cette efficacité n'est pas en relation croissante avec la gamme de dose. D'un point de vue plus général, le Conseil consultatif français d'éthique exprime en conclusion de son

avis 73 son souhait de voir "modifier les modalités méthodologiques des essais de phase I en cancérologie en cherchant, en dépit des difficultés ..., à diminuer le risque de toxicité et à rechercher conjointement toxicité et efficacité". Si de telles données sont disponibles, cela change le but de l'étude. Observer trop de toxicités ou pas assez d'efficacité sont dans les deux cas un indicateur de l'échec du traitement. Le but d'une telle étude est alors de localiser sous-contraintes la dose pour laquelle le taux de succès du traitement est le plus élevé. Cette dose est alors décrite dans Zohar and O'Quigley (2006) non plus comme la MTD mais comme la MSD (Most Successful Dose).

On note de plus que les données sur l'efficacité ne sont pas toujours simples. Elle peuvent être des données appauvries de ce que l'on considérerait naturellement comme une réussite du traitement. Ainsi, on dispose toujours des données de survie du patients durant l'étude. Aussi faibles soit-ils, tous les indicateurs de l'efficacité doivent être pris en compte pour la décision finale. Un traitement bien toléré mais qui ne semble pas avoir d'effet thérapeutique doit être rejeté aussi qu'un traitement efficace provoquant de lourdes et fréquentes réactions toxiques.

Enfin, les phases I utilisent de plus en plus souvent des extensions de cohortes. C'est à dire que la MTD est d'abord localisée (phase I traditionnelle) puis l'on teste de nouvelles cohortes de patient à cette dose dans le but d'évaluer des signes précoces de son efficacité vis à vis de différents types de maladies (Iasonos and O'Quigley, 2013). Dans ce cadre, et parce qu'aucune des méthodes ne peut se prévaloir de résultats suffisamment précis sur les petits échantillons usuels en phase I, pas même la méthode optimale, il semble naturel de continuer à vouloir se positionner vis à vis de la toxicité et de l'efficacité conjointement. La gamme de doses peut être réduite sur la base des informations déjà disponibles mais il est de toute façon probable que des méthodes fonctionnant bien réduiront assez vite le champ des explorations.

1.3.5 Étude simultanée de la toxicité de deux substances

Le dernier cas spécifique que nous abordons survient lors de l'étude simultanée de deux substances A1 et A2. Il n'est alors plus possible d'ordonner les doses par quantité de produit toxique en accord avec une relation dose/réponse croissante. L'ordre et la monotonie de la relation sont partiels.

Soient (i, j) représentant la combinaison de la i -ième dose de la substance A1 et la j -ième dose de la substance A2, avec $i \in \{1, \dots, I\}$ et $j \in \{1, \dots, J\}$. L'étude se déroule alors sur l'espace produit $D = \{1, \dots, I\} \times \{1, \dots, J\}$ muni de son ordre partiel: $(i, j) < (r, s)$ si et seulement si $i < r$ et $j < s$. La relation de monotonie s'exprime alors selon l'hypothèse suivante:

Hypothèse 1.3.1. $(i, j) < (r, s) \Rightarrow \beta_{(i,j)} < \beta_{(r,s)}$.

Il existe dans ce contexte de nombreuses et récentes publications. Dans le cadre des méthodes basées sur un modèle, on cite en autres Wages et al. (2011) qui proposent d'utiliser une CRM à ordre partiel (po-CRM) concevant que l'ordre total existe mais doit être estimé. Wang and Ivanova (2005) et Braun and Jia (2013) proposent eux aussi des extensions de la CRM. Yin and

Yuan (2009a) et Yin and Yuan (2009b) basent leur modèles sur des copules. Thall et al. (2003) introduisent un modèle logistique à 6 paramètres. Une publication atypique et intéressante dans ce paysage est celle de Mander and Sweeting (2015) qui conçoivent une méthode sur une base non-paramétrique en utilisant des a priori indépendants sur la probabilité de chacune des doses (PIPE). Comme on le verra dans le cours de cette thèse, de tels a priori ont fonction de modèles.

1.4 Nouvelles études et résultats proposés

Les travaux que l'on présente se concentrent, en majorité, sur la situation canonique des phases I. Il y a une unique substance à tester et un nombre fini de doses prédéterminées. Pour une dose d testée, la réponse toxique possède une loi de Bernoulli de paramètre β_d . La relation dose/réponse suit un ordre stochastique croissant: $\forall i < j, \beta_i < \beta_j$. Le but est de trouver la dose d associée à une probabilité de toxicité β_d la plus proche possible d'un seuil α fixé par les praticiens. Pour des raisons éthiques et pragmatiques, on veut localiser le plus vite possible la bonne dose et éviter de sous-traiter ou sur-traiter les patients. Les méthodes obtenant les meilleures performances pour ces critères sont les méthodes basées sur un modèle, en particulier la CRM qui sert de référence. Les méthodes locales utilisées avec les règles ad hoc possèdent des performances très comparables. Toutes ces méthodes prennent pour objectif de traiter le prochain patient au mieux de ses intérêts au regard des informations déjà disponibles. Nous soutenons qu'elles participent toutes à un même cadre découlant naturellement du problème posé. Pour en exposer les grandes lignes, nous rappelons les principales qualités de ces différentes méthodes:

- L'unicité du paramètre de la CRM est en relation avec la question posée, trouver une dose parmi une gamme ordonnée.
- L'EWOC donne un accès direct au paramètre d'intérêt, la MTD.
- La méthode "curve free" de Gasparini and Eisele (2000) permet une plus grande adaptabilité du modèle aux données observées.
- La simplicité des méthodes locales est à l'image du choix séquentiel posé: monter, descendre ou rester à la même dose.

Le but étant de trouver la valeur d'un paramètre, la MTD, qui est une fonction marginale de ce qui peut être observé, nous pensons que ce problème est par nature semi-paramétrique. Ce qui revient à dire qu'il y a un unique paramètre d'intérêt à estimer, la MTD à valeurs dans un espace fini, et des paramètres surnuméraires qui n'ont pas à être estimés précisément mais servent à la bonne conduite de l'essai. Le paramètre d'intérêt, la MTD, est noté θ . On s'est placé dans une perspective bayésienne et θ est estimé par la mise à jour d'une loi Π . L'information contenue dans les observations est transférée à la loi Π par un modèle de la

fonction de régression, ou plutôt une série de modèles associés à chacune des valeurs possibles de θ . Un exemple simple est obtenue à la figure 1.7 où le modèle de la CRM a été utilisé pour associer à chacune des doses une fonction de régression passant exactement par le seuil α en la dose. Conditionnellement à des données, on calcule la probabilité respective de chacune de ces fonctions de régressions, ce qui revient, à multiplication de l'a priori près, à standardiser la vraisemblance associée à chacune des courbes.

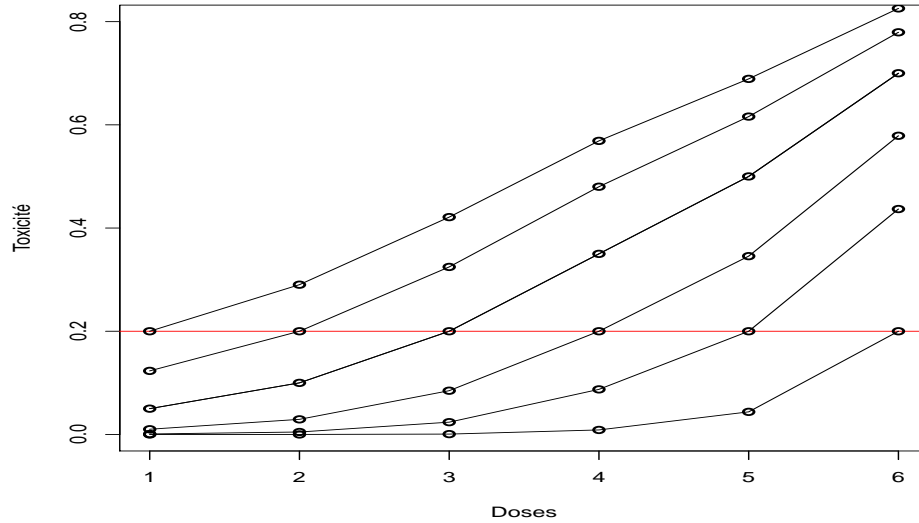


Figure 1.7: 6 fonctions de régression pour 6 MTD possibles.

On note $q \in [0, 1]^{\#D}$ le vecteur des valeurs prises par une fonction de régression sur l'ensemble des doses $D = \{1, \dots, m\}$. On a donc un simple paramètre θ à estimer et l'ajustabilité du modèle, permettant d'éviter les cas de mauvaises spécifications, est fournie par un a priori sur la plus vaste classe possible de vecteurs de régression. Cet a priori Λ est un noyau de transition ou plus simplement une famille de lois de probabilité indexées par θ . On l'appellera en anglais le "prior model", ce qui peut vouloir dire "l'a priori sur le modèle" ou "le modèle constitué par des a priori" (voir figure 1.8). La mise à jour de $\Lambda \otimes \Pi$ a lieu comme dans un modèle bayésien hiérarchique classique sauf que le paramètre d'intérêt est l'hyper-paramètre. Il s'agit donc d'intégrer marginalement selon q . Cela revient au final à une sélection parmi m modèles, chacun associant une dose particulière à la MTD.

$$\Pi_n(\theta) = \frac{\left[\int L_n(q) \Lambda(dq|\theta) \right] \Pi(\theta)}{\sum_{\theta=1}^m \left[\int L_n(q) \Lambda(dq|\theta) \right] \Pi(\theta)}, \quad (1.2)$$

avec $L_n(q)$ la vraisemblance du vecteur de régression q au rang n . La dose recommandée au rang n est celle qui maximise la loi a posteriori Π_n . Cette modélisation par sa souplesse et son

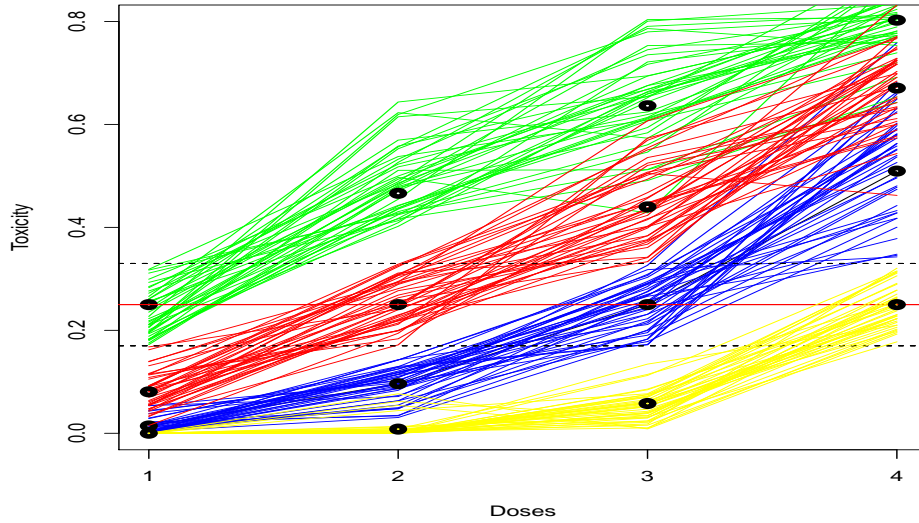


Figure 1.8: Échantillon généré sous un "prior model".

caractère général permet d'aborder des questions transversales aux méthodes de phase I. Que signifie la consistance pour les méthodes de phase I? Que sont-elles en mesure de localiser? Quelle est l'information minimale à extraire de l'échantillon? Ces questions se structurent autour d'un théorème d'impossibilité énoncé par David Azriel, Micha Mandel et Yosef Rinott dans leur article (2011) *The treatment versus experimentation dilemma in dose finding studies* : il n'existe pas de méthodes de Phase I dont l'estimateur courant, celui qui recommande la prochaine dose, converge presque sûrement vers la MTD pour tous les scénarios.

Théorème 1.4.1. (Azriel et al., 2011) Soit \mathcal{M} une méthode de phase I dont l'estimateur courant au rang n est $\mathcal{M}(X_1^n, Y_1^n) = X_{n+1}$. Un scénario β satisfaisant 1.1.1 existe tel que:

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

Ce résultat se comprend bien par une argumentation heuristique. Être la MTD est une qualité relative des doses; une dose d est la MTD si sa probabilité de toxicité β_d est plus proche du seuil α que ne le sont celles des autres doses. Si une méthode converge presque sûrement vers une dose d dans l'espace discret D , elle ne recommande plus que cette dose à partir d'un certain rang N . Les observations collectées sur les autres doses sont alors en nombre fini et il existe une probabilité non nulle que ces observations soient trompeuses et aient conduit la méthode à se focaliser sur la dose d .

Cela a amené à introduire un comportement asymptotique des méthodes de phases I qui dépendent du scénario engendrant les données. Ce comportement est en lien avec les intervalles d'indifférence définis par Cheung and Chappell (2002) pour la CRM.

Définition 1.4.1. Soit $I_\epsilon = [\alpha - \epsilon; \alpha + \epsilon]$ défini pour un certain $\epsilon \geq 0$. On considère l'ensemble des doses associées avec une probabilité de toxicité dans I_ϵ : $\mathcal{E}(I_\epsilon, \beta) = \{j \in D : \beta_j \in I_\epsilon\}$. Une méthode \mathcal{M} est dite ' ϵ -sensitive' si pour tout β tel que $\mathcal{E}(I_\epsilon, \beta) \neq \emptyset$, on a:

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

Si pour un certain scénario β , on a : $\mathcal{E}(I_\epsilon, \beta) \neq \emptyset$, une méthode ' ϵ -sensitive' converge vers une dose (ou plusieurs) dont le risque de toxicité appartient à l'intervalle I_ϵ . Et cette dose (ou cet ensemble de doses) n'est pas forcément la MTD. En revanche, les deux doses consécutives associées à des probabilités de part et d'autre du seuil α ont, grâce à l'hypothèse de croissance de la fonction de régression, un statut particulier. Elles forment l'ensemble minimal de doses dont il est nécessaire de connaître les probabilités de toxicité afin de déterminer la MTD. Cet ensemble n'est pas toujours un couple de doses, notamment dans le cas d'un ordre partiel induit par deux substances testées simultanément. Cela donne lieu à la définition suivante:

Définition 1.4.2. L'ensemble minimal M_D est le plus petit ensemble de doses pour lequel la MTD s'exprime comme une fonction des probabilités de toxicité de cet ensemble. Pour tout scénario β , on a:

- $\exists f$, une fonction telle que: $d^* = f[(d, \beta_d)_{d \in M_D}]$.
- $\forall M \subsetneq M_D$, $\nexists f$, une fonction telle que: $d^* = f[(d, \beta_d)_{d \in M}]$.

On note qu'une séquence $(X_n)_{n \in \mathbb{N}}$ converge vers un ensemble B , noté $X_n \xrightarrow{S} B$, si:

$$\sup_{x \in B} (\liminf_{n \rightarrow +\infty} d(X_n, x)) = 0.$$

La convergence de l'estimateur courant d'une méthode vers la totalité de l'ensemble minimal est un caractère souhaitable. En particulier, car c'est un comportement complémentaire et antagoniste de la propriété ' ϵ -sensitive', qui nous permet de décrire plus pleinement les caractéristiques asymptotiques des méthodes.

Définition 1.4.3. (a) Une méthode est ' ϵ -balanced', si pour tout β tel que $\mathcal{E}(I_\epsilon, \beta) = \emptyset$, on a:

$$X_n \xrightarrow{S} M_D, \text{ p.s. .}$$

(b) Une méthode est 'balanced' si la propriété précédente est vérifiée quelque soit le scénario.

Les principaux résultats asymptotiques démontrés par la suite pour les méthodes de phase I sont relatifs aux notions 'balanced' et ' ϵ -sensitive'. Après cet exposé transversal des notions intriduites, on donne un bref résumé de chacun des chapitres qui vont suivre. Ils sont rédigés comme des articles indépendants. Le premier est un développement de l'outil hiérarchique

dans un contexte général. Il est en français. Les autres portent sur les phases I et sont en anglais.

Chapitre 1: Modèle hiérarchique subordonné à une équation estimatrice

Broniatowski and Keziou (2012) introduisent une méthode d'estimation par minimisation de divergence pour des modèles satisfaisants des contraintes linéaires de paramètres inconnus. Ces contraintes linéaires sont résumées par l'équation estimatrice $(E): \int g(x, \theta) \mathbb{P}(dx|q) = 0$, où θ est inconnu. Cela amène à paramétrer un espace de lois S par θ :

$$S = \bigcup_{\theta \in \Theta} S_{\theta}, \text{ with } S_{\theta} = \left\{ q \in S : \int g(x, \theta) \mathbb{P}(dx|q) = 0 \right\}.$$

Dans ce chapitre, cet analyse est reprise et exportée à l'estimation bayésienne de lois à support fini. Ces lois sont les éléments d'un simplexe et les S_{θ} des intersections du simplexe avec des espaces vectoriels. L'estimation est hiérarchique et peut être vue comme une sélection de modèle. La loi Π est celle du paramètre d'intérêt et les Λ_{θ} des lois dans les classes S_{θ} . L'estimation est donnée par:

$$\Pi_n(\theta) = \frac{\left[\int L_n(q) \Lambda(dq|\theta) \right] \Pi(\theta)}{\sum_{\theta=1}^m \left[\int L_n(q) \Lambda(dq|\theta) \right] \Pi(\theta)},$$

avec $L_n(q)$ la vraisemblance d'un élément du simplexe. Les qualités de cette méthode sont explorées et un théorème de consistance est démontré pour des lois indépendantes mais pas forcément identiquement distribuées.

Chapitre 2: Semi-Parametric Dose Finding Methods

La CRM est une méthode dont les performances font référence en phases I, mais elle n'est pas unanimement utilisée. Ces réserves sont dues à la difficulté de sa paramétrisation et à des craintes de mauvaises spécifications en résultant. Cela est corroborée par les résultats théoriques de cette méthode qui converge sous conditions portant sur le mécanisme de génération des données. Sur la base des résultats de Cheung and Chappell (2002), et en prenant une approche bayésienne, une nouvelle méthode, la SPM, est introduite qui dérive naturellement de la CRM et en constitue une généralisation. Plusieurs résultats de consistance sont démontrés dépendants du choix des supports des lois a priori. Cette méthode peut à la fois être ' ϵ -sensitive' et ' ϵ -balanced' ou uniquement 'balanced', ce qui signifie qu'elle converge vers l'ensemble minimal et oscille sur les doses a et b de de cet ensemble. Les ratios asymptotiques

d'observations sur les doses de l'ensemble minimal sont donnés et, moyennant une mise à jour des supports des lois a priori on passe de $n_a/n_b \rightarrow D_{KL}(\beta_b, \alpha)/D_{KL}(\beta_a, \alpha)$ *p.s.*, pour la SPM classique à $n_a/n_b \rightarrow 0$ *p.s.*, si b est la MTD et vice et versa. Les performances obtenues par simulations sur de larges échantillons générés aléatoirement sont bonnes, comparables à la CRM et supérieures dans certains cas identifiés. La SPM peut être facilement implémenté et fournir des résultats très rapides lorsqu'elle est utilisée avec des lois Bêta conjuguées avec la vraisemblance.

Chapitre 3: From CRM to local methods

Dans ce chapitre, la SPM est décrite comme un cadre général de méthodes de phases I. Les méthodes locales, la CCD, le mTPI et le BOIN, qui recommandent la prochaine dose à partir des observations effectuées sur la dose en cours peuvent aussi être vue comme des cas particuliers de la SPM. L'analyse partielle qu'elles font des données peut facilement être relaxée. Cette modification leur permet de mieux tenir compte de l'ensemble des observations et d'obtenir de meilleurs résultats. Les performances comparées de ces différentes paramétrisations de la SPM sont exposées au moyen de simulations sur de grands échantillons de scénarios. Les scénarios sont obtenus par génération d'espacements de lois uniformes.

Chapitre 4: The Role of Minimal Sets in Early Phase Dose Finding Studies

Lorsque la SPM est paramétrée pour obtenir un comportement asymptotique 'balanced', elle obtient, dans les simulations effectuées, de très bons résultats en terme de traitement à la MTD. Un tel résultat introduit l'idée que rechercher l'ensemble minimal de prime-abord n'est pas contradictoire avec l'objectif de localiser la MTD. Dans le cas standard traité, l'ensemble minimal est un couple de doses consécutives (ou une unique dose) dans lequel se situe forcément la MTD. Ces ensembles sont paramétrés par les cuts (MTC, pour maximum tolerated cut). La modélisation de la SPM est reprise en remplaçant la loi Π par une loi $\tilde{\Pi}$ sur les possibles MTC. Ce nouvel ensemble de méthodes (SPMc) fonctionne en deux temps : d'abord estimer l'ensemble minimal le plus probable puis sélectionner l'une des doses de cet ensemble pour traiter le ou les prochains patients. On introduit à cette fin une famille de stratégies de sélection mettant en balance le nombre d'observations sur les doses et la distance au sens de Kullback-Leibler de l'estimateur de toxicité $\hat{\beta}_d$ au seuil α . On montre la consistance de cette méthode, en particulier le comportement 'balanced' et la convergence de n_{d^*}/n vers 1, avec n_{d^*} le nombre d'observations à la MTD. D'un point de vue plus théorique, on montre que les modèles de la SPMc sont, sous certaines conditions, équivalents à ceux de la SPM, dans le sens où ils se génèrent réciproquement. On passe alors aisément de la loi $\tilde{\Pi}$ à la loi Π à n'importe quelle étape de l'essai. Les simulations sur la SPMc sont effectuées sur de larges échantillons de

scénarios. Elles sont l'occasion de tester les stratégies de sélection qui fonctionnent le mieux.

Chapitre 5: The Case of Partial Ordering (en collaboration avec Roxane Duroux)

Les deux substances testées simultanément introduisent un ordre partiel sur la gamme $D = \{1, \dots, I\} \times \{1, \dots, J\}$. La SPM peut être adaptée à cet ordre partiel et l'on présente différentes paramétrisations et prises en compte de l'information entre deux doses non-classées. Les propriétés asymptotiques 'balanced' et ' ϵ -sensitive' sont obtenues. L'ensemble minimal est plus complexe, il se résume rarement à deux doses. Pour mieux le comprendre, on se sert de la notion de contour introduit dans (Mander and Sweeting, 2015) et on formalise cette notion. Les contours et les ensembles minimaux sont en bijection. Cela nous permet d'étendre le fonctionnement de la SPMc en proposant une modélisation dont le paramètre d'intérêt est le contour. Cette nouvelle modélisation peut être vue comme une extension de la méthode PIPE introduite par Mander and Sweeting (2015). Pour cette méthode, les auteurs n'avaient pas dissocié la loi a priori sur le paramètre d'intérêt et celle sur le modèle de sorte que les a priori possibles étaient très restreints. On démontre que la SPMc est balanced, ce qui est aussi le cas pour la méthode PIPE en tant que cas particulier.

Chapter 2

Modèle hiérarchique subordonné à une équation estimatrice

2.1 Introduction

L'espace des lois, S , est sous-paramétré par une équation estimatrice $(E): \int g(x, \theta) \mathbb{P}(dx|q) = 0$, ce qui conduit à l'indexation de cet espace par θ :

$$S = \bigcup_{\theta \in \Theta} S_\theta, \text{ with } S_\theta = \left\{ q \in S : \int g(x, \theta) \mathbb{P}(dx|q) = 0 \right\}. \quad (2.1)$$

Le paramètre θ_T est inconnu et les données observées sont engendrées sous ce paramètre. La méthode d'estimation proposée repose sur une sélection bayésienne de modèle. Celle-ci peut être vue comme un cas d'application de la modélisation hiérarchique. Dans notre circonstance, la variable d'intérêt θ indexant le modèle est l'hyper-paramètre. On se munit d'un a priori $\Lambda \otimes \Pi$ sur la totalité de l'espace S , telle que $\Lambda(\cdot|\theta)$ soit une loi sur la classe S_θ et Π la loi sur le paramètre d'intérêt θ . La loi a posteriori Π_n sur θ découle de l'intégration de la vraisemblance par rapport à la mesure $\Lambda(\cdot|\theta)$, l'a priori sur le paramètre q de premier niveau conditionnellement à θ . Sous des conditions classiques qui seront précisées par la suite, on a :

$$\begin{aligned} \Lambda(dq|\theta, x) \Pi(d\theta|x) \mathbb{P}(x) &= \mathbb{P}(x|\theta, q) \Lambda(dq|\theta) \Pi(d\theta) \\ \Pi(d\theta|x) \mathbb{P}(x) &= \int \mathbb{P}(x|\theta, q) \Lambda(dq|\theta) \times \Pi(d\theta), \end{aligned}$$

car $\int \Lambda(dq|\theta, x) = 1$. On présente la méthode dans le cas simple des variables aléatoires à valeurs dans un espace discret fini. Des résultats asymptotiques sont obtenus. Si l'espace des lois S est convexe alors les sous-espaces S_θ le sont aussi. C'est l'hypothèse principale qui nous permet d'étendre les résultats du cas iid (indépendants identiquement distribués) au cas inid (indépendants non identiquement distribués).

2.2 Contexte et notations

Le phénomène aléatoire observé est à valeurs dans un espace discret fini $\mathcal{X} = \{\dot{x}_1, \dots, \dot{x}_m\}$ muni de la tribu des parties de \mathcal{X} , notée $\mathcal{P}(\mathcal{X})$. Une loi sur ce support est un vecteur q à m composantes q_1, \dots, q_m , avec $q_j = q(\dot{x}_j)$, et l'ensemble des lois de probabilité sur ce support est en bijection canonique avec le simplexe $S_{\mathcal{X}}$, muni de sa tribu borélienne \mathcal{B}_S , plongé dans un espace réel de dimension m :

$$S_{\mathcal{X}} = \left\{ q \in \mathbb{R}^m / \sum_{i=1}^m q_i = 1 \right\}.$$

- Pour $m = 2$, tous les éléments du simplexe représentent des lois de Bernouillis. Ces lois ne possèdent qu'un seul paramètre tout comme la dimension du simplexe qui est alors représenté par le segment d'extrémité $(0, 1)$ et $(1, 0)$ dans le plan ; la lecture des coordonnées de points du segment fournissant des exemples de lois.
- Pour $m = 3$, le simplexe est de dimension 2 et vit dans l'espace à trois dimensions. Dans le tétraèdre de sommets $(0, 0, 0), (1, 0, 0), (0, 1, 0)$ et $(0, 0, 1)$, il est la base opposée à l'origine. Il n'est pas restreint à une famille de lois connues.

On note $(\Theta, \mathcal{B}_{\theta})$ l'espace des paramètres muni de sa tribu borélienne. Soit g une fonction de $\mathcal{X} \times \Theta$ dans \mathbb{R}^n et g_m sa version vectorielle sur les éléments du simplexe: $g_m(\cdot) = [g(\dot{x}_1, \cdot), \dots, g(\dot{x}_m, \cdot)]$. L'équation estimatrice (E) s'écrit: $\sum_{i=1}^m g(\dot{x}_i, \theta) q_i = \langle g_m(\theta); q \rangle = 0$, avec $\langle \cdot; \cdot \rangle$ le produit scalaire usuel. Une telle structure s'applique à une grande diversité de problème.

Exemple 2.2.1. *Quelle probabilité a-t-on de ne pas dépasser un seuil y ?* Si l'on suppose le phénomène aléatoire engendré par une unique loi de probabilité, ce problème se ramène à l'étude locale en y de la fonction de répartition. On estimera la solution en θ de l'équation estimatrice (E) avec $g(x, \theta) = \mathbf{1}_{\{x \leq y\}} - \theta$.

Exemple 2.2.2. *Quelle est la valeur moyenne d'une statistique S de l'échantillon?* Ce problème se traduit par l'étude de la fonction estimatrice (E) avec $g(x, \theta) = S(x) - \theta$. Ce type de problème dont l'exemple 1.1 fait partie, possède des propriétés simples sur lesquelles nous reviendrons.

Exemple 2.2.3. *Quelles valeurs, liées par la fonction h , prennent les moyennes des statistiques S_1 et S_2 de deux phénomènes aléatoires observés conjointement?* On note $x = (x_1, x_2)$ la variable résumant les informations des deux phénomènes aléatoires. La fonction g prend cette fois ses valeurs dans \mathbb{R}^2 : $g(x, \theta) = (S_1(x_1) - \theta, S_2(x_2) - h(\theta))$.

Exemple 2.2.4. *Qu'estime le maximum de la vraisemblance d'un modèle de loi donné?* On note $L(\cdot, \theta)$ la fonction de vraisemblance de l'échantillon associée à un θ donné. On s'intéresse alors à l'équation estimatrice (E) pour laquelle: $g(x, \theta) = \partial \ln(L(x, \theta)) / \partial \theta$.

Dans la description de ces quatre exemples d'équation estimatrice, les termes "variable aléatoire" ou "loi" ne sont pas utilisés afin de conserver la généralité de la structure. Deux phénomènes aléatoires sont discernables, au sens du modèle posé par l'équation estimatrice, si ils sont attachés à deux valeurs distinctes du paramètre θ , et c'est ce paramètre que l'on cherche à estimer. Or, différentes lois de probabilités du simplexe peuvent posséder le même paramètre; elles seront de notre point de vue équivalentes. Et l'on ne suppose pas qu'un phénomène aléatoire est engendré par une loi déterminée du simplexe mais par une classe de loi $S_\theta = \{q \in S_{\mathcal{X}} : \sum_{i=1}^m g(\dot{x}_i, \theta)q_i = 0\}$.

Exemple 2.2.1. (suite) L'équation estimatrice suivante étant posée: $(E) : \sum_{i=1}^m (\mathbf{1}_{\{\dot{x}_i \leq y\}} - \theta)q_i = 0$ pour $\dot{x}_1 < \dots < \dot{x}_k \leq y < \dot{x}_{k+1} < \dots < \dot{x}_m$, Les classes S_θ sont les lois du simplexe dont la fonction de répartition pris en y vaut θ , ce qui s'écrit: $S_\theta = \{q \in S_{\mathcal{X}} : \sum_{i=1}^k q_i = \theta\}$. Pour $\theta \in (0, 1)$ et $y \in (\dot{x}_1, \dot{x}_m)$, S_θ est un espace indénombrable.

Ainsi, la survenue d'une série de mesures sur le phénomène aléatoire considéré est modélisée par la réalisation d'une suite de variables aléatoires indépendantes dont les lois appartiennent toutes à une unique classe S_θ . Pour concevoir une telle situation, on peut imaginer le relevé d'un échantillon dont les paramètres globaux du phénomène qui l'engendre évolue sans pour autant perturber la statistique d'intérêt, ou, pour le dire autrement, l'équation estimatrice s'intéresse à ce qu'il y a de constant dans le phénomène observé sans tenir compte de ses aspects évolutifs. On pensera entre autres à des mesures biologiques variant d'un individu à l'autre tout en conservant un lien constant entre elles ou à des lois économiques pérenne en dépit des évolutions temporelles ou contextuelles influant sur les variables gérées.

Soit Ω un espace abstrait muni de sa tribu \mathcal{F} . Les données sont représentées par une suite de variables aléatoires indépendantes $(X_n)_{n \in \mathbb{N}}$; cette suite est définie sur (Ω, \mathcal{F}) et prend ses valeurs dans \mathcal{X} . S_{θ_T} est la classe à laquelle appartiennent les lois générant $(X_n)_{n \in \mathbb{N}}$. Il existe donc une suite déterministe inconnue $(Q_n)_{n \in \mathbb{N}}$ d'éléments de S_{θ_T} telle que X_i suit la loi Q_i , pour tout $i \in \mathbb{N}$.

Hypothèse 2.2.1. $(X_n)_{n \in \mathbb{N}}$ est une suite de variables aléatoires indépendantes à valeur dans \mathcal{X} et il existe $(Q_n)_{n \in \mathbb{N}} \in S_{\theta_T}^\infty$ telle que, pour tout $n \in \mathbb{N}$, X_n suit la loi $Q_n : X_n \sim Q_n$.

Pour tout entier naturel n , l'espace produit \mathcal{X}^n est muni de la tribu produit $\mathcal{P}(\mathcal{X})^{\otimes n}$. Ce sont les espaces dans lesquels vivent les échantillons $X_1^n = (X_1, \dots, X_n)$ dont la loi est notée Q_1^n . De même, on note que la suite $(X_n)_{n \in \mathbb{N}}$ prend ses valeurs dans l'espace \mathcal{X}^∞ muni de la tribu $\mathcal{P}(\mathcal{X})^{\otimes \infty}$ engendré par les cylindres de la forme $B \times \mathcal{X} \times \dots \times \mathcal{X} \times \dots$ avec $B \in \mathcal{P}(\mathcal{X})^{\otimes n}$. La famille de lois $(Q_1^n)_{n \in \mathbb{N}}$ vérifie les conditions d'application du théorème de prolongement de Kolmogorov. Pour $B \in \mathcal{P}(\mathcal{X})^{\otimes n}$, on a bien: $Q_1^{n+1}(B \times \mathcal{X}) = Q_1^n(B) \times Q_{n+1}(\mathcal{X}) = Q_1^n(B)$. Cela implique l'existence d'une unique loi, Q_1^∞ , telle que pour tout $B \in \mathcal{P}(\mathcal{X})^{\otimes n}$, on ait: $Q_1^\infty(B \times \mathcal{X} \times \dots \times \mathcal{X} \times \dots) = Q_1^n(B)$. Q_1^∞ est la loi de la suite $(X_n)_{n \in \mathbb{N}}$.

Cette modélisation n'exclut pas le cas où la suite de lois engendrant les données est constante. Dans cette circonstance, un estimateur naïf de θ_T revient à utiliser la résolution de l'équation

estimatrice en approximant la vraie loi q par la loi empirique e_n , qui comme tout autre élément du simplexe peut être vue comme une fonction de $x \in \mathcal{X}$: $e_n(x) = \frac{1}{n} \sum_{j=1}^n \mathbb{1}_{\{X_j=x\}}$, $\forall x \in \mathcal{X}$. De plus, les ensembles $\{q \in \mathbb{R}^n : \langle g_m(\theta); q \rangle = 0\}$ sont des noyaux d'applications linéaires. Les S_θ peuvent donc être considérés comme des intersections d'espace vectoriel avec le simplexe et plus simplement ils sont fermés et convexes au sens des mélanges de lois.

Définition 2.2.1. C est un ensemble de lois convexe si pour toute lois μ sur C , on a: $\int q \mu(dq) \in C$.

Ainsi, la moyenne de plusieurs éléments d'une classe S_θ appartient à la classe S_θ . Cela permet d'adapter la loi des grands nombres pour des variables indépendantes (Breiman, 1992) afin d'obtenir un premier résultat de convergence. Pour des raisons de commodité, cette convergence est présentée au sens de la norme infinie, mais elle est évidemment valable pour toutes les normes.

Définition 2.2.2. La suite des lois empiriques $(e_n)_{n \in \mathbb{N}}$ converge presque sûrement, au sens de Q_1^∞ , vers un sous-ensemble S de $S_{\mathcal{X}}$, si il existe une suite $(s_n)_{n \in \mathbb{N}}$ à valeurs dans S telle que:

$$\|e_n - s_n\|_\infty \xrightarrow[n \rightarrow +\infty]{} 0 \quad , \quad Q_1^\infty - p.s.$$

Théorème 2.2.1. Sous la condition 2.2.1, la suite des lois empiriques $(e_n)_{n \in \mathbb{N}}$ converge vers S_{θ_T} , Q_1^∞ -presque sûrement.

Démonstration. Pour tout $x \in \mathcal{X}$, on a:

$$Q_n(x) = \mathbb{E}(\mathbb{1}_{\{X_n=x\}}), \forall n \in \mathbb{N} \quad ,$$

où Q_n est ici la version fonctionnelle de la loi de X_n . On considère la suite de variables aléatoires indépendantes centrées suivantes : $(\mathbb{1}_{\{X_n = x\}} - Q_n(x))_{n \in \mathbb{N}}$. On a :

$$\mathbb{E}[(\mathbb{1}_{\{X_n = x\}} - Q_n(x))^2] = Q_n(x)(1 - Q_n(x)) < \infty \quad , \quad \forall n \in \mathbb{N}$$

et

$$\sum_{n=1}^{\infty} \mathbb{E}\left[\frac{(\mathbb{1}_{\{X_n=x\}} - Q_n(x))^2}{n^2}\right] < \infty$$

La loi des grands nombres pour des variables aléatoires indépendantes (Breiman, 1992) amène:

$$\frac{\sum_{j=1}^n (\mathbb{1}_{\{X_j = x\}} - Q_j(x))}{n} = e_n(x) - \bar{q}_n(x) \longrightarrow 0, \quad p.s.$$

avec $\bar{q}_n(x) = 1/n \sum_{j=1}^n Q_j(x)$. Comme S_{θ_T} est convexe, on a: $\bar{q}_n \in S_{\theta_T}$. Le support de \mathcal{X} est fini, ce qui achève la démonstration. \square

Dans le cadre de l'estimation bayésienne, ce résultat est un argument central des théorèmes de consistance. Il implique l'existence d'un estimateur de θ_T . Soit θ_{e_n} l'estimateur plug-in obtenu en résolvant en θ l'équation estimatrice approximée à l'aide de la loi empirique:

$\langle g_m(\theta_{e_n}); e_n \rangle = 0$, où $e_n = (e_n(\dot{x}_1), \dots, e_n(\dot{x}_m))$. L'existence d'un tel estimateur dépend de l'équation estimatrice.

Hypothèse 2.2.2. *La fonction $f : (\theta, q) \mapsto \vec{g}_m(\theta)^t \cdot q$ est C^1 et, $q \in S_{\mathcal{X}}$, sa dérivée partielle par rapport à θ est non nulle. De plus, pour tout $q \in S_{\mathcal{X}}$, il existe un unique $\theta_q \in \Theta$ tel que $\langle g_m(\theta_q); q \rangle = 0$*

Sous cette condition, on obtient la convergence presque sûre de θ_{e_n} vers θ_T .

Corollaire 2.2.1. *Sous les hypothèses 2.2.1, 2.2.2, l'estimateur θ_{e_n} converge Q_1^∞ -presque sûrement vers θ_T .*

Démonstration. On reprend les notations du théorème 1. L'hypothèse 2.2.2 correspond à celles du théorème des fonctions implicites. Ainsi, il est possible de créer localement des fonctions continues associant à tout élément $q \in S_{\mathcal{X}}$ le paramètre θ_q de la classe à laquelle il appartient. D'après l'unicité de θ_q , ces fonctions peuvent être recollées pour n'en former qu'une seule, $\varphi : q \mapsto \theta_q$, conservant la propriété de continuité. Le théorème 1 implique alors le résultat. \square

2.3 Heuristique du modèle hiérarchique

Comme dans le cadre de la statistique bayésienne classique, nous prenons pour méthodologie l'estimation d'une loi sur l'espace des paramètres d'intérêt Θ . Pour autant, il n'existe pas de vraisemblance associée à cet espace. Cependant, il est possible d'adjoindre, à toute loi sur l'espace des paramètres, une loi sur le simplexe. Pour des nécessités de définition de ces lois, on a besoin de supposer que les classes S_θ ne s'intersectent pas. C'est une conséquence immédiate de l'hypothèse 2.2.2, entraînant aussi l'identifiabilité du modèle proposé.

Dans un premier temps, on propose une estimation bayésienne naïve afin de souligner en contrepoint les qualités de la modélisation hiérarchique. On suppose pour cela que les distributions Q_1^∞ de l'échantillon sont générées par un mécanisme aléatoire dont on connaît la loi, Λ_θ , conditionnellement à la classe S_θ qui, elle, est inconnue. L'estimation de θ passe alors par l'a posteriori d'une loi Π sur l'espace des paramètres. La vraisemblance naturelle associée à la réalisation x est celle fournie par le modèle a priori, la famille $(\Lambda_\theta)_{\theta \in \Theta}$. Il suffit pour cela de prendre l'espérance, relativement à Λ_θ , de la probabilité conditionnelle d'obtenir x sachant que la donnée est générée par un élément q de la classe S_θ . En effet, cette famille d'espérance peut être normalisée afin d'obtenir une loi de probabilité sur $(\mathcal{X}, \mathcal{P}(\mathcal{X}))$. On obtient :

$$\Pi(d\theta|x) \propto \frac{\mathbb{E}_{\Lambda_\theta}[q(x)]}{\int_{\Theta} \mathbb{E}_{\Lambda_\theta}(q(x))\Pi(d\theta)} \times \Pi(d\theta).$$

et pour un échantillon complet X_1^n :

$$\Pi(d\theta|X_1^n(\omega)) \propto \prod_{j=1}^n \frac{\mathbb{E}_{\Lambda_\theta}[q(X_j(\omega))]}{\int_{\Theta} \mathbb{E}_{\Lambda_\theta}[q(X_j(\omega))]\Pi(d\theta)} \times \Pi(d\theta).$$

Cependant, la consistance de cette méthode n'est assurée que dans le cas où l'on connaît le mécanisme aléatoire de génération des lois Q_1^∞ conditionnellement à la classe S_θ . On construit un exemple de mauvaise spécification en prenant pour Λ_θ des lois de Dirac sur des éléments q_θ et une vraie loi $q_T \in S_{\theta_T}$ telle que: $\theta_F = \arg \inf_{\theta \in \Theta} D_{KL}(q_\theta \| q_T) \neq \theta_T$, où D_{KL} est la divergence de Kullback-Leibler: $D_{KL}(q_\theta \| q_T) = \sum_{i=1}^m q_T(x_i) \log(q_T(x_i)/q_\theta(x_i))$. Pour éviter ce problème, il faudrait pouvoir actualiser les lois Λ_θ afin que l'apport d'information d'une réalisation x soit conditionnel aux données déjà observées. C'est ce que nous permet le modèle hiérarchique (Good, 1980). Soit Π une loi de probabilité sur l'espace des paramètres $(\Theta, \mathcal{B}_\Theta)$, et Λ un noyau de transition allant de $\mathcal{B}_S \times \Theta$ dans $[0, 1]$, muni de sa tribu borélienne. $\Lambda \otimes \Pi$ est une loi de probabilité sur $(S_{\mathcal{X}}, \mathcal{B}_S)$. À θ fixé, la loi $\Lambda(\cdot | \theta) = \Lambda_\theta(\cdot)$ a son support restreint à la classe S_θ .

$$\forall A \in \mathcal{B}_S \quad , \quad \Lambda \otimes \Pi(A) = \int_{\Theta} \int \mathbf{1}_{\{q \in A \cap S_\theta\}} \Lambda_\theta(dq) \Pi(d\theta).$$

Le calcul de l'a posteriori hiérarchique, conditionnellement à l'échantillon X_1^n , se fait selon la formule suivante:

$$\Pi(A | X_1^n(\omega)) = \frac{\int_{\Theta} [\int \prod_{j=1}^n q(X_1^n(\omega)) \Lambda_\theta(dq)] \Pi(d\theta)}{\int_{\Theta} [\int \prod_{j=1}^n q(X_1^n(\omega)) \Lambda_\theta(dq)] \Pi(d\theta)}. \quad (2.2)$$

À chaque étape, ce calcul peut être pensé en deux temps. D'abord la mise à jour bayésienne de la loi Π , comme elle avait été proposée pour l'estimation naïve, puis la mise à jour bayésienne des lois Λ_θ selon la formule classique: $\Lambda_\theta(dq | x) \propto q(x) \Lambda_\theta(dq)$. Ainsi, pour la survenue successive de deux données x_1 et x_2 , on obtient:

$$\begin{aligned} \Pi(d\theta | x_1, x_2) &\propto \mathbb{E}_{\Lambda_{\theta(\cdot | x_1)}}[q(x_2)] \times \Pi(d\theta | x_1) \\ &\propto \mathbb{E}_{\Lambda_{\theta(\cdot | x_1)}}[q(x_2)] \times \mathbb{E}_{\Lambda_\theta}[q(x_1)] \times \Pi(d\theta). \end{aligned}$$

L'information apportée par x_2 se trouve conditionnée à l'observation x_1 à travers la loi $\Lambda_\theta(\cdot | x_1)$. On note que l'ordre de survenue des observations n'influe pas sur le résultat:

$$\begin{aligned} \mathbb{E}_{\Lambda_{\theta(\cdot | x_1)}}[q(x_2)] \times \mathbb{E}_{\Lambda_\theta}[q(x_1)] \times \pi(\theta) &= \mathbb{E}_{\Lambda_\theta} \left[\frac{q(x_2) \times q(x_1)}{\mathbb{E}_{\Lambda_\theta}[q(x_1)]} \right] \times \mathbb{E}_{\Lambda_\theta}[q(x_1)] \times \pi(\theta) \\ &= \mathbb{E}_{\Lambda_\theta}[q(x_1) \times q(x_2)] \times \pi(\theta). \end{aligned}$$

La construction formelle de cette structure hiérarchique est fournie en annexe. Ce modèle qui a pour vocation d'estimer l'hyper-paramètre θ a une nature semi-paramétrique. Le paramètre de nuisance est alors la loi q ou le mélange de loi (ce qui revient au-même par convexité) dans la classe S_{θ_T} . Si l'on suppose, comme expliqué dans la section précédente, que le phénomène aléatoire est engendré par une séquence de loi dont la moyenne ne converge pas, il est toujours possible, à condition que ces lois appartiennent à une même classe, d'établir la consistance de la méthode sous des conditions générales.

2.4 Convergence au sens de Doob et convergence faible

Le premier résultat est un théorème de convergence au sens de Doob (1949), sous sa forme introduite par Schervish (2012). Le corollaire 2.2.1 prouve l'existence d'un estimateur de la classe θ sous laquelle sont engendrées les observations. Cela entraîne la convergence de l'estimateur de Bayes sur un ensemble de mesure pleine au regard de la loi jointe de (θ, X_1^∞) .

Théorème 2.4.1. *Soit d une métrique associée à l'espace des paramètres $(\Theta, \mathcal{B}_\Theta)$. Sous les hypothèses 2.2.1 et 2.2.2, on a, pour tout $C \in \mathcal{B}_\Theta$, la convergence presque sûre par rapport à la loi jointe de (θ, X_1^∞) :*

$$d(\Pi(C|X_1^n), \mathbf{1}_{\{\theta \in C\}}) \longrightarrow 0.$$

Démonstration. On se munit de la suite de sous-tribus \mathcal{P}_n de $\mathcal{P}(\mathcal{X})^{\otimes \infty} \otimes \mathcal{B}_\Theta$ engendrée par X_1^n . Le théorème limite sur les martingales de Levy (Liese and Miescke, 2008, theorem A 34) implique l'égalité presque sûre suivante:

$$\lim_{n \rightarrow \infty} \Pi(C|X_1^n) = \lim_{n \rightarrow \infty} \mathbb{E}(\mathbf{1}_{\{\theta \in C\}} | \mathcal{P}_n) = \mathbb{E}(\mathbf{1}_{\{\theta \in C\}} | \mathcal{P}_\infty),$$

où \mathcal{P}_∞ est la tribu engendrée par la suite $(X_n)_{n \in \mathbb{N}}$. Soit φ la fonction continue qui a toute loi q du simplexe associe le paramètre θ_q (voir corollaire 2.2.1) et e_n la loi empirique de l'échantillon X_1^n . Soit A un ensemble de mesure nulle en dehors duquel $\varphi(e_n)$ atteint sa limite. On pose:

$$G(\omega) = \begin{cases} \lim_{n \rightarrow \infty} \varphi(e_n)(\omega), & \omega \in \Omega \setminus A \\ \theta_0, & \text{sinon.} \end{cases}$$

Alors, pour un θ_0 fixé, G est une variable aléatoire mesurable dans $(\Theta, \mathcal{P}_\infty)$. Or, $\mathbf{1}_{\{\theta \in C\}}$ est égale presque sûrement à $\mathbf{1}_{\{G \in C\}}$, ce qui nous donne:

$$\mathbb{E}(\mathbf{1}_{\{\theta \in C\}} | \mathcal{P}_\infty) = \mathbf{1}_{\{\theta \in C\}}, \text{ p.s. au sens de } (\theta, X_1^\infty),$$

et achève la preuve. On note que comme Θ est un espace polonais, un estimateur consistant en probabilité suffit à la preuve (par extraction d'une sous-suite d'estimateurs convergeant presque sûrement). \square

Ce résultat de consistence est satisfaisant, si le paramètre θ est supposé par nature comme issu d'un mécanisme aléatoire engendré sous la loi a priori Π . Si, en revanche, le paramètre θ est considéré comme fixe ou que la loi Π est fixée par l'observateur pour les besoins de l'estimation, il est alors nécessaire de contrôler plus finement l'ensemble négligeable sur lequel la méthode échoue afin que le vrai paramètre n'y appartienne pas. C'est l'objet du résultat qui suit et dont la démonstration reprend les outils introduit dans Freedman (1963).

On suppose que tout élément x_n de l'échantillon a une probabilité a priori non nulle de survenir: $\mathbb{P}_n(x_n) > 0, \forall n \in \mathbb{N}$. Cette hypothèse sera aisément vérifiée par l'appartenance de la suite déterministe $(Q_n)_{n \in \mathbb{N}}$ au support topologique $F(\Lambda \otimes \Pi)$ de la mesure de probabilité $\Lambda \otimes \Pi$:

Hypothèse 2.4.1. $\forall n \in \mathbb{N}, \quad Q_n \in F(\Lambda \otimes \Pi).$

L'espace des probabilités sur \mathcal{B}_Θ est muni de la topologie faible $*$ et l'on notera $\Pi_n \xrightarrow{*} \Pi$ quand $\int_\Theta f(d\Pi_n) \rightarrow \int_\Theta f(d\Pi)$ pour toute fonction continue sur Θ . Le résultat de consistance vaut pour un couple $(\Lambda \otimes \Pi, (Q_n)_{n \in \mathbb{N}})$.

Définition 2.4.1. On dit que $(\Lambda \otimes \Pi, (Q_n)_{n \in \mathbb{N}})$, avec $(Q_n)_{n \in \mathbb{N}} \in S_{\theta_T}^\infty$, est consistant si $\Pi_{n,\omega} \xrightarrow{*} \mathbb{1}_{\theta_T} Q_1^\infty$ -presque sûrement.

On introduit les ensembles suivants: $\mathcal{X}_+ = \{x \in \mathcal{X} : \exists i \in \mathbb{N}, Q_i(x) > 0\}$ et $T_+ = \{q \in \mathbb{R}^{|\mathcal{X}_+|} : \sum_{i=1}^m q(x) \leq 1\}$. La projection canonique d'un élément q sur T_+ est notée q_+ .

Enfin la notion centrale nous permettant d'obtenir une concentration de la mesure sur une base de voisinage de Θ_T est l'entropie relative de deux lois p et q : $H(p|q) = \sum_{x \in \mathcal{X}} q(x) \log(p(x))$, avec $\log 0 = -\infty$ et $0 \cdot -\infty = 0$. La fonction d'entropie $H(\cdot|q)$ est convexe et atteint son maximum en q .

Théorème 2.4.2. Sous les conditions 2.2.1, 2.2.2 et 2.4.1, le couple $(\Lambda \otimes \Pi, (Q_n)_{n \in \mathbb{N}})$ est consistant.

Démonstration. On reprend les notations du théorème 2.2.1, \bar{q}_n étant la moyenne des lois sous lesquelles sont générées les données et e_n , la loi empirique usuelle. On introduit le voisinage de \bar{q}_n dans T_+ , avec $\delta > 0$:

$$V_n = \left\{ p \in T_{\mathcal{X}_+} / \sum_{x \in \mathcal{X}_+} \frac{\bar{q}_n(x)[p(x) - \bar{q}_n(x)]^2}{[\bar{q}_n(x) + (p(x) - \bar{q}_n(x))^+]^2} < 2\delta \right\},$$

où $x^+ = x\mathbb{1}_{x>0}$, puis le voisinage de \bar{q}_n dans $S_{\mathcal{X}}$: $\tilde{V}_n = \{p \in S_{\mathcal{X}} : p_+ \in V_n\}$. Ceci nous permet de définir le voisinage de θ_T : $W_n = \{\theta \in \Theta : \tilde{V}_n \cap S_\theta \neq \emptyset\}$. La preuve sera achevée si l'on parvient à montrer que:

$$\lim_{n \rightarrow \infty} \Pi_{n,\omega}(W_n) = 1, \quad Q_1^\infty - ps \quad . \quad (2.3)$$

En effet, supposons que (2.3) soit admis. Pour $k \in \mathbb{N}$, $V_{n,k}$, $\tilde{V}_{n,k}$ et $W_{n,k}$ sont définis en remplaçant δ par δ_k dans l'expression de V_n , avec $\delta_k \rightarrow 0$ quand $k \rightarrow \infty$, et $\tilde{V}_{q,k}$ est défini en centrant le voisinage sur q à la place de \bar{q}_n . Les boules ouvertes $B_{\epsilon,q}$ de rayon $\epsilon > 0$ et de centre q sont celles associées à la topologie de la norme infinie, notée $\|\cdot\|_\infty$.

La fonction φ définie au corollaire 2.2.1, associant à toute loi q le paramètre θ de la classe à laquelle elle appartient, est C^1 sur le compact $S_{\mathcal{X}}$. Elle est donc uniformément continue. Ainsi, pour tout voisinage, W , de Θ_T , il existe un $\epsilon > 0$ tel que: $W_\epsilon = \{\theta \in \Theta : \forall q \in S_{\theta_T}, S_\theta \cap B_{\epsilon,q} \neq \emptyset\} \subset W$. Soit $\epsilon' < \epsilon/2$ et $T_{\epsilon',q} = \{x \in \mathcal{X} : q(x) > \epsilon'/m\}$. Il existe k_0 suffisamment grand tel que, pour tout $q \in S_{\mathcal{X}}$, on ait: $\forall p \in \tilde{V}_{q,k_0}, \forall x \in T_{\epsilon',q} : |p(x) - \bar{q}_n(x)| < \epsilon'/m$. Comme $\sum_{T_{\epsilon',q}} q(x) > 1 - 2\epsilon'$, la propriété précédente reste vraie pour $x \in S_{\mathcal{X}}$. Ainsi, à partir d'un certain rang k_0 , et pour tout $n \in \mathbb{N}$, \tilde{V}_{n,k_0} est inclus dans B_{ϵ,\bar{q}_n} . Finalement, pour tout

voisinage W de θ_T , il existe un $\epsilon > 0$ et un rang k_0 pour lequel, on a :

$$\forall n \in \mathbb{N}, W_{n,k} \subset W_\epsilon \subset W,$$

ce qui implique la consistance du couple $(\Lambda \otimes \Pi, (Q_n)_{n \in \mathbb{N}})$.

Il reste à montrer la proposition (2.3). On montre d'abord que :

$$\sup_{p \in T_+ \setminus V_n} -H(p|e_n) < -H(\bar{q}_n) - \delta \quad (2.4)$$

où l'entropie est naturellement étendue à l'ensemble T_+ des sous-probabilités du simplexe.

Comme la fonction $p \mapsto -H(p|e_n)$ est concave et atteint son maximum en e_n , on a :

$$\sup_{p \in T_+ \setminus V_n} -H(p|e_n) = \max_{p \in \partial V_n} -H(p|e_n)$$

On utilise l'inégalité suivante: $\forall x \geq -1, \log(1+x) \leq x - \frac{1}{2}(1+x^+)^{-2}x^2$. Cela donne, pour $p \in \partial V_n$:

$$\begin{aligned} -H(p, \bar{q}_n) + H(\bar{q}_n) &= \sum_{x \in \mathcal{X}_+} \bar{q}_n(x) \log[1 + \bar{q}_n(x)^{-1}(p(x) - \bar{q}_n(x))] \\ &\leq \sum_{x \in \mathcal{X}_+} p(x) - \sum_{x \in \mathcal{X}_+} \bar{q}_n(x) - \frac{1}{2} \sum_{x \in \mathcal{X}_+} \frac{\bar{q}_n(x)[p(x) - \bar{q}_n(x)]^2}{[\bar{q}_n(x) + (p(x) - \bar{q}_n(x))^+]^2} \\ &\leq -2\delta \end{aligned}$$

car le dernier terme est égal -2δ , tandis que les deux premiers ensembles sont négatifs.

Ainsi, on a :

$$\max_{p \in \partial V_n} -H(p|\bar{q}_n) \leq 2\delta$$

L'adaptation de la loi des grands nombres présentée dans le théorème 2.2.1 nous permet d'obtenir (2.4).

Soit n_x , le cardinal des occurrences de x dans l'échantillon au rang n . Pour $p \in S_{\mathcal{X}_+} \setminus \tilde{V}_n$, on a :

$$\prod_{x \in \mathcal{X}} p(x)^{n_x} \leq \exp(-n[H(\bar{q}_n) + \delta]), Q_1^\infty - ps.$$

Or, si $p \in S_\theta$, avec $\theta \notin W_n$ alors $p_+ \notin V_n$, ce qui amène :

$$\int_{\Theta \setminus W_n} \int \prod_{j=1}^n p(X_j(\omega)) \Lambda(dp|\theta) \Pi(d\theta) \leq \exp(-n[H(\bar{q}_n) + \delta]), Q_1^\infty - ps \quad (2.5)$$

Soient $U_n = \{p \in V_n : H(p|\bar{q}_n) < H(\bar{q}_n) + \delta\}$, $\tilde{U}_n = \{p \in S_{\mathcal{X}} : p_+ \in U_n\}$ et $W_n^U = \{\theta \in \Theta :$

$\tilde{U}_n \cap S_\theta \neq \emptyset$. Pour tout $p \in \tilde{U}_n$ et pour presque tout ω , en appliquant le théorème 1, on a :

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^n \log p[X_j(\omega) > -H(\bar{q}_n) - \delta] .$$

L'ensemble Q_1^∞ -négligeable sur lequel cette inégalité n'est pas vraie dépend de p . Le théorème de Fubini nous permet d'inverser la dépendance et d'obtenir, pour tout ω fixé l'existence d'un ensemble $\Lambda \otimes \Pi$ -presque sûr dans \tilde{U}_n tel que, sur cet ensemble, l'inégalité précédente soit vraie. On utilise alors le lemme de Fatou sur W_n^U , et $Q_1^\infty - ps$:

$$\liminf_{n \rightarrow \infty} \int_{W_n^U} \int \prod_{j=1}^n p(X_j(\omega)) \exp(n[H(\bar{q}_n) + \delta]) \Lambda(dp|\theta) \Pi(d\theta) > \Lambda \otimes \Pi(\tilde{U}_n) .$$

Comme $W_n^U \subset W_n$, $Q_1^\infty - ps$:

$$\int_{W_n} \int \prod_{j=1}^n p(X_j(\omega)) \Lambda(dp|\theta) \Pi(d\theta) \geq \Lambda \otimes \Pi(\tilde{U}_n) \exp(-n[H(\bar{q}_n) + \delta]) . \quad (2.6)$$

D'après *M3* et comme la fonction $p \mapsto H(p|\bar{q}_n)$ est continue, $\Lambda \otimes \Pi(\tilde{U}_n)$ est strictement positif. Les inégalités (2.5) et (2.6) impliquent alors :

$$\frac{\Pi_{n,\omega}(\Theta \setminus W_n)}{\Pi_{n,\omega}(W_n)} \leq \Lambda \otimes \Pi(\tilde{U}_n)^{-1} \exp(-n\delta) ,$$

ce qui prouve la proposition (2.3) et achève la démonstration. \square

On a obtenu un résultat de convergence pour l'a posteriori Π_n dans le cas inid. Le principal travail pour obtenir ce type de résultat se situe au niveau des liens entre la topologie ou la quasi-métrique liée à la divergence de Kullback-Leibler sur l'espace S des lois et la topologie envisagée sur l'espace paramétrique Θ . De tels liens peuvent être faits sous des conditions imposées à l'équation estimatrice (hypothèse 2.2.2).

2.5 Conclusion et perspectives

Ce travail offre un point de vue générale sur l'estimation statistique. En effet, de nombreuses questions pratiques peuvent être traduites par une vision semi-paramétrique du problème: question de seuil, de moment, de corrélation ... La méthode proposée, en définissant les ensembles S_θ (2.1), permet d'élaborer le modèle à partir de la question posée. Les modèles obtenus sont suffisamment souples pour permettre d'obtenir une consistance dans le cas inid. De plus, la structure hiérarchique permet de gérer séparément l'a priori sur le paramètre d'intérêt, θ , et les a priori dans les classes constituées par le noyau de transition Λ . La famille $(\Lambda_\theta)_{\theta \in \Theta}$ peut alors être vue comme une alternative adaptative aux modèles paramétriques classiques. Dans la suite, nous l'appellerons le modèle a priori (prior model).

Les résultats sont proposés dans le cas discret. Il reste à les étendre au cas continu (travail en cours). Un tel projet doit être mené en lien avec des méthodes de calcul efficace pour cette structure. Cela rejoint pour une part la constitution des a priori Λ_θ . L'idée d'intégrer de l'information à des a priori non-informatifs sous la forme de données ou de pseudo-données permet de générer le modèle à partir de données antérieures ou de même nature. Les lois conjuguées offrent alors un cadre d'implémentation efficace.

Dans la suite de cette thèse, un premier exemple d'application est développé. Les résultats valables pour des équations estimatrices le sont aussi pour des systèmes d'inéquations. Il s'agit de trouver la dose $\theta \in D$ associée à une probabilité de toxicité β_θ , la plus proche d'un seuil α prescrit par les investigateurs. La segmentation de l'espace des lois est alors donnée par:

$$F = \bigcup_{\theta \in \Theta} S_\theta, \text{ avec } S_\theta = \left\{ \beta \in [0, 1]^{|D|} : \forall j \in D, |\beta_\theta - \alpha| < |\beta_j - \alpha| \right\}.$$

Dans ce cas de figure, les données sont séquentielles. C'est un cas intéressant d'applications de la méthodologie posant des problèmes de consistance spécifiques.

2.6 Annexe: Construction formelle du modèle hiérarchique

Cette partie reprend la construction de l'a posteriori présentée dans (Liese and Miescke, 2008). On s'intéresse, pour cela, à la loi du triplet aléatoire (X, \mathbf{q}, θ) à valeurs dans l'espace probabilisable $(\mathcal{X} \times S_{\mathcal{X}} \times \Theta, \mathcal{P}(\mathcal{X}) \otimes \mathcal{B}_S \otimes \mathcal{B}_{\Theta})$. On définit cette loi jointe par la donnée de la loi de θ , de la loi conditionnelle de \mathbf{q} sachant θ puis de la loi conditionnelle de X sachant (\mathbf{q}, θ) . Soit $\mathbb{P} : \mathcal{P}(\mathcal{X}) \times S_{\mathcal{X}} \times \Theta \mapsto [0, 1]$ le noyau de transition tel que pour tout $\theta \in \Theta$:

$$\mathbb{P}(x|q, \theta) = \mathbb{P}_{(q, \theta)}(x) = \begin{cases} q(x) & , q \in S_{\theta} \\ 0 & , q \in S_{\theta^c}, \end{cases}$$

ce qui correspond à l'idée que la probabilité d'obtenir x est fixée dès lors qu'on a choisi une loi q dans une classe S_{θ} . On obtient ainsi la loi de probabilité sur $(\mathcal{X} \times S_{\mathcal{X}} \times \Theta, \mathcal{P}(\mathcal{X}) \otimes \mathcal{B}_S \otimes \mathcal{B}_{\Theta})$:

$$\mathbb{P} \otimes \Lambda \otimes \Pi(C) = \int \int \sum_{x \in \mathcal{X}} \mathbf{1}_C(x, q, \theta) \mathbb{P}(x|q, \theta) \Lambda(dq|\theta) \Pi(d\theta) \quad (2.7)$$

qui est la loi du triplet (X, \mathbf{q}, θ) . On peut généraliser cette définition à un échantillon indépendant de taille n . En effet, la loi conditionnelle jointe d'un tel échantillon est bien un noyau de transition (Kallenberg, 2006, lemme 1.38) :

$$\mathbb{P}_n(B|q, \theta) = \mathbb{P}_{(q, \theta)}^{\otimes n}(B) = \mathbb{P}_{(q, \theta)} \otimes \cdots \otimes \mathbb{P}_{(q, \theta)}(B).$$

L'existence de la loi conditionnelle de θ sachant X est assurée par le fait que $(\Theta, \mathcal{B}_{\Theta})$ et $(S_{\mathcal{X}}, \mathcal{B}_S)$ soient des ensembles de Borel, c'est à dire sont en bijection bimesurable avec un borélien inclus dans $[0, 1]$ (Liese and Miescke, 2008, théorème A37). On a alors :

$$\begin{aligned} \mathbb{P}_n \otimes \Lambda \otimes \Pi(C) &= \int \int \sum_{x \in \mathcal{X}} \mathbf{1}_C(x_1^n, q, \theta) \mathbb{P}_n(x_1^n|q, \theta) \Lambda(dq|\theta) \Pi(d\theta) \\ &= \int \sum_{x \in \mathcal{X}} \int \mathbf{1}_C(x_1^n, q, \theta) \Lambda_{\mathbb{P}_n}(dq|x_1^n, \theta) \mathbb{P}_n \Lambda(x_1^n|\theta) \Pi(d\theta) \end{aligned}$$

où, à θ fixé, le noyau de transition, $\Lambda_{\mathbb{P}_n}$, est la loi conditionnelle régulière de \mathbf{q} sachant X_1^n indexée par θ et $\mathbb{P}_n \Lambda$ est la loi marginale de X_1^n sur la classe S_{θ} . De même, on peut écrire :

$$\mathbb{P}_n \otimes \Lambda \otimes \Pi(C) = \int \sum_{x \in \mathcal{X}} \int \mathbf{1}_C(x_1^n, q, \theta) \Lambda_{\mathbb{P}_n}(dq|x_1^n, \theta) \mathbb{P}_n \Lambda(x_1^n|\theta) \Pi(d\theta) \quad (2.8)$$

$$= \sum_{x \in \mathcal{X}} \int \int \mathbf{1}_C(x_1^n, q, \theta) \Lambda_{\mathbb{P}_n}(dq|x_1^n, \theta) \Pi_{\mathbb{P}_n \Lambda}(d\theta|x) \mathbb{P}_n \Lambda(x_1^n), \quad (2.9)$$

où le noyau de transition, $\Pi_{\mathbb{P}_n \Lambda}$, est la loi conditionnelle régulière de θ sachant X_1^n que l'on voulait obtenir et $\mathbb{P}_n \Lambda \Pi$ est la loi marginale de X_1^n sur le simplexe. Le modèle bayésien obtenu

est résumé par l'égalité intuitive suivante:

$$\mathbb{P}_n \Lambda \otimes \Pi(x_1^n, d\theta) = \mathbb{P}_n \Lambda(x_1^n | \theta) \times \Pi(d\theta) = \Pi_{\mathbb{P}_n \Lambda}(d\theta | x_1^n) \times \mathbb{P}_n \Lambda \Pi(x_1^n). \quad (2.10)$$

On remarque que l'information apportée par l'échantillon sur θ est la loi conditionnelle de l'intégralité de l'échantillon sachant θ vu à travers le prisme de la loi Λ sur la classe S_θ ($\mathbb{P}_n \Lambda(x_1^n | \theta)$), ce qui est en adéquation avec le modèle hiérarchique introduit à la section 2.3. On peut alors étendre le résultat à un échantillon. Tout comme l'on avait défini la loi Q_1^∞ dans la première partie, on pose pour $B \in \mathcal{P}(\mathcal{X})^{\otimes n}$:

$$\mathbb{P}_{(q,\theta)}^{\otimes \infty}(B \times \mathcal{X} \times \dots \times \mathcal{X} \times \dots) = \mathbb{P}_{(q,\theta)}^{\otimes n}(B).$$

Cela permet de définir uniquement la loi $\mathbb{P}_{(q,\theta)}^{\otimes \infty}$. De plus, les ensembles $C \in \mathcal{P}(\mathcal{X})^{\otimes \infty}$, pour lesquels les fonctions $\theta \mapsto \mathbb{P}_{(q,\theta)}^{\otimes \infty}(C)$ sont mesurables, forment une classe monotone contenant l'algèbre des cylindres $B \times \mathcal{X} \times \dots \times \mathcal{X} \times \dots$. Le théorème de la classe monotone (Kallenberg, 2006) montre que $\mathbb{P}_{(q,\theta)}^{\otimes \infty}$ est bien un noyau de transition. Il est alors possible de reprendre la construction précédente en l'appliquant à l'espace probabilisé $(\mathcal{X}^\infty \times S_{\mathcal{X}} \times \Theta, \mathcal{P}(\mathcal{X})^{\otimes \infty} \otimes \mathcal{B}_S \otimes \mathcal{B}_\Theta, \mathbb{P}_{(q,\theta)}^{\otimes \infty} \otimes \Lambda \otimes \Pi)$.

Considérant que pour un $\omega \in \Omega$ on a, pour tout n , $x_n = X_n(\omega)$, le résumé du modèle bayésien fourni par l'égalité (8):

$$\mathbb{P}_n \Lambda \otimes \Pi(x_1^n, d\theta) = \mathbb{P}_n \Lambda(x_1^n | \theta) \times \Pi(d\theta) = \Pi_{\mathbb{P}_n \Lambda}(d\theta | x_1^n) \times \mathbb{P}_n \Lambda \Pi(x_1^n)$$

entraîne alors la formule bayésienne de l'a posteriori (2.2):

$$\Pi_{n,\omega}(A) = \Pi(A | X_1^n(\omega)) = \frac{\int_A [\int \prod_{j=1}^n q(X_1^j(\omega)) \Lambda_\theta(dq)] \Pi(d\theta)}{\int_\Theta [\int \prod_{j=1}^n q(X_1^j(\omega)) \Lambda_\theta(dq)] \Pi(d\theta)}.$$

Chapter 3

Semi-Parametric Dose Finding Methods

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 (?) - 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 the information obtained through sampling as well as to satisfy large sample convergence of the recommended dose. 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 to 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 parametrization of the regression function, so much so that it is often described as an under parametrized model. The basic rationale and operating behavior of mTPI and BOIN is the same as that of CRM and these approaches can be viewed as being non-parametric. In this paper we complete the picture describing a general approach to dose finding that can be characterized as being semi-parametric. The semi-parametric model corresponds to a collection of distributions $\{\Lambda_\theta, \theta \in D\}$, indexed by the parameter θ . This parameter assumes values on the set of doses under study. Note that, since this set is finite, we may anticipate a connection with bandit methods. This follows when we view the Bayesian sequential design for Phase I clinical trials in terms of a K-armed bandit problem. We wish to minimize the distance from a threshold α , on the basis of some chosen metric, indexed by the doses themselves, under an ethical constraint that the patients not be assigned to dose levels that are too highly toxic. "Even without the ethical constraint, we would have a non-standard bandit problem, because there is stochastic monotonicity in responses at the dose levels" (Rosenberger and Haines, 2002). Monotonicity has been seen as a key assumption in dose-finding designs and this restriction is often naturally built into the models that are used. Our intuition, and experience, convinces us that at least for the majority of applications, the monotonicity requirement ought form the basis of any initial attempts at modelling. Taking our cue from Cheung and Chappell's work on the CRM, we introduce a semi-parametric characterization of the method (Section 2). Generalizing this formula allows us to create a framework in which we can fit other methods. We immediately gain some theoretical advantages such as almost sure convergence to the MTD under weaker conditions than those currently admitted. It is also much easier to anticipate large sample behaviour in more general situations. Perhaps no less importantly the general structure allowed for by semi-parametric models enables deeper study of the various methods currently available.

3.1 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 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.

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

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 is called the MTD (maximum tolerated dose) and we denote it d^* . It is assumed that the range D has been chosen by clinical expertise so that the doses are ordered in terms of the probability of toxic response.

D 2.
$$\forall (d, d') \in D^2, \quad d < d' \Rightarrow \beta_d < \beta_{d'}.$$

This hypothesis is all the more important since 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 and hypothesis formalize this objective.

Definition 3.1.1. *A method, \mathcal{M} , is said to be coherent (Cheung, 2005) 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) .

It does not follow from this definition that the dose at step n must be entirely determined as a function of what has been already observed. Nonetheless the methods considered here verify the following hypothesis:

M 1. *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:

$$\widehat{\beta}_{d,n} = \frac{n_d^1}{n_d^0 + n_d^1} = \frac{n_d^1}{n_d}, \text{ avec } n_d^i = \sum_{j=1}^n \mathbb{1}_{\{X_n=d, Y_n=i\}}, i \in \{0, 1\}. \quad (3.1)$$

Lemma 3.1.1. *For all methods satisfying M1, we have:*

(i) *Law of large numbers:*

$$\widehat{\beta}_{d,n} \xrightarrow[n_d \rightarrow +\infty]{} \beta_d, \text{ a.s.}$$

(ii) *Law of the iterated logarithm: for all $\beta_d \in]0, 1[$,*

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

$$\text{with: } \sigma_d = \sqrt{\beta_d(1 - \beta_d)}.$$

Proof. (i) This is shown in Azriel, Mandel, and Rinott (2011, lemma 3). (ii) Let

$$M_n = \sum_{i=1}^n \mathbb{1}_{\{X_i=d\}} (Y_i - \beta_d),$$

a martingale adapted to the filtration $\sigma(X_1^n, Y_1^n)$. For $n \geq 1$:

$$T_n = M_n - M_{n-1}, \quad s_n^2 = \sum_{i=1}^n \mathbb{E} [T_n^2 | \sigma(X_1^n, Y_1^n)] = n_d \beta_d (1 - \beta_d)$$

and $u_n = \sqrt{2 \log(\log(s_n^2))}$. Let $K_n = s_{n-1}^{-a}$, for some arbitrary a belonging to $]0, 1[$. We then have:

$$K_n \in \sigma(X_1^{n-1}, Y_1^{n-1}), \quad K_n \xrightarrow[n_d \rightarrow +\infty]{} 0, \quad |T_n| \leq K_n \frac{s_n}{u_n},$$

results that allow us to apply Theorems 1 and 2 shown in Stout (1970). \square

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 final 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 3.1.1. *Let \mathcal{M} be a method satisfying M1. A scenario β satisfying D2 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 becomes asymptotically negligible and, of course, we may then make an incorrect recommendation however large n .

Example 3.1.1. *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: $n_3^1 = 2$, $n_3 = 5$ et $\hat{\beta}_{3,n} = 0.4$. 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 hypothesis of monotonicity D2 allow us to eliminate almost surely dose 1 from the candidate doses for being the MTD.*

According to the monotonicity property D2, 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 proposed in this article arise in a natural way from a critical analysis of the asymptotic properties of the CRM (O’Quigley et al., 1990). The CRM can be seen as a starting point from which we generalize in order to construct a flexible context within which sequential decision making concerning the MTD can be made. 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 implementation in difficult situations. In the next two sections we study this in the context of a generalization of the CRM. In further unpublished work we study more deeply this generalization as it relates to the CCD, mTPI and BOIN methods.

3.2 Parametric methods: Continual Reassessment Method

In this section we recall the principle features of the continual reassessment method. This helps the flow of our presentation, helps motivate our development and, finally, enables us

to clarify those features that are key and those that are not really essential. We steer this description in a direction that allows the further developments of the following section to arise quite naturally.

The CRM is a sequential inferential method introduced by O'Quigley, Pepe, and Fisher (1990) in order to address the problems in the design and analysis of Phase I dose finding studies. The 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 . In practice, the range D is often replaced by a range D' where the elements belong to $[0, 1]$, known as the "skeleton" (Cheung, 2011). The range D' results from an increasing monotonic transformation of D . It allows the construction of a more intuitive fit between the working model and the reality.

Example 3.2.1. *The power model is defined by $f(x|a) = t(x)^{\exp(a)}$, where t is an increasing monotone transformation called the skeleton: $t(D) = D'$. Suppose that $D' = \{0.05, 0.10, 0.20, 0.35, 0.55, 0.70\}$ and $\beta = (0.05, 0.10, 0.20, 0.35, 0.55, 0.70)$, such that for $a = 0$, the model reproduces the particular scenario chosen (reality).*

From a probabilistic point of view, the situation described in this example never happens. In principle, the regression model would not allow a perfect overall fit to the reality and a good approximation on the whole range of dose levels is rare. This is a consequence of the very simple parameterization. O'Quigley, Pepe and Fisher (1990) studied a richer parametrization corresponding to the 2-parameter logistic model. However, the very limited amount of information early in the trial and the requirement of the method to furnish regular sequential decisions leads to difficulties of over fitting. Specifically, for two-parameter designs, the occurrence of a toxicity at a dose d , in the initial steps of the study, can lead to the impossibility of ever visiting any other dose $d' \geq d$ regardless of the outcomes, i.e., an unbounded number of non-toxicities at level d will never result in escalation to dose $d + 1$ (Cheung, 2011). In view of the overall poor level of parameterisation, the CRM should generally be considered to provide a poor model specification. Nonetheless, the obtained results are very good with respect to the criteria PCS and TR. (O'Quigley, Paoletti, and Maccario (2002) ; O'Quigley and Zohar (2006)). The CRM, denoted \mathcal{M}_1 in its functional form, is a method that responds to a problem itself of a sequential nature. It finds a natural functional expression in terms of an iterative algorithm.

Step 1. *The current amount of information (X_1^n, Y_1^n) , is used to update our estimate a_n of the parameter a .*

The updated estimate of the parameter may be carried out by maximizing the likelihood L_n :

$$L_n(a) = \prod_{i=1}^n f(X_i|a)^{Y_i} (1 - f(X_i|a))^{1-Y_i}.$$

The natural estimator of β according to the CRM design is: $\tilde{\beta}_{d,n} = f(d|a_n)$. A Bayesian approach is also possible. Denote the prior distribution and density by G and g respectively. Let g_n be the posterior density given the observations (X_1^n, Y_1^n) and G_n the associated probability's distribution:

$$g_n(a) = \frac{L_n(a) g(a)}{\int_{a \in [A, B]} L_n(a) g(a) da}.$$

The expected value of a with respect to g_n provides an estimate of the parameter. Following the same principle, it is also possible to directly obtain estimators of the vector β : $\tilde{\beta}_{d,n} = \mathbb{E}_{G_n}[f(d, a)]$.

Step 2. *The function $f(\cdot|a_n)$ enables us to select the dose whose estimated probability of toxicity is the closest to the desired target: $X_{n+1} = \mathcal{M}_1(X_1^n, Y_1^n) = \arg \min_{d \in D} |\tilde{\beta}_{d,n} - \alpha|$.*

In order for the algorithm's dynamics to provide a satisfactory solution to the problem, and taking account of the ordering expressed via D2, some additional assumptions are required on the family of functions f and their ability to adapt locally to the unknown reality β :

M 2. *For all $t \in [0, 1]$ and all x , the function $a \mapsto t \times (f'/f)(x|a) + (1-t) \times (f'/(f-1))(x|a)$ is continuous and strictly monotone.*

M 3. (a) *For every value of the parameter a , the function $f(\cdot|a)$ is strictly increasing.*
 (b) *The function $f(x|\cdot)$ is continuous and strictly monotone in a , in the same direction for all x .*

F 1. *The model is rich enough so that, for any dose d , there exists a value of the parameter $a_d \in [A, B]$ such that: $f(d|a_d) = \beta_d$.*

The hypothesis M3(a) translates the physical aspect of the ordering whereas M3(b), by ensuring that the functions $f(\cdot|a)$ do not cross, assures the coherence and the identifiability of the method. M2 expresses the usual constraints on the regularity of the model. The local fit of the model to the reality is guaranteed by F1. This hypothesis translates the goal of the CRM: local convergence of the model to the MTD.

Theorem 3.1.1 rules out the existence of a method providing almost sure convergence to the MTD regardless of the circumstances. However, such a property can be obtained by restricting the degree to which overall fit can depart from reality. The first such theorem of strong consistency was given by Shen and O'Quigley (1996). In order to obtain

their result they hypothesized some restrictions on the differences between the model and the reality so that fitting locally any β_d leads to choosing the MTD: $\forall d \in D, a_d \in S = \{a : |f(d^*|a) - \alpha| < |f(d|a) - \alpha|, \forall d \neq d^*\}$. This hypothesis puts limits on the degree of poor model specification. Cheung and Chappell (2002) provided an interpretation of the poor model specification, leading to a relaxing of certain conditions and still obtaining strong consistency. Their hypothesis leans on the definition of the intervals $H_1 = [A, b_1)$, $H_m = (b_{m-1}, B]$ and $H_d = (b_{d-1}, b_d)$ for $d = 1, \dots, m - 1$, where the values b_d solves: $f(d|b_d) + f(d+1|b_d) = 2\alpha$. Cheung and Chappell show that the parameter belongs to H_d if and only if the model recommends dose d as the MTD: $H_d = \{a : |f(d|a) - \alpha| < |f(d'|a) - \alpha|, \forall d' \neq d\}$. This breakdown of the parameter space together with the sequential nature of CRM leads to the following assumption.

F 2.
$$a_{d^*} \in H_{d^*} \quad \text{and} \quad \forall d \in D \setminus \{d^*\}, a_d \notin H_d.$$

Strong consistency of the CRM is shown under this condition in Azriel (2012). This method works by extrapolation (see step 2). The estimator of the MTD is the closest dose to the specified target according to the model. This procedure based on extrapolation enables the CRM to obtain very good results with respect to the ethical and pragmatical criterion, but this same extrapolation can be a cause of difficulty in cases of specifications outside of the hypothesis F2. This assumption can not be verified because it requires control over the reality

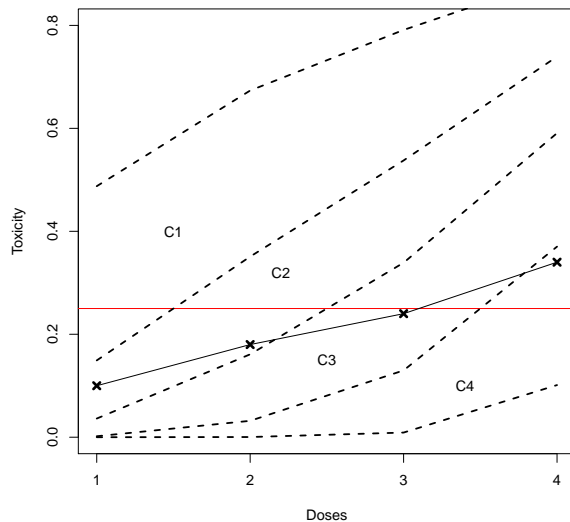


Figure 3.1: A case of poor specification according to F2: 4 doses ; $\alpha = 0.25$; $\beta = (0.1, 0.18, 0.24, 0.34)$ (solid line); $f(x|a) = x^{\exp(a)}$; $D' = (0.04, 0.17, 0.35, 0.60)$; $[A, B] = [-1.5, 1.5]$. The collection of curves C_d in the figure corresponds to the set of functions $f(\cdot|a)$ when a belongs to the interval H_d . When a_n belongs to H_d , the curve of $f(\cdot|a_n)$ is included in the particular collection C_d and, at step $n + 1$, the CRM will recommend the dose d . The model based estimate of toxicities at doses 2 and 4 hamper final recommendation of the MTD

expressing itself via the parametric elements a_d . 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 $(H_d)_{d \in D}$. In this way, in a Bayesian setting, it would be possible to replace Step 2 by a step that expresses more directly the goal:

Step 2'. At step $n + 1$, the selected dose X_{n+1} is the most probable according to the posteriori G_n : $X_{n+1} = \arg \max_{d \in D} G_n(H_d)$.

Now, the distribution G is equal to the tensor product $\Lambda \otimes \Pi$, where Π is a measure on the range D such that: $\Pi(\theta) = G(H_\theta)$, and Λ is a markov kernel that can be seen as the family indexed by θ of the restrictions of the distribution G to the sets H_d : $\Lambda(\cdot|\theta) = G(\cdot|H_\theta)$. The posterior Π_n , following the observations (X_1^n, Y_1^n) is defined by:

$$\Pi_n(\theta) = \frac{[\int L_n(a)\Lambda(da|\theta)] \Pi(\theta)}{\sum_{\theta=1}^m [\int L_n(a)\Lambda(da|\theta)] \Pi(\theta)}.$$

We then have: $\Pi_n(d) = G_n(H_d)$. This hierarchical structure has no particular value in this form but it allows a greater conceptual understanding of the distributions $\Lambda(\cdot|\theta)$ and their topological support. The resulting model structure can then, in this sense, be viewed as a generalization of the CRM. In this setting, the flexibility of the parameterization enables us to obtain methods that benefit from improved asymptotic properties while still conserving behavior that is very close to that of the CRM.

3.3 Semi-Parametric Models

3.3.1 General setting

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 3.3.2). The initial topological support of this prior is a broad one, corresponding to a non-informative situation. Subsequently, we model the accumulating information via an m -tuple of Bernoulli laws.

The purpose of SPM is to find a useful compromise between a rigid parametrization that may show poor performance when too far removed from the truth and a fully non parametric approach that lacks efficiency.

To this end we introduce $F = [0, 1]^m$ denoting the vector space of Bernoulli parameters needed to cover 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 connection with the classical CRM can be expressed as: $\forall j \in D, q_j = f_t(j, a)$ although t is not necessarily increasing. We proceed by partitioning the set F in terms of the main parameter of interest, θ . The result of this partition is to create distinct classes, each individual class containing an infinite set of members sharing the same MTD.

$$F = \bigcup_{\theta \in D} F_\theta, \text{ and } F_\theta = \{q \in F : \forall j \in D, |q_\theta - \alpha| \leq |q_j - \alpha|\}. \quad (3.2)$$

For all β in F_θ , the MTD is indexed by θ . The vector $(f(d, a))_{d \in D}$, with $a \in H_\theta$ (see the preceding section), is then included in F_θ . The likelihood is only evaluated on consistent sequences within the design (i.e. $X_i = \mathcal{M}(X_1^{i-1}, Y_1^{i-1})$ and $\mathbb{P}_q(X_i | X_1^{i-1}, Y_1^{i-1}) = 1$). By setting arbitrarily the first dose, this leads to:

$$\begin{aligned} \mathbb{P}_q(X_1^n, Y_1^n) &= \mathbb{P}_q(Y_1 | X_1) \times \prod_{i=2}^n [\mathbb{P}_q(Y_i | X_i) \mathbb{P}_q(X_i | X_1^{i-1}, Y_1^{i-1})] \\ &= \prod_{i=1}^n (q_{X_i})^{Y_i} (1 - q_{X_i})^{1 - Y_i} = \prod_{j=1}^m q_j^{n_j^1} (1 - q_j)^{n_j^0}, \end{aligned}$$

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 topological support of the measure $\Lambda(\cdot | \theta) = \Lambda_\theta(\cdot)$ is included in the class F_θ . When the discrete set D is endowed with its borel algebra, then the family of probability measures indexed by θ can be viewed as a kernel. This simple condition enables us to calculate the posterior distribution of θ given (X_1^n, Y_1^n) :

$$\Pi_n(\theta) = \Pi(\theta | X_1^n, Y_1^n) = \frac{\left[\int \prod_{j=1}^m q_j^{n_j^1} (1 - q_j)^{n_j^0} \Lambda_\theta(dq) \right] \Pi(\theta)}{\sum_{\theta=1}^m \left[\int \prod_{j=1}^m q_j^{n_j^1} (1 - q_j)^{n_j^0} \Lambda_\theta(dq) \right] \Pi(\theta)}. \quad (3.3)$$

Suppose the measures Λ_θ to have densities λ_θ with respect to a measure ν . The posterior distribution on F_θ given (X_1^n, Y_1^n) is then defined as:

$$\lambda_{\theta, n}(q) = \frac{\prod_{j=1}^m q_j^{n_j^1} (1 - q_j)^{n_j^0} \lambda_\theta(q)}{\int \prod_{j=1}^m q_j^{n_j^1} (1 - q_j)^{n_j^0} \lambda_\theta(q) \nu(dq)}. \quad (3.4)$$

By replacing (3.4) in (3.3), we have:

$$\Pi_n(\theta) = \frac{\left[\int q_{X_n}^{Y_n} (1 - q_{X_n})^{1 - Y_n} \lambda_{\theta, n-1}(q) \nu(dq) \right] \Pi_{n-1}(\theta)}{\sum_{\theta=1}^m \left[\int q_{X_n}^{Y_n} (1 - q_{X_n})^{1 - Y_n} \lambda_{\theta, n-1}(q) \nu(dq) \right] \Pi_{n-1}(\theta)}. \quad (3.5)$$

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 . Estimation of the MTD is carried out by updating Π .

Indeed, the dose for which Π_n attains its maximum value is an estimator of the MTD.

$$\hat{\theta}_n = \arg \max_{\theta \in D} \Pi_n(\theta). \quad (3.6)$$

An estimator of θ based on expectation makes little sense without a quantitative interpretation of the dose. Average estimators of parameters β_j are easily obtained by integrating q_j with respect to the posterior distribution $(\Lambda \otimes \Pi)_n = \Lambda_n \otimes \Pi_n$ given (X_1^n, Y_1^n) :

$$\tilde{\beta}_j = \mathbb{E}_{(\Lambda \otimes \Pi)_n} [q_j] = \sum_{\theta=1}^m \left[\int q_j \Lambda_{\theta,n}(dq_j) \right] \Pi_n(\theta). \quad (3.7)$$

The sequential decision for the next dose to be used on the basis of the estimator $\hat{\theta}_n$ corresponds to Step 2' of CRM, while the one based on the distance between $\tilde{\beta}_j$ and the threshold α corresponds to Step 2. The Beta distribution linked to the likelihood has a key role. Let g be the following function: $g : [0, 1] \times \mathbb{R}_+^2 \rightarrow [0, 1]$, with $g(x, a, b) = x^a(1-x)^b$. Let $B(a+1, b+1)$ be the Beta distribution with shape parameters $a+1$ and $b+1$. The associated probability density function is $g(\cdot, a, b) / \int_0^1 g(x, a, b) dx$. The posterior distribution Π can be written as follows:

$$\Pi_n(\theta) = \frac{I_{\theta,n} \times \Pi(\theta)}{\sum_{\theta=1}^m I_{\theta,n} \times \Pi(\theta)}, \quad \text{with } I_{\theta,n} = \int \prod_{j=1}^m g(q_j, n_j^1, n_j^0) \Lambda_{\theta}(dq). \quad (3.8)$$

As the set of possible value for θ is discrete, the comparison of the integrals $I_{\theta,n}$ provides the estimator $\hat{\theta}_n$. The different parametrizations of the distributions $\Lambda \otimes \Pi$ and the range of possible estimators make SPM a large class of sequential decision rules. Almost all of the designs currently in use, including the algorithmic designs, can be seen in this setting. This broad generalization allows us to investigate more involved questions, to more readily allow comparison between competing designs and to look for ways to improve on any given design.

3.3.2 A simple prior model specification

Here, we focus on the prior distributions inside the classes, $(\Lambda_{\theta})_{\theta \in D}$. A wide variety of distributions could be used for Λ . In the Bayesian discrete case, the choice of topological support determines consistency (Freedman, 1963). It is therefore a central step in the parametrisation of the SPM. 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 (D2). 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, we introduce a natural partition of the interval $[0, 1]$ into 3 sets: $I_{\epsilon} = [\alpha - \epsilon_1, \alpha + \epsilon_2]$, $A = [\alpha + \epsilon_2, 1]$ et $B = [0, \alpha - \epsilon_1]$. The following structure encompasses

the topological support for each class:

$$S_\theta \subset B^{\theta-1} \times I_\epsilon \times A^{m-\theta} \quad (3.9)$$

Thus, we have : $q \in S_\theta \Rightarrow q_j \leq q_\theta \leq q_{j'}, \forall j < \theta < j'$; and this local property provides the basis for inference on θ . For $j < \theta < j'$, we have: $q_\theta \geq q_\theta'' \geq q_\theta'$, $\forall (q, q', q'') \in S_j \times S_{j'} \times S_\theta$. This make it possible to distinguish between the three classes by using only the observation on the current dose θ . Note that the inequalities involve a problem of identifiability when one or more of the true parameters β_j are equal to $\alpha - \epsilon_1$ or $\alpha + \epsilon_2$. But, at least for now, we will ignore these cases. In the case of classic CRM, these 3 intervals vary with dose (see Figure 3.1) and the support is one-dimensional. Here, the topological support of distributions that we choose are in m dimension. The Bernoulli parameters at each dose are considered as independent from the point of view of a single class θ . The family $(\Lambda_\theta)_{\theta \in D}$ is called the prior model and is used to make inference in place of the usual parametric model. Thus, the monotonicity hypothesis M3(b) about $(f(\cdot, a))_{a \in [A, B]}$ becomes a property of stochastic ordering on the prior model.

H 1. Λ_θ is a product of unidimensionnal distributions:

$$\Lambda_\theta(dq) = \Lambda_\theta^1(dq_1) \times \dots \times \Lambda_\theta^m(dq_m).$$

When possible, we write $\Lambda_\theta(dq_j)$ in place of $\Lambda_\theta^j(dq_j)$.

H 2. Let r and t be two members of D such that: $r < t$. For each dose j_0 , the posterior $\Lambda_{r,n}^{j_0}(\cdot)$ is stochastically greater than $\Lambda_{t,n}^{j_0}(\cdot)$:

$$\Lambda_{r,n}^{j_0}([0, x]) \leq \Lambda_{t,n}^{j_0}([0, x]), \forall x \in [0, 1], \forall n \in \mathbb{N}.$$

When the topological supports are not modified during the trial, this last assumption involves the coherence of the design. The next example fulfills this hypothesis.

Theorem 3.3.1. *If the assumptions H1 et H2 hold and $(\Lambda_\theta)_{\theta \in D}$ is the single prior model used, then SPM is coherent.*

Proof. Suppose that: $Y_{n+1} = 1$. A similar argument holds when $Y_{n+1} = 0$. We have:

$$\Pi_{n+1}(\theta) \propto \left[\int q \Lambda_{\theta,n}(dq) \right] \Pi_n(\theta) \quad .$$

If $\hat{\theta}_n = r$ then: $\forall \theta \in D, \Pi_n(\theta) \leq \Pi_n(r)$. Let $t > r$. The posterior $\Lambda_{r,n}^r(\cdot)$ is then stochastically greater than $\Lambda_{t,n}^r(\cdot)$ and,

$$\int q_r \Lambda_{\theta,n}(dq) = \int \left[\int \mathbb{1}_{\{0 \leq x < q_r\}} \mu(dx) \right] \Lambda_{\theta,n}(dq_r) = \int \Lambda_{\theta,n}^r([x, 1]) \mu(dx).$$

The hypothesis H2 leads to:

$$\int \Lambda_{r,n}^r(\cdot|x, 1)\mu(dx) \geq \int \Lambda_{t,n}^r(\cdot|x, 1)\mu(dx),$$

hence,

$$\int q_r \Lambda_{r,n}(dq) \geq \int q_r \Lambda_{t,n}(dq), \text{ so that : } \Pi_{n+1}(r) \geq \Pi_{n+1}(t).$$

□

The following example shows use of a conjugate prior for the likelihood.

Example 3.3.1. *The prior model is defined by a triplet $((S_\theta)_{\theta \in D}, (q^\theta)_{\theta \in D}, c)$. The sets S_θ are the topological supports of the distributions Λ_θ :*

$$S_\theta = \{q \in [0, 1]^m : q_i \leq \alpha - \epsilon_1 \leq q_\theta \leq \alpha + \epsilon_2 \leq q_j, \forall i < \theta, \forall j > \theta\} = B^{\theta-1} \times I_\epsilon \times A^{m-\theta} \subset F_\theta.$$

The vectors $q^\theta \in [0, 1]^m$ are the modes of distributions Λ_θ . Thus q_j^θ takes the maximum value of the density function λ_θ^j : $q^\theta = \arg \max_{q \in F_\theta} \Lambda_\theta(q)$. For reasons of identifiability and coherence, these modes are non-decreasing ($q_1^\theta < \dots < q_m^\theta, \forall \theta \in D$) and do not cross each other ($q_j^{\theta_1} < q_j^{\theta_2}, \forall \theta_1 > \theta_2, \forall j \in D$). This replaces the assumption M3 for the CRM. The positive real value c is the dispersion parameter of the distributions. Uniform priors on the topological supports are updated. For all $\theta \in D$ and all $q \in F_\theta$, we have:

$$\lambda_\theta(dq) \propto \prod_{j=1}^{\theta-1} g(q_j, cq_j^\theta, c(1 - q_j^\theta)) \mathbb{1}_{\{q_j \in [0, \alpha]\}} \mu(dq_j) \times \prod_{j=\theta+1}^m g(q_j, cq_j^\theta, c(1 - q_j^\theta)) \mathbb{1}_{\{q_j \in [\alpha, 1]\}} \mu(dq_j) \times \nu_\theta(dq_\theta),$$

where ν is the counting measure and μ the Lebesgue measure. This is a product of incomplete beta priors on $[0, \alpha[$ and $]\alpha, 1]$, except in θ which is a Dirac measure. These distributions are conjugate for the likelihood function.

Porzio and Ragozini (2009) laid out the following criterion concerning stochastic order.

Theorem 3.3.2. *A distribution Λ_1 is stochastically greater than Λ_2 if the ratio λ_1/λ_2 is non-decreasing.*

The shape parameters of the Beta distributions used here are all greater than 1. Theorem 3.3.2 implies that: if $a \geq a'$ and $b \leq b'$, then the distribution $B(a, b)$ is stochastically greater than $B(a', b')$. This result helps us to check the assumption H2. We note that the use of a single value c for all distributions and marginals is not required. Figure 3.2 represents a sample with respect to Λ_3 when the prior model is chosen with the aim of emulating the experimental design of the CRM.

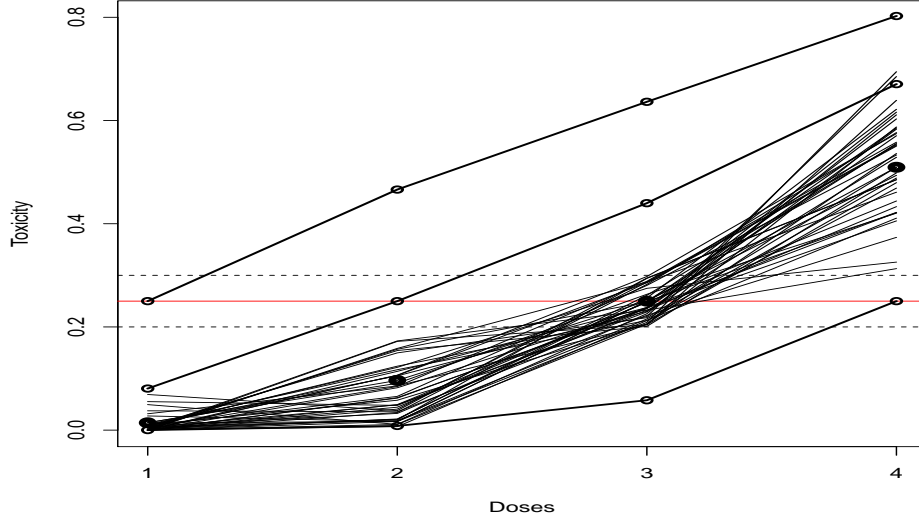


Figure 3.2: prior model: 4 doses, $\alpha = 0.25$, $\epsilon = 0.05$, $c = 15$. The symbols \circ locate the modes. These are based on the models of Figure 3.1: $q^d = (f(j, \alpha_d))_{j \in D}$, with $f(d, \alpha_d) = \alpha$. The figure shows a sample of size 40 from the class $\theta = 3$ following Λ_3^ϵ .

3.4 Large sample theory

Asymptotic behavior helps provide invaluable insights into general properties and the impact of different configurations on those properties. At the same time, in practice, sample sizes are typically small so that finite sample properties such as coherence are also important. Simulations complete the picture and these are presented in the next section. 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 3.1.1 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 to α within the distance ϵ of the target α , then we shall obtain the 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 assumptions on the prior model.

H $_\epsilon$. Let $\epsilon \geq 0$. The topological support S_θ of the distribution Λ_θ is:

$$S_\theta = \{q \in [0, 1]^m : q_i \leq \alpha - \epsilon \leq q_\theta \leq \alpha + \epsilon \leq q_j, \forall i < \theta, \forall j > \theta\} = B^{\theta-1} \times I_\epsilon \times A^{m-\theta} \subset F_\theta,$$

and for each dose j , S_θ^j , Λ_θ^j and λ_θ^j denote the marginal supports and their density functions.

Asymptotically, a parametrization satisfying this assumption corresponds to locating the dose

associated to some toxicity α with ϵ -accuracy. The second leans on the regularity of the prior model.

H 3. 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 \mathbb{R}_+^* , such that, for all $(j, \theta) \in D^2$, we have:

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

However in example 6.3.1, the assumption H3(b) is not verified. Indeed, the density function can not be bound into the neighborhoods of 0 or 1. We can resolve this problem by using uniform priors on the small intervals $[0, \delta[$ and $]\delta, 1]$:

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

The following two definitions describe the asymptotic behavior of SPM. ϵ -sensitivity is a property connected to indifference intervals. (Cheung and Chappell, 2002).

Definition 3.4.1. 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 reality is such that a unique dose is associated with a toxicity in the interval I_ϵ , then a method that is ϵ -sensitive converges almost surely to the MTD. When no dose has a toxicity located within I_ϵ , the SPM will assume an oscillating behavior between two doses with toxicities either side of the target α .

Definition 3.4.2. 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 3.1.1), 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 in the goal of determining almost surely the MTD.

Theorem 3.4.1. *Under the conditions H1, H3, H_ϵ , with $\epsilon \geq 0$ fixed, SPM is ϵ -sensitive and ϵ -balanced (see remark 1).*

Proof. The proof is given in the supplementary material. \square

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 of presenting the results.*

Large sample behavior of SPM is established by Theorem 3.4.1. 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 choice ϵ is made on the interval $[0, \max\{\alpha, 1 - \alpha\}]$. For values that are not too high the obtained performances are comparable. 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 Example 6.3.1, 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 :

$$\tilde{\theta}_n = \arg \max_{j \in \hat{D}_n} |\alpha - \hat{\beta}_{j,n}|, \quad (3.10)$$

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 3.4.1. *Under the conditions H1, H3, H_ϵ , 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 3.1.1) leads to an immediate conclusion. \square

Strong consistency of the estimator based on isotonic regression of the observations could be obtained in the same way. 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: $n_a \rightarrow \infty$ et $n_b \rightarrow \infty$. 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:

$$D_{KL}(p||q) = H(q|p) - H(p) = p \log\left(\frac{p}{q}\right) + (1-p) \log\left(\frac{1-p}{1-q}\right). \quad (3.11)$$

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 3.4.2. *Under the conditions H1, H3, H_ϵ , with $\epsilon = 0$, and when at least one of the toxicities β_a et 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). We can also obtain that n_a/n_b converges to 0 or $+\infty$ when b or a is the MTD respectively. We proceed by updating the support of the distribution Λ (see supplementary material). The theoretical advantage resulting from having variable topological support is not easily seen in practice for the usual applications of the CRM with small samples (< 50). We do not provide simulations for this. The circumstances in which the variability of the support could improve the behavior of the SPM with respect to the criteria (PCS) and (TR) remain to be investigated. The method arising from the variability of support and its theoretical results are presented in supplementary material. The purpose of the following section is to highlight the practical performance of the usual SPM through a comparison with the CRM.

3.5 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 the circumstances described below;

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 is $(S_\theta, (q^\theta)_{\theta \in D}, c = 48)$. S_θ verifies the hypothesis H_ϵ with: $\epsilon = 0.015$. The modes $(q^\theta)_{\theta \in D}$ are chosen close to the model of the CRM (3.1). They are given by table 3.1. The law Π can be used as an alternative way to prescribe initial dose escalation and

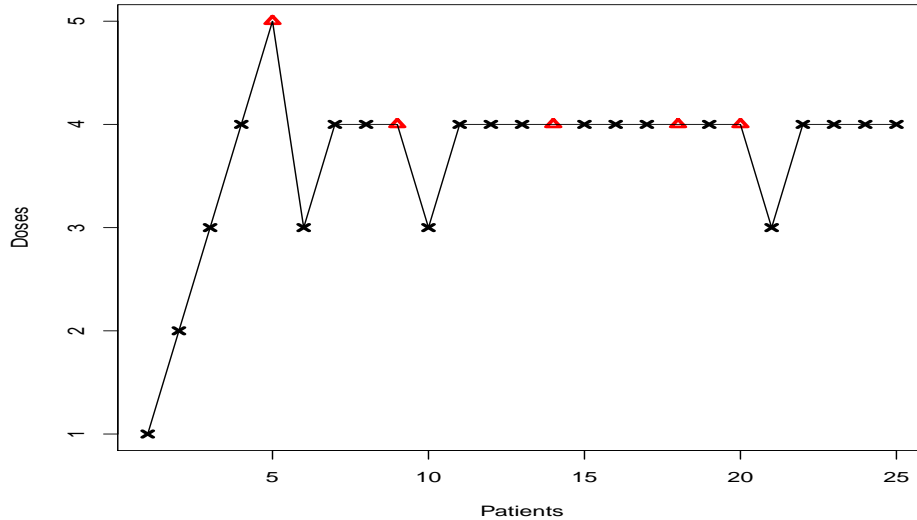
Table 3.1: The modes of the prior model

q_j^θ	q^1	q^2	q^3	q^4	q^5	q^6
1	0.20	0.12	0.02	0.01	0.00	0.00
2	0.29	0.20	0.07	0.05	0.00	0.00
3	0.42	0.36	0.20	0.08	0.02	0.00
4	0.57	0.48	0.35	0.20	0.09	0.01
5	0.69	0.62	0.50	0.34	0.20	0.04
6	0.82	0.78	0.70	0.58	0.44	0.20

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 laws 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 law \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 the intersection of a neighbourhood of it that is arbitrarily small with $\mathcal{B}(s)$. This intersection is not empty. It is then possible to find an approximation as accurate as we wish of the least informative distribution providing the sequence. Table 3.2 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 law corresponding to the sequence d is used. The CRM also produces this sequence while awaiting to observe the first toxicity.

Table 3.2: Prior II, walk through the levels awaiting the first observed toxicity.

Sequences without toxicity	1	2	3	4	5	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 3.3: An example of a sequence for SP-CRM (scénario 4); \triangle : toxicity, \times : non-toxicity.

Model and prior model

Table 3.3 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 hypothesis F2 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 this convergence difficulty manifesting itself in practice (see figure 3.4). 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

Table 3.3: 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

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 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 3.5). For this purpose the following algorithm, called *the pseudo-uniform scenario* is used.

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

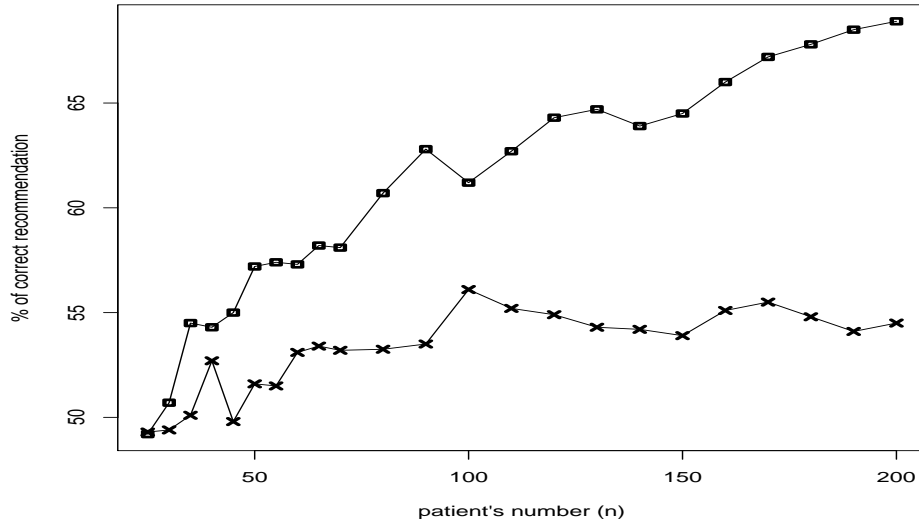


Figure 3.4: For scenario 6, (PCS) as a function of the number of included patients in the study. \square : SP-CRM ; \times : CRM.

value k .

2. 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 : $M \sim B(\max\{m - k; 0.5\}, 1)$.
3. The random scenario β has the law of an ordered sample of m uniform laws on $[0, B_s]$ conditioned by the event $\{MTD = k\}$.

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 3.4 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 3.4.2); (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

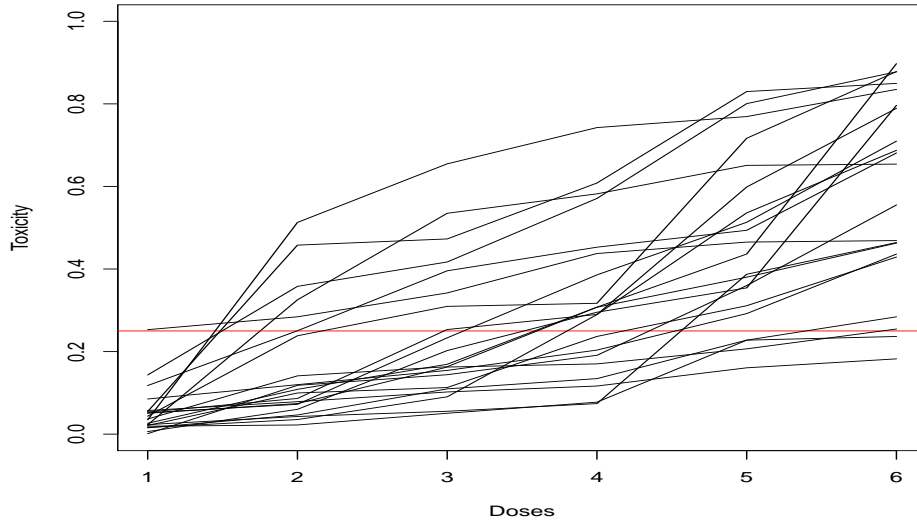


Figure 3.5: 20 randomly generated scenarios

criterion (Δ) : $\Delta(\text{Opt}) = 9.75$. The base reference is calculated for the CRM.

$$\text{R-}\Delta(\mathcal{M}) = \frac{\Delta(\mathcal{M}) - \Delta(\text{OPT})}{\Delta(\text{CRM}) - \Delta(\text{OPT})} \quad (3.12)$$

As $\text{R-}\Delta$ gets closer to 0, all the more the considered method gets close to the optimal design. In

Table 3.4: 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

all categories, the SP-CRM obtains the best results. Regarding the criteria (Δ), the difference between the SP-CRM and the optimal method is 40% smaller than what exists between this method and the CRM. This large difference can be explained in part by the fact that we are very close to the performance of the optimal method.

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. We have not studied how to possibly go about obtaining yet more advantageous parametrisations of SPM. 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 Figures 3.5 , the target is $\alpha = 0.2$ and the

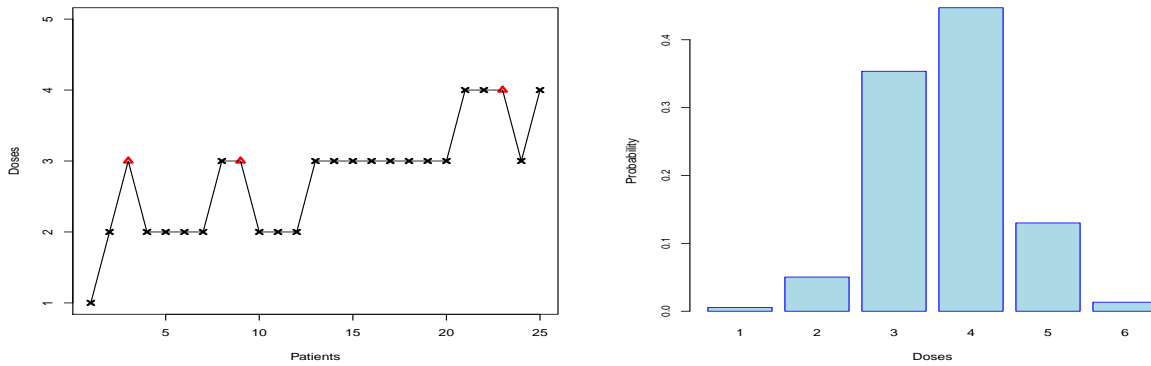


Figure 3.6: A trial: scenario = $(0.04, 0.08, 0.16, 0.24, 0.35, 0.45)$, \triangle =toxic, \times =non-toxic; and the final posterior Π .

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 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

The central feature of the SPM is the direct modelling of the key parameter of interest, the doses, structured around a regression function that is not fully specified. This allows us to incorporate the main working hypothesis, that of monotonicity as well as constraints guided by ethical principles. The approach is expressed via a hierarchical Bayesian structure. So far, we have not tried to increase the flexibility of the 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 itself. 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 it 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 non-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 some cases better, than all other available methods.

3.6 Appendix 1: Material for the proof of asymptotic results

3.6.1 Progressive elimination of unsuitable candidate doses

We focus on the asymptotic behaviour of the posterior odds of two classes: $I_{n,r}/I_{n,t}$ (see 3.8). For dose j , different situations describe the distance between the reality β_j and the marginal supports of the considered distributions. Under the assumption $H_{\epsilon,n}$, in the circumstances studied in the lemma 3.6.1, the order induced by the prior-distance of Kullback-Leibler is the same as the one induced by Euclidean distance. Finally, considering the hypothesis H1, we shall focus on each dose separately. We provide two competing values r and t of parameter θ . As usual, Λ_r and Λ_t denote the priors and S_r and S_t their topological supports. In order to unburden the notation we do not refer to dose j on which the observations are based. Thus, $\beta, q, n, n^1, n^0, \Lambda_r, \Lambda_t, S_r$ and S_t denote respectively $\beta_j, q_j, n_j, n_j^1, n_j^0, \Lambda_r^j, \Lambda_t^j, S_r^j$ and S_t^j . The ratio for which we would like to evaluate the limit is:

$$\frac{M_{n,r}}{M_{n,t}} = \frac{\int g(q, n^1, n^0) \Lambda_r(dq)}{\int g(q, n^1, n^0) \Lambda_t(dq)}. \quad (3.13)$$

An infinite sequence of observations will lead to two types of asymptotic behavior. We refer to these as "differentiated" and "non-differentiated".

Differentiated: A differentiated situation is one where the infinite sequence of observations about the dose tested allow us to exclude one of the classes. The ratio $M_{n,r}/M_{n,t}$ then converge almost surely to 0 or diverges almost surely, according to the class favored by the data. This appears when the marginal supports S_r and S_t are not equal and we distinguish the following generic cases:

D₁: β is closer to S_t than S_r : there exists $c > 0$ such that $d(\beta, S_t) + c < d(\beta, S_r)$, where the Euclidean distance, d , is extended to the distance between a point x and a set S : $d(\beta, S) = \inf\{d(x, y) : y \in S\}$. We exclude the special circumstances where the two supports are on the both side of the central interval, I_ϵ , and β lies in that interval. According to the demonstration of Theorem 3.4.1, this case does not need to be treated.

D₂: β is equal to α and S_t is the centered interval $I_\epsilon = \{\alpha\}$.

Non-Differentiated: A non-differentiated situation is one where the infinite sequence of observations about the dose tested is not decisive for the comparison of the two classes. In this case, it will be possible to bound the ratio $M_{n,r}/M_{n,t}$. This occurs when the marginal supports, S_r and S_t , are equal.

Lemma 3.6.1. *We assume that H_ϵ and H3 are satisfied.*

(i) *In the differentiated situation **D₁** and **D₂**, we have:*

$$\frac{M_{n,r}}{M_{n,t}} \xrightarrow{n \rightarrow +\infty} 0, \text{ a.s.}$$

(ii) In the non-differentiated situation, we have:

$$0 < \liminf_{n \rightarrow +\infty} \frac{M_{n,r}}{M_{n,t}} \leq \limsup_{n \rightarrow +\infty} \frac{M_{n,r}}{M_{n,t}} < +\infty, \text{ a.s.}$$

Proof. (i) **D₁**: Without loss of generality, we assume that $S_t = I_\epsilon = [\alpha \pm \epsilon]$ and $S_r = [\alpha + \epsilon, 1]$. First, we suppose that β lies in S_t . We have then: $\epsilon > c/2$. For n large enough, $\widehat{\beta}_n$ is in S_t (see Lemme 3.1.1). We introduce the interval of length $c/3$: $I_n = [\widehat{\beta}_n, \widehat{\beta}_n + c/3]$. For n large enough, I_n is included in S_t . For $q \in [0, 1]$, the functions $g(\cdot, q, 1 - q)$ are decreasing on $[q, 1]$. We have then, for $q \in I_n$:

$$g(q, \widehat{\beta}_n, 1 - \widehat{\beta}_n) > g(\widehat{\beta}_n + c/3, \widehat{\beta}_n, 1 - \widehat{\beta}_n),$$

and for $q \in S_{r,n}$:

$$g(q, \widehat{\beta}_n, 1 - \widehat{\beta}_n) < g(\widehat{\beta}_n + c/2, \widehat{\beta}_n, 1 - \widehat{\beta}_n).$$

As the functions $q \mapsto g(q + c/3, q, 1 - q) - g(q + c/2, q, 1 - q)$ and $q \mapsto g(q + c/2, q, 1 - q)$ are strictly positive and continuous on the compact set $[0, 1 - c]$, their minima are strictly positive:

$$\exists \delta > 0, \quad g(\widehat{\beta}_n + c/3, \widehat{\beta}_n, 1 - \widehat{\beta}_n) - g(\widehat{\beta}_n + c/2, \widehat{\beta}_n, 1 - \widehat{\beta}_n) > \delta,$$

and there exists $\delta' > 0$ such that: $g(\widehat{\beta}_n + c/2, \widehat{\beta}_n, 1 - \widehat{\beta}_n) > \delta'$.

The hypothesis H3 allows us to find a lower bound for the measure of interval I_n . Then, we have:

$$\frac{\int g(q, n^1, n^0) \Lambda_r(dq)}{\int g(q, n^1, n^0) \Lambda_t(dq)} \leq \frac{3}{cs} \left(\frac{g(\widehat{\beta}_n + c/2, \widehat{\beta}_n, 1 - \widehat{\beta}_n)}{g(\widehat{\beta}_n + c/3, \widehat{\beta}_n, 1 - \widehat{\beta}_n)} \right)^n \leq \frac{3}{cs} \left(\frac{\delta'}{\delta' + \delta} \right)^n,$$

which leads to the desired result. A similar argument can be given when β is less than $\alpha - \epsilon$. In this case, we set: $I_n = [\alpha - \epsilon, \alpha - \epsilon + c/3]$.

D₂: The question then reduces to comparing a Dirac law with one having a continuous support. By symmetry, we suppose that $S_{r,n}$ is equal to $[\alpha, 1]$. For $\epsilon_0 > 0$, we set:

$$\tilde{\sigma}_n = \frac{\sigma \sqrt{\log(\log(n))}}{\sqrt{n}} (1 + \epsilon_0), \quad \text{with } \sigma = \sqrt{\beta(1 - \beta)}.$$

Let I_n be the interval $[\alpha, \alpha + 3\tilde{\sigma}_n]$. The ratio of integrals can be written:

$$\frac{M_{n,r}}{M_{n,t}} = A_n + B_n = \frac{\int_{I_n} g(q, n^1, n^0) \Lambda_r(dq)}{g(\alpha, n^1, n^0)} + \frac{\int_{[\alpha, 1] \setminus I_n} g(q, n^1, n^0) \Lambda_r(dq)}{g(\alpha, n^1, n^0)}.$$

Our aim is to constrain the terms A_n et B_n by sequences that tend to zero. When $\widehat{\beta}_n$ is inferior to α , A_n is bounded by $3\tilde{\sigma}_n S$ and converges to 0. If this is not the case, then $\widehat{\beta}_n$ is included

in I_n for n large enough (lemma 3.1.1). We then have:

$$A_n \leq 3\tilde{\sigma}_n S \left(\frac{g(\hat{\beta}_n, \hat{\beta}_n, 1 - \hat{\beta}_n)}{g(\alpha, \hat{\beta}_n, 1 - \hat{\beta}_n)} \right)^n.$$

For n sufficiently large, a Taylor expansion to the second degree (as in Lemma 3.7.1 section 3.7) leads to the following inequality with $K > 0$:

$$\left(\frac{g(\hat{\beta}_n, \hat{\beta}_n, 1 - \hat{\beta}_n)}{g(\alpha, \hat{\beta}_n, 1 - \hat{\beta}_n)} \right)^n \leq \left(\frac{g(\hat{\beta}_n, \hat{\beta}_n, 1 - \hat{\beta}_n)}{g(\hat{\beta}_n - \tilde{\sigma}_n, \hat{\beta}_n, 1 - \hat{\beta}_n)} \right)^n \leq \exp [nK(\tilde{\sigma}_n^2 + o(\tilde{\sigma}_n^2))].$$

The bound of the sequence A_n converges to zero. For the other term, we have, for n large enough:

$$B_n \leq S(1 - \alpha) \left(\frac{g(\alpha + 3\tilde{\sigma}_n, \alpha + \tilde{\sigma}_n, 1 - (\alpha + \tilde{\sigma}_n))}{g(\alpha, \hat{\beta}_n, 1 - \hat{\beta}_n)} \right)^n, \text{ a.s.}$$

If $\alpha < 1/2$, then $g(\alpha, \alpha + \tilde{\sigma}_n, 1 - (\alpha + \tilde{\sigma}_n))$ is strictly less than $g(\alpha, \alpha - \tilde{\sigma}_n, 1 - (\alpha - \tilde{\sigma}_n))$. This gives:

$$B_n \leq S(1 - \alpha) \left(\frac{g(\alpha + 3\tilde{\sigma}_n, \alpha + \tilde{\sigma}_n, 1 - (\alpha + \tilde{\sigma}_n))}{g(\alpha, \alpha + \tilde{\sigma}_n, 1 - (\alpha + \tilde{\sigma}_n))} \right)^n.$$

A Taylor expansion to the second degree gives, for some $K_2 > 0$:

$$\left(\frac{g(\alpha + 3\tilde{\sigma}_n, \alpha + \tilde{\sigma}_n, 1 - (\alpha + \tilde{\sigma}_n))}{g(\alpha, \alpha + \tilde{\sigma}_n, 1 - (\alpha + \tilde{\sigma}_n))} \right)^n = \exp [-nK_2(\tilde{\sigma}_n^2 + o(\tilde{\sigma}_n^2))].$$

When $\alpha < 1/2$, the sequence $(B_n)_{n \in \mathbb{N}}$ tends to 0. We proceed in the same way when $\alpha \geq 1/2$.

(ii) This case is easily resolved by using the hypothesis H3. Equality of the supports implies that $S_r = S_t \neq [\alpha \pm \epsilon]$. By symmetry, we consider that the current case is $[0, \alpha - \epsilon]$. We then have:

$$\frac{s(\alpha - \epsilon) \int_0^{\alpha - \epsilon} g(q, n^1, n^0) dq}{S(\alpha - \epsilon) \int_0^{\alpha - \epsilon} g(q, n^1, n^0) dq} \leq \frac{M_{n,r}}{M_{n,t}} \leq \frac{S}{s},$$

which ends the demonstration. □

The results of Lemma 3.6.1 describe asymptotic behavior in the case of an infinite sequence. This covers the class of situations that we describe in the following subsections.

3.6.2 Proof of theorem 3.4.1

Proof. The following arguments deal with the asymptotic behavior of the method and establish convergence to 0 or divergence to $+\infty$. The doses that are tested a finite number of times can

be ignored since, for each fixed N and each dose j , we have:

$$0 < \frac{\int g(q_j, N_j^1, N_j^0) \Lambda_1(d(q_j))}{\int g(q_j, N_j^1, N_j^0) \Lambda_2(d(q_j))} < +\infty \quad .$$

We proceed then as though those doses outside of \tilde{D} are never tested, which formally amounts to replacing D by \tilde{D} in the integrals.

In order to establish the property of ϵ -sensitivity, we need consider the case where $\mathcal{E}(I_\epsilon, \beta)$ is not empty. The increasing dose-toxicity function implies that $\mathcal{E}(I_\epsilon, \beta)$ is a set of successive doses and possibly only a single dose: $\mathcal{E}(I_\epsilon, \beta) = \llbracket t^-, t^+ \rrbracket$, with t^- et t^+ the minimum and the maximum of this set.

We obtain a contradiction by assuming that \tilde{D} is not included in $\mathcal{E}(I_\epsilon, \beta)$. Let r be an element of $\tilde{D} \setminus \mathcal{E}(I_\epsilon, \beta)$. By symmetry we can choose r strictly less than t^- . We then compare the integrals $I_{n,r}$ and I_{n,t^-} .

$$\frac{I_{n,r}}{I_{n,t^-}} = \frac{\int \prod_{j \in \tilde{D}} g(q_j, n_j^1, n_j^0) \Lambda_r(dq)}{\int \prod_{j \in \tilde{D}} g(q_j, n_j^1, n_j^0) \Lambda_{t^-}(dq)} = \prod_{j \in \tilde{D}} \frac{\int g(q_j, n_j^1, n_j^0) \Lambda_r(dq_j)}{\int g(q_j, n_j^1, n_j^0) \Lambda_{t^-}(dq_j)} .$$

Using, Lemma 3.6.1, we study the location of the toxicities with respect to their marginal support. We obtain results that are valid for the marginal integrals of q_j , with $j \in \tilde{D}$.

When j belongs to $(\llbracket 1, r-1 \rrbracket \cup \llbracket t^-+1, m \rrbracket) \cap \tilde{D}$ then the situation is non-differentiated. When j belongs to $\llbracket r, t^- \rrbracket \cap \tilde{D}$ the situation is differentiated and corresponds to the cases \mathbf{D}_1 or \mathbf{D}_3 if $\beta_{t^-} = \alpha$. Furthermore this last set is not empty because r belongs to \tilde{D} . All of the marginals are bounded and, at least one of them corresponds to a case where

$$\frac{I_{n,r}}{I_{n,t^-}} \xrightarrow{n \rightarrow +\infty} 0, \text{ a.s.} \quad .$$

This contradicts our hypothesis since $I_{n,r}$ must be greater than $I_{n,j}$, for all $j \in D$, infinitely often. ϵ -sensitivity is thereby established. In order to show that SPM is ϵ -balanced, we need consider the case where $\mathcal{E}(I_\epsilon, \beta)$ is empty. We use the notation Definition 3.4.2. We obtain a contradiction by assuming that $\tilde{D} \setminus \{a, b\}$ is not empty. The same argument (Lemma 3.6.1), we have:

$$\forall j < a \Rightarrow \frac{I_{n,j}}{I_{n,a}} \xrightarrow{n \rightarrow +\infty} 0, \text{ a.s.} \text{ and } \forall j > b \Rightarrow \frac{I_{n,j}}{I_{n,b}} \xrightarrow{n \rightarrow +\infty} 0, \text{ a.s.},$$

from which we deduce: $\mathbb{P}_\beta [\exists N, \forall n > N : \mathcal{M}(X_1^n, Y_1^n) \in \{a, b\}] = 1$.

It remains to show that $\tilde{D} = \{a, b\}$ almost surely. Supposing that this were not true. By symmetry of the problem, we can restrict ourselves to $\tilde{D} = \{b\}$. We consider the marginal

integrals with the help of Lemma 3.6.1 and we find :

$$\frac{I_{n,b}}{I_{n,a}} \xrightarrow{n \rightarrow +\infty} 0, \text{ a.s. } .$$

which is a contradiction. □

3.6.3 Proof of corollary 3.4.2

Proof. In the case where $\epsilon = 0$, the SPM is ϵ -balanced and we have: $\hat{D}_n \rightarrow \tilde{D} = \{a, b\}$. Suppose that: $B_n = [0, \alpha - \epsilon_n]$, $I_{\epsilon,n} = \{\alpha\}$ et $A_n = [\alpha + \epsilon_n, 1]$. Recall that the beta function satisfies:

$$B(x, y) = \int_0^1 t^{x-1} (1-t)^{y-1} = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)} .$$

Let W be the set of full measure on which we have: $\tilde{D} = \{b, a\}$ and for $j \in \tilde{D}$, $\hat{\beta}_{j,n} \rightarrow \beta_j$. Our study is restricted to this set, and asymptotic behavior with respect to the weights of the classes a and b :

$$\frac{I_{a,n}}{I_{b,n}} = \prod_{j \in \tilde{D}} \frac{\int g(q_j, n_j^1, n_j^0) \Lambda_a(dq_j)}{\int g(q_j, n_j^1, n_j^0) \Lambda_b(dq_j)} = L_n \times K_n ,$$

with:

$$K_n = \frac{g(\alpha, n_a^1, n_a^0)/B(n_a^1 + 1, n_a^0 + 1)}{g(\alpha, n_b^1, n_b^0)/B(n_b^1 + 1, n_b^0 + 1)} ,$$

and

$$L_n = \frac{\int_{B_n} g(q_b, n_b^1, n_b^0)/B(n_b^1 + 1, n_b^0 + 1) \Lambda_a(dq_a)}{\int_{A_n} g(q_a, n_a^1, n_a^0)/B(n_a^1 + 1, n_a^0 + 1) \Lambda_b(dq_b)} .$$

Now, for a Bernoulli law, the posterior distribution constructed from a uniform prior on $[0, 1]$ is consistent (Freedman, 1963, theorem 1). This means that, almost surely, it converges weakly to a Dirac law at the true value of the parameter. This result is based on the asymptotic properties of the sample by virtue of the law of large numbers and, which in our situation, are attained (Theorem 3.1.1). Note that, γ is the counting measure. We then have for every function f , continuous on $[0, 1]$:

$$\int f(q_j) \frac{g(q_j, n_j^1, n_j^0)}{B(n_j^1 + 1, n_j^0 + 1)} \mu(dq_j) \xrightarrow{n_j \rightarrow +\infty} \int f(q_j) \mathbb{1}_{\{\beta_j\}}(q_j) \gamma(dq_j) = f(\beta_j) , \text{ a.s.}$$

The hypothesis H3 allows us to suppose the following structure:

$$s/S \leq \liminf L_n \leq \limsup L_n \leq S/s .$$

When n increases without bound, Stirling's formula leads to the following equivalent structure:

$$\frac{n+1}{B(n^1+1, n^0+1)} = \binom{n}{n\hat{\beta}_n} \sim \frac{1}{\hat{\beta}_n^{n\hat{\beta}_n} (1-\hat{\beta}_n)^{n(1-\hat{\beta}_n)}} \frac{1}{\sqrt{2\pi n\hat{\beta}_n(1-\hat{\beta}_n)}}$$

from which, $\log(1/B(n_j^1, n_j^0)) \sim n_j H(\hat{\beta}_{j,n})$. Therefore, when n tends to positive infinity

$$K_n \sim \tilde{K}_n = \frac{\exp\left[-n_a D_{KL}(\hat{\beta}_{a,n}|\alpha)\right]}{\exp\left[-n_b D_{KL}(\hat{\beta}_{b,n}|\alpha)\right]}.$$

We argue by contradiction, in supposing that there exists a subsequence $(n')_{n \in \mathbb{N}}$ and real $c > 0$ such that:

$$n'_a > n'_b (c + D_{KL}(\beta_b|\alpha) / D_{KL}(\beta_a|\alpha)).$$

We then have:

$$\tilde{K}_{n'} < \frac{\exp\left[-n'_b D_{KL}(\hat{\beta}_{a,n'}|\alpha) \left(\frac{D_{KL}(\beta_b|\alpha)}{D_{KL}(\beta_a|\alpha)}\right)\right]}{\exp\left[-n'_b D_{KL}(\hat{\beta}_{b,n'}|\alpha)\right]} \exp\left[-c n'_b D_{KL}(\hat{\beta}_{a,n'}|\alpha)\right].$$

Now, on W , we have: $D_{KL}(\hat{\beta}_{b,n}|\alpha) \xrightarrow[n_b \rightarrow +\infty]{} D_{KL}(\beta_b|\alpha)$ et $D_{KL}(\hat{\beta}_{a,n}|\alpha) \xrightarrow[n_b \rightarrow +\infty]{} D_{KL}(\beta_a|\alpha)$. Therefore, there exists $\delta > 0$ such that, for n large enough, we have: $D_{KL}(\hat{\beta}_{a,n}|\alpha) > \delta$. The sequence $(K_{n'})_{n \in \mathbb{N}}$ and the ratio $I_{a,n'}/I_{b,n'}$ converge to 0. If we accept that the sequence $(I_{a,n}/I_{b,n})_{n \in \mathbb{N}}$ can not have an accumulation point in 0, this leads to the final result:

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

It remains then to show that $(I_{a,n}/I_{b,n})_{n \in \mathbb{N}}$ can not have a sub-sequence converging to zero. Let n' be the infinite subsequence of n such that $\hat{\theta}_{n'} \neq \hat{\theta}_{(n-1)'}$. The SPM being ϵ -balanced, there exists $l > 0$ such that the sequence $(I_{a,n'}/I_{b,n'})_{n \in \mathbb{N}}$ falls an infinite number of times in each of the intervals $[0, l]$ and $[l, +\infty]$. The variations $(I_{a,n+1}/I_{b,n+1}) \times (I_{b,n}/I_{a,n})$ being bounded below by a strictly positive number V , we obtain the expected result for the subsequence $(I_{a,n'}/I_{b,n'})_{n \in \mathbb{N}}$.

Let $n_{a,k} = n_a - \lfloor n_a/k \rfloor$, where $\lfloor \cdot \rfloor$ is the function that produces the largest integer less than or equal to the argument. $n_{a,k}^1$ and $n_{a,k}^0$ are defined as usual: $n_{a,k}^1 = \sum_{i \in N(n,a,k)} \mathbb{1}_{\{Y_i=1\}}$ with $N(n,a,k) = \{i \in \mathbb{N}, X_i = a, Y_i = 1, n_{a,k} \leq i \leq n_a\}$ and $n_{a,k} = n_{a,k}^1 + n_{a,k}^0$. The law of large number is valid for the frequentist estimator $n_{a,k}^1/n_{a,k}$. Thus, as in Lemma 3.6.1 (i)D1, we

can shown that, for n large enough:

$$\frac{\int g(q_a, n_{a,k}^1, n_{a,k}^0) \Lambda_a(dq_a)}{\int g(q_a, n_{a,k}^1, n_{a,k}^0) \Lambda_b(dq_a)} < 1/V \text{ a.s.} \quad (3.14)$$

We set $\varphi(n) = \max_{n' \in E_n} (n' - (n-1)')$, where $E_n = \{n' < n, \hat{\theta}_{(n-1)'} = a\}$. We suppose that there exist a strictly positive number k such that for n large enough: $\varphi(n) - 1 > n/k$. Then, for all N , there exists $n' > N$ such that $\hat{\theta}_{(n-1)'} = a$ and $n' - (n-1)' > n/k$. Using the inequality (5.18), there exists N large enough, such that if $n' > N$, $\hat{\theta}_{(n-1)'} = a$ and $n' - (n-1)' > n/k$ then we can find $n \in \llbracket (n-1)'; n' \rrbracket$ with $\hat{\theta}_n = b$, which is a contradiction. Thus, we have: $n' - (n-1)' \in o(n)$. This leads to the final result:

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

□

3.7 Appendix 2: Consistency and oscillations in the case of variable support

3.7.1 Asymptotic results

The accumulation of information about a dose should ultimately lead to disqualify this dose if it does not correspond to our definition of the MTD. Thus, when θ_0 is not the MTD, there needs to exist a positive distance between β_{θ_0} and the topological support of the distribution Λ_{θ_0} . This condition cannot be verified in advance since it depends on the unknown reality β . However, our expectations may be adjusted to the observations during the trial by updating the ϵ in the hypothesis H_ϵ . The following hypothesis includes H_ϵ as a degenerate case.

H $_{\epsilon, \mathbf{n}}$. Let $(\epsilon_n)_{n \in \mathbb{N}}$ be a positive process such that: $\epsilon_n \in \sigma(X_1^n, Y_1^n)$. The topological supports $S_{\theta, n}$ of the distributions $\Lambda_{\theta}(\cdot | \epsilon_n)$ satisfy:

$$S_{\theta, n} = \{q \in [0, 1]^m : q_i \leq \alpha - \epsilon_n \leq q_\theta \leq \alpha + \epsilon_n \leq q_j, \forall i < \theta, \forall j > \theta\} = B^{\theta-1} \times I_{\epsilon_n} \times A^{m-\theta} \subset F_\theta .$$

$S_{\theta, n}^j$ denotes the marginal supports.

Here, the prior inside the classes are modified in the course of the trial. We describe the updating of the process $(\epsilon_n)_{n \in \mathbb{N}}$. Let

$$\hat{D}_n = \{\hat{\theta}_n, \hat{\theta}_{n'}\}, \text{ with } n' = \max\{j < n : \hat{\theta}_j \neq \hat{\theta}_n\}, \quad (3.15)$$

and,

$$j_n^* = \operatorname{arginf}_{\theta \in \hat{D}_n} |\alpha - \hat{\beta}_{n,\theta}|, \text{ and } \epsilon_n = \left(|\alpha - \hat{\beta}_{n,j_n^*}| - \delta_n \right) \mathbb{1}_{\{|\alpha - \hat{\beta}_{n,j_n^*}| - \delta_n > 0\}}, \quad (3.16)$$

where δ_n forms a determinist and non-negative sequence converging to 0. The following lemma allows us to obtain a necessary and sufficient criterion to avoid this case. In this lemma, we assume that the recommended dose does not vary, hence it does not appear in the index of $\hat{\beta}_n$.

Lemma 3.7.1. *Let $(\delta_n)_{n \in \mathbb{N}}$ be a sequence of real numbers converging to 0 and c denote a strictly positive constant. For any $\epsilon > 0$ and for n large enough, we have:*

(i) *When $\beta \in]0, 1[$, a strictly positive constant K_1 exists such that:*

$$\frac{g(\hat{\beta}_n + (1+c)\delta_n, n\hat{\beta}_n, n(1-\hat{\beta}_n))}{g(\hat{\beta}_n + \delta_n, n\hat{\beta}_n, n(1-\hat{\beta}_n))} \leq \exp[-n(K_1 - \epsilon)(\delta_n^2 + o(\delta_n^2))], \text{ a.s.},$$

(ii) *When β is equal to 0 or 1, a strictly positive constant K_2 exists such that :*

$$\frac{g(\hat{\beta}_n + (1+c)\delta_n, n\hat{\beta}_n, n(1-\hat{\beta}_n))}{g(\hat{\beta}_n + \delta_n, n\hat{\beta}_n, n(1-\hat{\beta}_n))} \leq \exp[-n(K_2 - \epsilon)(\delta_n + o(\delta_n))], \text{ a.s.}.$$

Proof. (i) The sequence that we would like to study is:

$$\begin{aligned} A_n &= \frac{g(\hat{\beta}_n + (1+c)\delta_n, n\hat{\beta}_n, n(1-\hat{\beta}_n))}{g(\hat{\beta}_n + \delta_n, n\hat{\beta}_n, n(1-\hat{\beta}_n))} \\ &= \left(1 + \frac{c\delta_n}{\hat{\beta}_n + \delta_n} \right)^{n\hat{\beta}_n} \left(1 - \frac{c\delta_n}{1 - \hat{\beta}_n - \delta_n} \right)^{n(1-\hat{\beta}_n)}. \end{aligned}$$

We set: $a_n = \hat{\beta}_n + \delta_n$ et $b_n = 1 - \hat{\beta}_n - \delta_n$. According to Lemma 3.1.1, the sequences a_n and b_n converge respectively to β and $1 - \beta$. Thus, for n large enough A_n exists and we have:

$$\begin{aligned} A_n &= \exp \left(n \left[(a_n - \delta_n) \log \left(1 + \frac{c\delta_n}{a_n} \right) + (b_n + \delta_n) \log \left(1 - \frac{c\delta_n}{b_n} \right) \right] \right) \\ &= \exp \left(-n \left[c(2-c) \frac{b_n + a_n}{2a_n b_n} \delta_n^2 + o(\delta_n^2) \right] \right). \end{aligned}$$

The consistency of sequence a_n and b_n leads to an upper bound that we need to establish, with:

$$K_1 = c(2-c) \frac{1}{2\beta(1-\beta)}.$$

The proof of (ii) can also be made by using a Taylor's expansion. \square

Thus, if the dose prescribes j such that $\beta_j < \alpha$, then, for any $\theta > j$, the marginal support $S_\theta^{\epsilon_n, j}$ is equal to $[0, \hat{\beta}_{n,j} + \delta_n]$ and $S_j^{\epsilon_n, j}$ is equal to $[\hat{\beta}_{n,j} + \delta_n, 2\alpha - (\hat{\beta}_{n,j} + \delta_n)]$. Under an

assumption of regularity of the prior (see H3), we have: $s\delta_n \leq \Lambda_\theta^j([\widehat{\beta}_{n,j}, \widehat{\beta}_{n,j} + \delta_n] | \epsilon_n) \leq S\delta_n$. In order to promote uncertainty for current dose j , the decrease of the interval's measure shall be compensated by a faster convergence of the likelihood ratio presented in the Lemma 3.7.1. This requirement leads to the following hypothesis:

H 4. *The sequence $(\epsilon_n)_{n \in \mathbb{N}}$ is specified with the sequences $(\widehat{\beta}_n)_{n \in \mathbb{N}}$ and $(\delta_n)_{n \in \mathbb{N}}$, as described in (3.15) and (3.16). Moreover, a positive real number C exists such that: $C < 1/2$ and $n^{-C} = o(\delta_n)$.*

The following definition describes the asymptotic behavior of SPM.

Definition 3.7.1. *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 balanced, if for all β , we have:*

$$\tilde{D} = \{a, b\}, \text{ a.s.}$$

This leads to the following theorem which illustrates the behavior of SPM in the setting of variable support.

Theorem 3.7.1. *In the setting of variable support defined by $H_{\epsilon,n}$ and under the hypotheses H1, H3 and H4, we have:*

(i) *The SPM is ϵ -balanced and the sequence $(\epsilon_n)_{n \in \mathbb{N}}$ converges to $|\alpha - \beta_{d^*}|$.*

(ii) *Almost surely, $\frac{n_a}{n_b} \xrightarrow{n \rightarrow +\infty} \begin{cases} 0 & , \text{ if } MTD = b \\ +\infty & , \text{ if } MTD = a. \end{cases}$*

(see remark 2).

Remark 2. *In the theorem and proof we assume that there exists no dose such that β_{j_0} is equal to α . If we allow for the occurrence of such a possibility, it is possible to deal with it in a simpler way: the SPM is then ϵ -sensitive. This assumption makes for a clearer presentation. The proof is given in the following sections.*

The running estimate oscillates between the MTD and its alternative. As with Corollary 3.4.1 it is possible to deduce the strong consistency of the estimator $\tilde{\theta}_n$. Furthermore, point (ii) guarantees that the choice of dose alternative to the MTD becomes ever rarer as the study progresses. This is the main property of a convergence in probability in the discrete case. It is also possible to fix a low bound δ , to the sequence $(\epsilon_n)_{n \in \mathbb{N}}$ in order to obtain the property of ϵ -sensitivity for the interval $[\alpha - \delta, \alpha + \delta]$. The properties of Theorem 3.7.1 are then preserved when no component of the vector β belongs to this interval.

3.7.2 Progressive elimination of unsuitable candidate doses in case of variable support

The definitions and lemma given in this Section are a generalisation of the ones of section 3.6.1 to the case of variable supports. We provide two competing values r and t of parameter

θ and we focus on the asymptotic behaviour of the posterior odds of these two classes. As usual, $\Lambda_r(\cdot|\epsilon_n)$ and $\Lambda_t(\cdot|\epsilon_n)$ denote the priors and $S_{r,n}$ and $S_{t,n}$ their topological supports. In order to unburden the notation we do not refer to dose j on which the observations are based. Thus, $q, n, n^1, n^0, \Lambda_r(\cdot|\epsilon_n), \Lambda_t(\cdot|\epsilon_n), S_{r,n}$ and $S_{t,n}$ denote respectively $q_j, n_j, n_j^1, n_j^0, \Lambda_r^j(\cdot|\epsilon_n), \Lambda_t^j(\cdot|\epsilon_n), S_{r,n}^j$ and $S_{t,n}^j$. The ratio for which we would like to evaluate the limit is:

$$\frac{M_{n,r}}{M_{n,t}} = \frac{\int g(q, n^1, n^0) \Lambda_r(dq|\epsilon_n)}{\int g(q, n^1, n^0) \Lambda_t(dq|\epsilon_n)}. \quad (3.17)$$

Two types of asymptotic behavior occur when the dose tested is learned through an infinite sequence of observations:

Differentiated: A differentiated situation is one where the infinite sequence of observations about the dose tested allows us to exclude one of the classes. The ratio $M_{n,r}/M_{n,t}$ then converges almost surely to 0 or $+\infty$ almost surely, according to the classe favored by the data. This occurs when the marginal supports $S_{r,n}$ and $S_{t,n}$ are not equal and we distinguish the following generic cases:

D₁: One of the two supports is the centered interval $I_{\epsilon_n} = [\alpha \pm \epsilon_n]$. For n large enough, the frequentist estimator $\widehat{\beta}_n$ is closer to $S_{t,n}$ than $S_{r,n}$ and a positive number c exists such that, for the Euclidean distance d , we have:

$$d(\widehat{\beta}_n, S_{t,n}) + c < d(\widehat{\beta}_n, S_{r,n}).$$

D₂: In this case, β is not equal to α and the sequence $(\epsilon_n)_{n \in \mathbb{N}}$ is not constant. For n large enough, the frequentist estimator $\widehat{\beta}_n$ is in $\liminf(S_t^{\epsilon_n})$ but not in $\liminf(S_r^{\epsilon_n})$. However, the Euclidean distance between $S_r^{\epsilon_n}$ and $\widehat{\beta}_n$ converges to 0, for a subsequence indexed by $(\varphi(n))_{n \in \mathbb{N}}$:

$$d(\widehat{\beta}_{\varphi(n)}, S_{r,\varphi(n)}) \xrightarrow{n \rightarrow +\infty} 0, \text{ a.s. .}$$

D₃: β is equal to α and $S_{t,n}$ is the centered interval I_{ϵ_n} .

Non-Differentiated: A non-differentiated situation is one where the infinite sequence of observations about the dose tested is not decisive for the comparison of the two classes. In this case, it will be possible to bound the ratio $M_{n,r}/M_{n,t}$. This occurs when the marginal supports, $S_{r,n}$ and $S_{t,n}$, are equal.

Lemma 3.7.2. *We assume that $H_{\epsilon,n}$, $H4$ and $H3$ are satisfied.*

(i) *In the differentiated situation **D₁**, **D₂** and **D₃**, we have:*

$$\frac{M_{n,r}}{M_{n,t}} \xrightarrow{n \rightarrow +\infty} 0, \text{ a.s.}$$

(ii) In the non-differentiated situation, we have:

$$0 < \liminf_{n \rightarrow +\infty} \frac{M_{n,r}}{M_{n,t}} \leq \limsup_{n \rightarrow +\infty} \frac{M_{n,r}}{M_{n,t}} < +\infty, \text{ a.s.}$$

Proof. (i) **D₁**: Without loss of generality, we assume that $S_{t,n} = [\alpha \pm \epsilon_n]$ and $S_{r,n} = [\alpha + \epsilon_n, 1]$. First, we suppose n large enough such that $\widehat{\beta}_n$ is in $S_{t,n}$. We have then: $\liminf \epsilon_n > c/2$. We introduce the interval of length $c/3$: $I_n = [\widehat{\beta}_n, \widehat{\beta}_n + c/3]$. For n large enough, I_n is included in $S_{t,n}$. For $q \in [0, 1]$, the functions $g(\cdot, q, 1 - q)$ are decreasing on $[q, 1]$. We have then, for $q \in I_n$:

$$g(q, \widehat{\beta}_n, 1 - \widehat{\beta}_n) > g(\widehat{\beta}_n + c/3, \widehat{\beta}_n, 1 - \widehat{\beta}_n),$$

and for $q \in S_{r,n}$:

$$g(q, \widehat{\beta}_n, 1 - \widehat{\beta}_n) < g(\widehat{\beta}_n + c/2, \widehat{\beta}_n, 1 - \widehat{\beta}_n).$$

As the functions $q \mapsto g(q + c/3, q, 1 - q) - g(q + c/2, q, 1 - q)$ and $q \mapsto g(q + c/2, q, 1 - q)$ are strictly positive and continuous on the compact set $[0, 1 - c]$, their minima are strictly positive:

$$\exists \delta > 0, \quad g(\widehat{\beta}_n + c/3, \widehat{\beta}_n, 1 - \widehat{\beta}_n) - g(\widehat{\beta}_n + c/2, \widehat{\beta}_n, 1 - \widehat{\beta}_n) > \delta,$$

and there exists $\delta' > 0$ such that: $g(\widehat{\beta}_n + c/2, \widehat{\beta}_n, 1 - \widehat{\beta}_n) > \delta'$.

The hypothesis H3 allows us to find a lower bound for the measure of interval I_n . Then, we have:

$$\frac{\int g(q, n^1, n^0) \Lambda_r^{\epsilon_n}(dq)}{\int g(q, n^1, n^0) \Lambda_t^{\epsilon_n}(dq)} \leq \frac{3}{cs} \left(\frac{g(\widehat{\beta}_n + c/2, \widehat{\beta}_n, 1 - \widehat{\beta}_n)}{g(\widehat{\beta}_n + c/3, \widehat{\beta}_n, 1 - \widehat{\beta}_n)} \right)^n \leq \frac{3}{cs} \left(\frac{\delta'}{\delta' + \delta} \right)^n,$$

which leads to the desired result. A similar argument can be given when $\widehat{\beta}_n$ is less than $\alpha - \epsilon_n$. In this case, we have: $I_n = [\alpha - \epsilon_n, \alpha - \epsilon_n + c/3]$.

D₂: It is sufficient to think in terms of the selected indices by the sub-sequence taken from φ , since the other case of sub-sequence can be reduced to either this case or the case **D₁**. In what follows, notation $\varphi(n)$ is replaced by n . By symmetry we consider: $\beta < \alpha$. According to the form of the prior laws given by the hypotheses $H_{\epsilon,n}$ and H4, we have $S_{r,n} = [\alpha \pm \epsilon_n]$, $S_{t,n} = [0, \alpha - \epsilon_n]$ and the sequence $(\alpha - \epsilon_n - \widehat{\beta}_n)_{n \in \mathbb{N}}$ converges to 0 and is strictly positive from step N and greater. For the rest of our argument we take n to be greater than N .

The function $g(\cdot, \widehat{\beta}_n, 1 - \widehat{\beta}_n)$ is decreasing on $[\widehat{\beta}_n, 1]$, from which we have:

$$\int g(q, n^1, n^0) \Lambda_1^{\epsilon_n}(dq) \leq g(\widehat{\beta}_n + \delta_n, n\widehat{\beta}_n, n(1 - \widehat{\beta}_n)),$$

and,

$$\frac{\int g(q, n^1, n^0) \Lambda_r^{\epsilon_n}(dq)}{\int g(q, n^1, n^0) \Lambda_t^{\epsilon_n}(dq)} \leq \left(\Lambda_t^{\epsilon_n}([\widehat{\beta}_n, \widehat{\beta}_n + \delta_n/2]) \right)^{-1} \frac{g(\widehat{\beta}_n + \delta_n, n\widehat{\beta}_n, n(1 - \widehat{\beta}_n))}{g(\widehat{\beta}_n + \delta_n/2, n\widehat{\beta}_n, n(1 - \widehat{\beta}_n))}.$$

The hypothesis H3 then gives:

$$\frac{M_{n,r}}{M_{n,t}} \leq (s\delta_n)^{-1} \frac{g(\widehat{\beta}_n + \delta_n, n\widehat{\beta}_n, n(1 - \widehat{\beta}_n))}{g(\widehat{\beta}_n + \delta_n/2, n\widehat{\beta}_n, n(1 - \widehat{\beta}_n))} ,$$

with $s > 0$.

Lemma 3.7.1 confirm that $M_{n,r}/M_{n,t}$ tends to 0.

D₃: The hypothesis H4 and the law of the iterated logarithm imply that ϵ_n is zero for n large enough. The proof works as in Lemma 3.6.1, D₃.

(ii) This case is easily resolved by using the hypothesis H3. Equality of the supports implies that $S_{r,n} = S_{t,n} \neq [\alpha \pm \epsilon_n]$. By symmetry, we consider that the current case is $[0, \alpha - \epsilon_n]$. We then have:

$$\frac{s(\alpha - \epsilon_n) \int_0^{\alpha - \epsilon_n} g(q, n^1, n^0) dq}{S(\alpha - \epsilon_n) \int_0^{\alpha - \epsilon_n} g(q, n^1, n^0) dq} \leq \frac{M_{n,r}}{M_{n,t}} \leq \frac{S}{s} ,$$

which ends the demonstration. □

3.7.3 Proof of theorem 3.7.1

Proof. (i) In a similar way to the proof of Theorem 3.4.1, we restrict our attention to the ratio of integrals on the set \tilde{D} :

$$\frac{I_{n,r}}{I_{n,t}} = \prod_{j \in \tilde{D}} \frac{\int g(q_j, n_j^1, n_j^0) \Lambda_r(dq_j)}{\int g(q_j, n_j^1, n_j^0) \Lambda_t(dq_j)} .$$

The following reasoning shows that \tilde{D} is included in $\{b, a\}$. Suppose that:

$$\limsup \epsilon_n = |\alpha - \beta_r|, \quad r \notin \{b, a\} .$$

Then, there exists r' belonging to \tilde{D} such that: $|\alpha - \beta_r| = |\alpha - \beta_{r'}|$. By symmetry, we choose $\beta_{r'} < \alpha$. We study the ratio $I_{n,r'}/I_{n,r'+1}$ according to the marginals indexed by j . When j belongs to $\llbracket 1, r' - 1 \rrbracket \cap \llbracket r' + 2, m \rrbracket \cap \tilde{D}$, the situation considered is non-differentiated. When j is equal to r' , the considered situation corresponds to the case **D₁** or **D₂**, according to the values of n . If $j = r' + 1$ belongs to \tilde{D} , we are in the case **D₁**. All of the marginals are bounded and there is at least one corresponding to a distinct case, from which

$$\frac{I_{n,r}}{I_{n,r+1}} \xrightarrow{n \rightarrow +\infty} 0, \quad a.s. \quad .$$

This contradicts the fact that r' belongs to \tilde{D} , for all r' such that $|\alpha - \beta_r| = |\alpha - \beta_{r'}|$, and therefore: $\tilde{D} \subset \{b, a\}$ almost surely. We then need to show that \tilde{D} is equal to $\{b, a\}$ almost

surely. By symmetry, we suppose that $\tilde{D} = \{b\}$. We make use of the marginal integrals via Lemma 3.7.2 and we find :

$$\frac{I_{n,b}}{I_{n,a}} \xrightarrow{n \rightarrow +\infty} 0, \text{ a.s. } .$$

We then have, almost surely: $\tilde{D} = \{b, a\}$, which implies: $\lim \epsilon_n = |\alpha - \beta_{d^*}|$, where d^* designates the MTD.

(ii) By symmetry suppose that: $d^* = b$. Let: $B_n = [0, \alpha - \epsilon_n]$, $I_{\epsilon,n} = \{\alpha\}$ and $A_n = [\alpha + \epsilon_n, 1]$. Recall that the beta function verifies:

$$B(x, y) = \int_0^1 t^{x-1} (1-t)^{y-1} = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)} .$$

Let W the set of full measure on which we have: $\tilde{D} = \{b, a\}$ et pour $j \in \tilde{D}$, $\hat{\beta}_{j,n} \rightarrow \beta_j$. We study on this set the asymptotic behavior of the ratio of the weights of classes a and b :

$$\frac{I_{a,n}}{I_{b,n}} = \prod_{j \in \tilde{D}} \frac{\int g(q_j, n_j^1, n_j^0) \Lambda_a(dq_j)}{\int g(q_j, n_j^1, n_j^0) \Lambda_b(dq_j)} = L_n \times K_n ,$$

with:

$$K_n = \frac{\int_{I_{\epsilon,n}} g(q_a, n_a^1, n_a^0) / B(n_a^1 + 1, n_a^0 + 1) \Lambda_a(dq_a)}{\int_{I_{\epsilon,n}} g(q_b, n_b^1, n_b^0) / B(n_b^1 + 1, n_b^0 + 1) \Lambda_b(dq_b)} ,$$

and

$$L_n = \frac{\int_{B_n} g(q_b, n_b^1, n_b^0) / B(n_b^1 + 1, n_b^0 + 1) \Lambda_a(dq_a)}{\int_{A_n} g(q_a, n_a^1, n_a^0) / B(n_a^1 + 1, n_a^0 + 1) \Lambda_b(dq_b)} .$$

As in the demonstration of Corollary 3.4.2, we have the following structure

$$s/S \leq \liminf L_n \leq \limsup L_n \leq S/s ,$$

and, on the basis of Stirling's formula: $\log(1/B(n_j^1 + 1, n_j^0 + 1)) \sim n_j H(\hat{\beta}_{j,n})$. We therefore have, when n tends to $+\infty$:

$$K_n \sim \tilde{K}_n = \frac{\int_{I_{\epsilon,n}} \exp \left[-n_a D_{KL}(\hat{\beta}_{a,n} || q_a) \right] \Lambda_a(dq_a)}{\int_{I_{\epsilon,n}} \exp \left[-n_b D_{KL}(\hat{\beta}_{b,n} || q_b) \right] \Lambda_b(dq_b)} .$$

We show the result as in Corollary 3.4.2. Proof follows by contradiction, in supposing that there exists a subsequence n' and $c > 0$ such that: $n'_a > c n'_b$. We then have:

$$\tilde{K}_{n'} < \frac{\int_{I_{\epsilon,n'}} \exp \left[-c n'_b D_{KL}(\hat{\beta}_{a,n'} || q_a) \right] \Lambda_a(dq_a)}{\int_{I_{\epsilon,n'}} \exp \left[-n'_b D_{KL}(\hat{\beta}_{b,n'} || q_b) \right] \Lambda_b(dq_b)} .$$

Now, on W , there exists $\delta > 0$ such that beyond some rank, we have:

$$D_{KL}(\widehat{\beta}_{a,n'} || \alpha + \epsilon_{n'}) > \delta \text{ and } D_{KL}(\widehat{\beta}_{b,n'} || \alpha - \epsilon_{n'}) \xrightarrow{n_b \rightarrow +\infty} 0.$$

Let $q_0 > \beta_b$ such that: $D_{KL}(\beta_b || q_0) = c\delta - \epsilon$, for $\epsilon \in]0, c\delta[$. Then, the sequence of functions $D_{KL}(\widehat{\beta}_{b,n'} || \cdot)$ converges uniformly to $D_{KL}(\beta_b || \cdot)$ on the compact space $[\beta_b, q_0]$. There therefore exists $l > 0$ and $\delta' \in]0, c\delta]$, such that beyond a certain rank, the interval $I_{n'} = [\alpha - \epsilon_{n'}, \alpha - \epsilon_{n'} + l]$ is included in $J_{2,n'} \cap [\beta_b, q_0]$ et:

$$\forall q \in I_{n'} , D_{KL}(\widehat{\beta}_{b,n'} || q) < \delta'.$$

We then have:

$$\tilde{K}_{n'} < \frac{\int_{J_{2,n'}} \exp[-cn'_b \delta] \Lambda_a(dq_a)}{s \times l \times \exp[-n'_b \delta']}.$$

Thus, almost surely, the sequence $(K_{n'})_{n' \in \mathbb{N}}$ converge to 0 and the ratio $I_{a,n'}/I_{b,n'}$ also which is impossible if we admit that $(I_{a,n}/I_{b,n})_{n \in \mathbb{N}}$ can not have a subsequence converging to 0.

We conclude as for Corollary 3.4.2 in order to obtain:

$$\frac{n_a}{n_b} \xrightarrow{n \rightarrow +\infty} 0, \text{ a.s. .}$$

□

Chapter 4

From CRM to Local Methods

4.1 Introduction

The semi-parametric method (SPM) may be viewed as some kind of a generalization of the continual reassessment method - certainly some restrictive conditions of the CRM are relaxed - and, from that standpoint, the immediate question to consider is whether, and in what way, this generalization leads to any improvement. As far as performance is concerned, in terms of accuracy of final recommendation and concentrating experimentation at the maximum tolerated dose (MTD), we showed that semi-parametric CRM (SP-CRM) was at least as good as the CRM and, in some cases, was able to obtain improvements. More importantly, the SPM is a method in its own right and can be seen to overcome some of the theoretical objections to the CRM such as the relatively narrow conditions under which almost sure convergence to the MTD can be obtained. However, not only is SPM very general, containing the CRM as a special case, it can also be readily calibrated to exactly mimic and reproduce the precise operating behaviour of many other available methods, in particular the CCD, the mTPI and the BOIN. All of these methods can now be viewed as special cases of the SPM. Furthermore, we can easily partially relax the restrictions on the SPM that results in CCD, mTPI and BOIN, and obtain immediate improvements on these designs. We call these SP-CCD, SP-mTPI and SP-BOIN, all of them coming under the heading of SPM. Our theoretical findings are backed up by simulations across a broad set of situations. In the following section we very briefly recall the semi-parametric method. In Section 3 we consider global models for the dose-toxicity regression function, under which heading we can place the CRM and SPM. In Section 4, we show that local methods, i.e., non-parametric methods that avoid modelling the whole range of doses, such as CCD, mTPI and BOIN, can be re-expressed in terms of SPM in conjunction with specific risk functions. These designs can then all be studied as special cases of the SPM and all of the earlier theoretical investigation of SPM can now be seen to also apply to CCD, mTPI and BOIN.

4.2 Semi-parametric dose finding methods

The aim is to estimate the root of an unknown dose-toxicity regression function in a situation where observations are accumulated sequentially. The predictor is the dose, $X_n \in D = \{0, \dots, 1\}$, and the variable to explain is the toxicity $Y_n \in \{0, 1\}$ (1 for a DLT and 0 otherwise). The conditional distribution of Y_n given $X_n = d$ is Bernoulli with parameter β_d which means that we assume that the toxicity probability at dose d does not depend of the way in which patients are included into the study.

Assumption 4.2.1. $\forall n \in \mathbb{N}, \quad \beta_d = \mathbb{P}(Y_n = 1 | X_n = d).$

Estimating the root of the regression function means determinating the dose among those available in the range D which suggests itself as having a probability of toxicity the closest to some target amount α chosen by the investigators. This dose is called the MTD (maximum tolerated dose) and we denote it d^* . The doses of range D have been chosen by clinical expertise so that they are ordered in terms of the probability of toxic response.

Assumption 4.2.2. $\forall (d, d') \in D^2, d < d' \Rightarrow \beta_d < \beta_{d'}.$

These first two assumptions are common in early phase dose finding studies. In chapter 3, we propose direct modelling of the MTD itself. The letter F denotes the vector space of Bernoulli parameters needed to cover a very wide range of situations: $F = [0, 1]^m$. 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 , creating m distinct classes, each class containing an infinite set of members sharing the same MTD.

$$F = \bigcup_{\theta \in D} F_\theta, \text{ and } F_\theta = \{q \in F : \forall j \in D, |q_\theta - \alpha| \leq |q_j - \alpha|\}. \quad (4.1)$$

For all β in F_θ , the MTD is indexed by θ . Let n_j^1 be the number of toxicities at dose j and $n_j = n_j^1 + n_j^0$, the number of patients treated at dose j . At step n , the likelihood function is $L_n(q) = \prod_{1 \leq j \leq m} q_j^{n_j^1} (1 - q_j)^{n_j^0}$. The set F is endowed with a probability measure $\Lambda \otimes \Pi$ such that Π is a measure on D and the topological support of the measure $\Lambda(\cdot | \theta) = \Lambda_\theta(\cdot)$ is included in the class F_θ . The posterior distribution of θ and $(q | \theta)$ given (X_1^n, Y_1^n) are:

$$\Pi_n(\theta) \propto \left[\int L_n(q) \Lambda_\theta(dq) \right] \times \Pi(\theta) \text{ and } \Lambda_{\theta, n}(A) \propto \int_A L_n(q) \Lambda_\theta(dq), \quad (4.2)$$

with A an element of the sigma-algebra associated with Λ_θ . 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 . The posterior of Π is obtained by weighting according to the expected value of the likelihood with respect to $(\Lambda_\theta)_{\theta \in D}$. The dose

for which Π_n attains its maximum value is the natural estimator of the MTD:

$$\hat{\theta}_n = \arg \max_{\theta \in D} \Pi_n(\theta). \quad (4.3)$$

For computational purpose, we assume that the Bernoulli parameters at each dose are independent from the point of view of a single class θ .

Assumption 4.2.3. Λ_θ is a product of unidimensionnal distributions:

$$\Lambda_\theta(dq) = \Lambda_\theta^1(dq_1) \times \dots \times \Lambda_\theta^m(dq_m),$$

when the context is not unambiguous, we write $\Lambda_\theta(dq_j)$ in place of $\Lambda_\theta^j(dq_j)$.

This assumption allow us to have a simple expression for the distribution of q_j , which provides an estimator of the true parameter β_j :

$$q_j \sim \Lambda^j \otimes \Pi \text{ and } \tilde{\beta}_{j,n} = \mathbb{E}_{(\Lambda \otimes \Pi)_n} [q_j] = \sum_{\theta=1}^m \left[\int q_j \Lambda_{\theta,n}^j(dq_j) \right] \Pi_n(\theta). \quad (4.4)$$

The choice of topological support is therefore a central step in the parametrisation of the SPM. Given that the main goal is that of estimation of the MTD, we use a natural partition of the interval $[0, 1]$ into 3 sets: $I_\epsilon = [\alpha - \epsilon, \alpha + \epsilon]$, $A = [\alpha + \epsilon, 1]$ and $B = [0, \alpha - \epsilon]$, with $\epsilon > 0$. The support S_θ of Λ_θ satisfies the following assumption:

Assumption 4.2.4. For all $\theta \in D$, we have: (i) $S_\theta \subset B^{\theta-1} \times I_\epsilon \times A^{m-\theta}$ or (ii) $S_\theta = B^{\theta-1} \times I_\epsilon \times A^{m-\theta}$.

When SPM works under assumption 4.2.3 and 4.2.4(ii), the main conjugate example of parametrization can be summarized by the couple $((q_j^\theta)_{\theta \in D}, (c_j^\theta)_{\theta \in D})$, where q_j^θ and c_j^θ are the mode and the dispersion of Λ_θ^j . Often, c_j^θ do not depend on θ and j . Let λ_θ be the density function of Λ_θ , we then have:

$$\lambda_\theta(q) \propto h(q_\theta, c_\theta^\theta, q_\theta^\theta) \mathbb{1}_{I_\epsilon}(q_\theta) \times \prod_{j=1}^{\theta-1} h(q_j, c_j^\theta, q_j^\theta) \mathbb{1}_B(q_j) \times \prod_{j=\theta+1}^m h(q_j, c_j^\theta, q_j^\theta) \mathbb{1}_A(q_j), \quad (4.5)$$

where $h(x, y, z) = x^{yz}(1-x)^{y(1-z)}$. This parametrization of the prior model corresponds to the construction of a global model very similar to the one used to describe the CRM. Such methods take into account the whole data on the range of doses in order to select the next dose.

4.3 Global models of the regression function: SPM and CRM

Many methods model the relation $Y \approx f(X)$ with the aim of using this relation to build an estimator of the MTD. The most well-known is the CRM. The purpose of this section is to

recall and illustrate quickly the link between CRM and SPM from the point of view of the respective topological support included in F_θ , which are the class of scenarios that have an MTD in θ . The CRM relies on a one dimensional parameter model which works by roughly approximating the dose-toxicity relationship $d \mapsto \beta_d$ by a family of continuous functions.

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

with $[A, B]$ the parameter set and \mathcal{X} a continuous space containing the range of doses D . Denote the prior distribution and density by G and g respectively. Let g_n be the posterior density given the observations (X_1^n, Y_1^n) :

$$g_n(a) = \frac{L_n(a) g(a)}{\int_{a \in [A, B]} L_n(a) g(a) da}.$$

The expected value of a with respect to g_n provides an estimate of the parameter. It is also possible to directly obtain estimators of the vector β : $\tilde{\beta}_{d,n} = \mathbb{E}_{G_n}[f(d, a)]$. With the goal of studying the consistency of this method, Cheung and Chappell (2002) introduced the intervals $H_1 = [A, b_1]$, $H_m = (b_{m-1}, B]$ and $H_d = (b_{d-1}, b_d)$ for $k = 1, \dots, m-1$, where the values b_d solves: $f(d|b_d) + f(d+1|b_d) = 2\alpha$. They show that the parameter belongs to H_d if and only if the model recommends dose d as the MTD:

$$H_d = \{a : |f(d|a) - \alpha| < |f(d'|a) - \alpha|, \forall d' \neq d\} \quad (4.6)$$

From the point of view of SPM, the distribution G can be seen as the tensor product $\Lambda \otimes \Pi$, where $\Pi(\theta) = G(H_\theta)$, and Λ is a family indexed by θ of the restrictions of the distribution G to the sets H_d : $\Lambda(\cdot|\theta) = G(\cdot|H_\theta)$. The posterior Π_n and $\Lambda_{\theta,n}$ given the observations (X_1^n, Y_1^n) are defined as in (4.2), which leads to the natural estimator $\hat{\theta}_n$ of the MTD (4.3). We then have: $\Pi_n(d) = G_n(H_d)$ and $\hat{\theta}_n = \arg \max_{d \in D} G_n(H_d)$. However, the sets H_d are determined indirectly. They are fixed by the choice of the model and a skeleton. Figure 4.1 (a) shows the topological support of Λ_θ for a power model ($f(x) = t(x)^a$) with $t(D) = (0.04, 0.17, 0.35, 0.60, 0.75)$ the value of the skeleton on the range of doses and the threshold $\alpha = 0.25$. One of the advantages of SPM is the ability to choose easily the prior model Λ_θ and its topological support. Figure 4.1 (b) shows the topological support of SPM satisfying the assumption 4.2.3 and 4.2.4(ii) with $\epsilon = 0.05$.

In that case, the topological supports are directly chosen as the whole prior model. This flexibility in the choice of Λ_θ and Π allows us to find parametrizations that satisfy theoretical properties which are difficult to obtain with CRM, such as asymptotic behavior, and Π calibration to constrain the method to follow a given sequence of initial dose escalation (chapter 3, simulations).

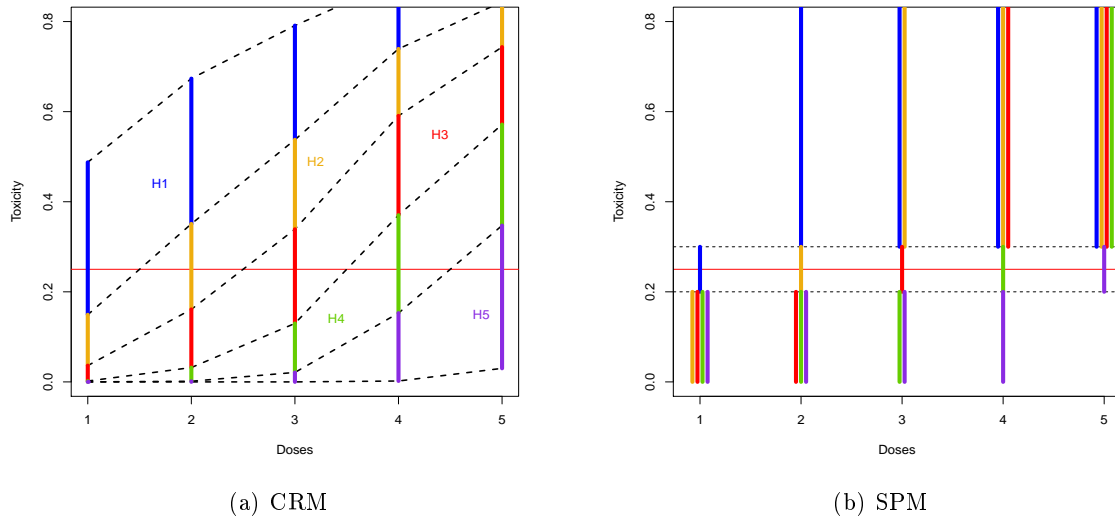


Figure 4.1: Topological support of Λ_θ . One colour for each $\theta \in D$.

4.4 From local methods to local models

We might view the standard 3+3 design as a local method but we do not consider it here and we limit our attention to designs that, even when they avoid the imposition of model structure, allow us to learn more as sample size increases. Several recently proposed methods can be viewed as being non-parametric in that no structure across the dose levels is assumed apart from that of monotonicity. Sequential decision making is based on a dynamically updated statistic related to the interval (B, I_e, A) (see 4.2.4 (ii)). These methods fall naturally under the heading SPM and expressing them within the context of SPM enables us to readily study their behaviour and principles upon which they are based.

Local decision making within SPM

A central principle upon which the non-parametric methods are based is to restrict attention to the current dose and to consider 3 possibilities: stay at the current level (\mathcal{S} , stay), escalate by one level (\mathcal{E} , escalate) or drop a level (\mathcal{D} , de-escalate). This local decision making makes sense in terms of being cautious. There exist several effective solutions to constrain methods coming under the SPM umbrella to behave in an identical way. The first amounts to making use of the risk functions $(0, 1, \infty)$, that we denote $(R_j)_{j \in D}$:

$$R_j(\theta, \hat{\theta}) = \begin{cases} 0 & , \text{ if } \theta = \hat{\theta} \\ 1 & , \text{ if } \theta \neq \hat{\theta} \text{ and } |\theta - j| \leq 1 \\ +\infty & , \text{ if } \theta \neq \hat{\theta} \text{ and } |\theta - j| > 1. \end{cases} \quad (4.7)$$

We use the function R_j when the last observation has been carried out at dose j (the current dose). The bayesian estimator corresponding to the statistic $\hat{\theta}_n$ (see 4.3) is then:

$$\hat{\theta}_n^l = \arg \min_{\theta \in D} \mathbb{E}_{\Pi_n} [R_{\hat{\theta}_{n-1}}(\theta, \hat{\theta}_n)] = \begin{cases} \hat{\theta}_{n-1} & , \text{ if } \hat{\theta}_n = \hat{\theta}_{n-1} \\ \hat{\theta}_{n-1} + 1 & , \text{ if } \hat{\theta}_n > \hat{\theta}_{n-1} \\ \hat{\theta}_{n-1} - 1 & , \text{ if } \hat{\theta}_n < \hat{\theta}_{n-1}. \end{cases} \quad (4.8)$$

The second solution is to modify the structure of the distribution Λ . We can suppress the ability to distinguish between classes r and t , situated the same side of θ , on the basis of observations obtained at the dose θ . If the marginal distributions Λ_θ are independent (see assumption 4.2.3), it is enough for the distributions to satisfy the following assumption:

Assumption 4.4.1. *For each pair (r, θ) and (t, θ') , we have:*

$$\text{sgn}(\theta - r) = \text{sgn}(\theta' - t) \Leftrightarrow \Lambda_\theta^r = \Lambda_{\theta'}^t,$$

where $\text{sgn}(\cdot)$ indicates the function sign and Λ_j^θ is the marginal at level θ of the law on the class j , denoted Λ_j .

This property means that, from the point of view of the prior model and for any given dose θ , any other dose can only be located: before, after or at that very same dose θ . For two doses r and t such that $r < t < \theta$ or $r > t > \theta$, if the family of laws verifies the assumptions assumptions 4.2.3 and 4.4.1, we have:

$$X_n = \theta \Rightarrow \frac{\Pi_n(r)}{\Pi_n(t)} = \frac{\Pi_{n-1}(r)}{\Pi_{n-1}(t)},$$

since

$$\begin{aligned} \frac{\Pi_n(r)}{\Pi_n(t)} &= \frac{\int g(q_\theta, n_\theta^1, n_\theta^0) \Lambda_r^\theta(dq_\theta)}{\int g(q_\theta, n_\theta^1, n_\theta^0) \Lambda_t^\theta(dq_\theta)} \times \prod_{j \in D \setminus \{\theta\}} \frac{\int g(q_j, n_j^1, n_j^0) \Lambda_r^j(dq_j)}{\int g(q_j, n_j^1, n_j^0) \Lambda_t^j(dq_j)} \\ &= \prod_{j \in D \setminus \{\theta\}} \frac{\int g(q_j, n_j^1, n_j^0) \Lambda_r^j(dq_j)}{\int g(q_j, n_j^1, n_j^0) \Lambda_t^j(dq_j)} = \frac{\Pi_{n-1}(r)}{\Pi_{n-1}(t)}. \end{aligned}$$

When r and s are located the same side of θ , testing at level θ provides no information on the dose r relative to the dose t . Another viewpoint would be to only make use of some subset of the information to decide on a dose. To facilitate this we make use of selection functions of the information on the sample:

$$\varphi_k(n_j^1, n_j^0, \theta) = \begin{cases} (n_j^1, n_j^0) & , \text{ if } |j - \theta| \leq k \\ (0, 0) & , \text{ otherwise.} \end{cases} \quad (4.9)$$

In this way, the function $\varphi_0(.,.,\theta)$ only makes use of those observations on the dose θ and the function $\varphi_1(.,.,\theta)$ makes use of information at θ and the adjacent doses. These different solutions can be used on a prior model such as that described for the SPM. They also give rise to a new perspective on the special cases described below.

Special Case 1: Modified Toxicity Probability Interval (mTPI)

The mTPI, introduced by Ji, Li, and Bekele (2010) amounts to making a sequence of local decisions based on the observations at the current dose: $(\mathcal{S}, \mathcal{E}, \mathcal{D})$. For dose j the Bernoulli parameter β_j is evaluated by a Bayesian approach. Let \mathcal{U}_I be the density of the uniform distribution on I , an interval includes in $[0, 1]$. For all $j \in D$, the prior \mathbb{P}_j for the probability of toxicity at dose j is a uniform law on the interval $[0, 1]$; the posterior distribution relative to the observations (n_j^1, n_j^0) satisfies:

$$\mathbb{P}_j(dq_j | n_j^1, n_j^0) \propto g(q_j, n_j^1, n_j^0) \mathbb{P}_j(dq_j),$$

where $g(x, y, z) = x^y \times (1 - x)^z$. For the current dose j and a given interval I , the statistic UPM (unit probability mass) is defined by calculating the posterior probability of the interval and dividing by its length $l(I)$:

$$\text{UPM}_j(I) = \frac{\mathbb{P}_j(I | n_j^1, n_j^0)}{l(I)} \propto \frac{\int_I g(q_j, n_j^1, n_j^0) \mathcal{U}_{[0,1]}(q_j) dq_j}{l(I)}.$$

The algorithm is based on the following 4 points that are used between subsequent cohorts:

1. The current dose is $j \in D$. We choose \mathcal{S} , \mathcal{E} or \mathcal{D} if the interval I_c , B or A (respectively) has the greatest UPM_j .
2. *Early termination*: The trial stops if $\mathbb{P}_1([\alpha, 1] | n_1^1, n_1^0) > \zeta$, with ζ close to 1 (for instance $\zeta = 0.95$).
3. *Dose Exclusion*: The set of doses to be explored are limited to $\{1, \dots, j - 1\}$ as soon as $\mathbb{P}_j([\alpha, 1] | n_j^1, n_j^0) > \zeta$, with ζ close to 1 (eg. $\zeta = 0.95$).
4. *Final recommendation*: Among the range of available doses, we select the one for which the isotonic estimator, \bar{p}_j , is the closest to α . If several doses satisfy the definition then we take the greatest when the isotonic estimator is smaller than α and the smallest otherwise.

The first point provides the real essence to the method. The intuitive interpretation on basing decisions on the maximization of the statistic UPM amounts to considering the question as

one of model choice:

$$\text{UPM}_j(I) \propto \frac{\int_I g(q_j, n_j^1, n_j^0) \mathcal{U}_{[0,1]}(q_j) dq_j}{l(I)} = \int g(q_j, n_j^1, n_j^0) \mathcal{U}_I(q_j) dq_j,$$

where $\mathcal{U}_I(\cdot)$ is the uniform distribution on the interval I . This first step then provides an answer to the question: which Bayesian model is the most probable for the parameter of the Bernoulli law at dose j : the uniform on B , the uniform on I_ϵ or on A ? Model selection is one of the interpretations of the methods coming under the SPM heading. It is therefore possible to provide a parametrization of SPM corresponding precisely to the main principal of mTPI. We proceed under the assumptions 4.2.3 and 4.4.1. The topological support of the laws Λ_θ are $B^{\theta-1} \times I_\epsilon \times A^{m-\theta}$ and the marginals are uniform on their supports. The law for Π is uniform on D . Bayesian updating is carried out following the results on each cohort of patients. The observations to be included in the estimation are selected by the function φ_0 . We have then, for the current dose j :

$$\begin{aligned} \Pi_n(\theta_0) &= \frac{\int \prod_{j'=1}^m g(q_{j'}, \varphi_0(n_{j'}^1, n_{j'}^0, j)) \Lambda_{\theta_0}(dq) \Pi(\theta_0)}{\sum_{\theta=1}^m \int \prod_{j'=1}^m g(q_{j'}, \varphi_0(n_{j'}^1, n_{j'}^0, j)) \Lambda_\theta(dq) \Pi(\theta)} = \frac{\int g(q_j, n_j^1, n_j^0) \Lambda_{\theta_0}(dq) \Pi(\theta_0)}{\sum_{\theta=1}^m \int g(q_j, n_j^1, n_j^0) \Lambda_\theta(dq) \Pi(\theta)} \\ &= \text{UPM}_j(B \mathbb{1}_{\{\theta_0 > j\}} + I_\epsilon \mathbb{1}_{\{\theta_0 = j\}} + A \mathbb{1}_{\{\theta_0 < j\}}). \end{aligned}$$

In this way, when we use $\hat{\theta}_n^l$ (see(4.8)) as the current estimator, the described parametrization is equivalent to the first point of the mTPI algorithm. This way of seeing the statistic UPM leads to an important question: what is the impact on behaviour resulting from making use of a function of restrictive selection. It may appear that we sacrifice efficiency by choosing the next dose to depend only on the current dose, j , and not, at least, on j and the adjacent doses $j-1$ and $j+1$. In the preceding formula this amounts to replacing the selection function φ_0 by φ_1 . One of the problems that we see in using φ_0 occurs when we escalate doses. The decision to reject the current dose does not take into account the information concerning the dose that we are moving to. Furthermore, investigation of this new dose level continues without regard for the information provided by our knowledge of monotonicity (assumption 4.2.2). As a result, we lose information and it seems likely that we will obtain poorer performance when working with φ_0 rather than φ_1 . That said, the mTPI can still obtain good results. This is due in part to the ad-hoc rules, corresponding to the algorithm's step (2), (3) and (4), which provide additional help to decision making based on the UPM. The two first are called *Conservative Rules*. Their purpose is to avoid selecting a dose that is too toxic given the observations. Note also that *Conservative Rules* and the principle *Final recommendation* can be adapted for use in any of the currently proposed dose finding methods.

Special Case 2: Local and Global BOIN designs.

The general construction of local and global BOIN, as well as CCD, is the same as that for mTPI and we only need to consider the main step of the method here. This corresponds to the first point of the algorithm. This step is based on making use of accumulated observations at the current dose. Use is made of a minimal sufficient statistic for the parameter of a Bernoulli law β_j in order to make a decision on the set $(\mathcal{S}, \mathcal{E}, \mathcal{D})$: $\hat{\beta}_j = n_j^1/n_j$. Locating $\hat{\beta}_j$ within the triplet of intervals $[0, \eta_L]$, $]\eta_L, \eta_U[$ and $[\eta_U, 1]$ corresponds respectively to the decisions \mathcal{E} , \mathcal{S} et \mathcal{D} . The lower and upper limits η_L and η_U are chosen in order to optimize different criteria that constitute the particularities of the method.

For BOIN (Bayesian Optimal INterval) introduced by Liu and Yuan (2015), the calibration of the intervals is carried out by minimizing the Bayesian decision error rate. This is defined for the current dose j by:

$$\tau(\eta_L, \eta_U) = \mathbb{P}[(\mathcal{E}^c) \cap \{\beta_j \in [0, \eta_L]\}] + \mathbb{P}[(\mathcal{S}^c) \cap \{\beta_j \in]\eta_L, \eta_U]\}] + \mathbb{P}[(\mathcal{D}^c) \cap \{\beta_j \in [\eta_U, 1]\}] \quad (4.10)$$

where the notation A^c indicates the complement of the event A . In order to provide an interpretation for τ , we need to give a probability distribution to β_j . Liu and Yuan suggest simulations for two versions of the prior. For local BOIN, the distribution of the parameter is a uniform mix of 3 Diracs at α_1 , α et α_2 ($\alpha_1 < \alpha < \alpha_2$): $1/3(\delta_{\alpha_1} + \delta_{\alpha} + \delta_{\alpha_2})$, whereas a mix of uniform laws on $[0, 1]$ is chosen for the global BOIN; $1/3(\mathcal{U}_B + \mathcal{U}_{I_e} + \mathcal{U}_A)$. The limits η_L et η_U obtained by minimization are fixed in the first case and vary as a function of the number of observations at the dose j in the second case.

An immediate corollary to Theorems 1 and 4 of Liu and Yuan (2015) is the inclusion of these methods within the framework of local model selection and, thereby, within that of SPM. In the case of local BOIN this means choosing the solution corresponding to maximum likelihood following observations at the current dose among the Dirac laws α_1 , α et α_2 . Global BOIN can be given a context within the framework of mTPI with the intervals A , I_e and B used in the prior being used to minimize the Bayesian error rate.

Special Case 3: Cumulative cohort design.

One of the corollaries of Theorem 1 (Liu and Yuan, 2015) is that any decision method that uses the placing of $\hat{\beta}_j$ within the triplet $([0, \eta_L],]\eta_L, \eta_U], [\eta_U, 1])$ reduces to a principle of maximum likelihood on a parameter space with 3 elements.

Proposition 4.4.1. *Let $(Y_i)_{1 \leq i \leq n}$, be a sample of i.i.d. Bernoulli variates with parameter α . Let $n^1 = \sum_{k=1}^n \mathbb{1}_{\{Y_k=1\}}$. For all values η_L, η_U and α such that: $0 < \eta_L < \alpha < \eta_U < 1$, there exists two numbers $\alpha_1 \in [0, \eta_L]$ and $\alpha_2 \in [\eta_U, 1]$ such that, whatever the value of $I \in$*

$\{[0, \eta_L[,]\eta_L, \eta_U[,]\eta_U, 1]\}$, we have:

$$\frac{n^1}{n} \in I \iff \arg \max_{a \in \{\alpha_1, \alpha, \alpha_2\}} \mathbb{P}_a [(n^1, n^0)] = \alpha_1 \mathbb{1}_I(\alpha_1) + \alpha \mathbb{1}_I(\alpha) + \alpha_2 \mathbb{1}_I(\alpha_2).$$

Proof. The maximum likelihood estimator on a set of parameters of Bernoulli laws is uniquely determined by the ratio $r = n^1/n$. Define the function f by:

$$f(a, r, n) = \frac{\mathbb{P}_a [(nr, n(1-r))]}{\mathbb{P}_\alpha [(nr, n(1-r))]} = \left[\left(\frac{a}{\alpha} \right)^r \left(\frac{1-a}{1-\alpha} \right)^{1-r} \right]^n.$$

We define α_1 and α_2 as the solutions in a of the following equations: $f(a, \eta_L, n) = 1$ for $a \in [0, \eta_L]$ and $f(a, \eta_U, n) = 1$ for $a \in [\eta_U, 1]$. The two numbers α_1 and α_2 verify the required properties since the function $f(\cdot, \eta_L, n)$ is increasing on $[0, \eta_L[$ and the function $f(\cdot, \eta_U, n)$ is decreasing on $[\eta_U, 1]$. \square

The CCD (Cumulative Cohort Design), introduced by Ivanova et al. (2007), fixes upper and lower limits η_L and η_U with the help of a markovian view of the dose finding problem together with a prior distribution on the possible scenarios. Proposition 4.4.1 implies that this method can be reduced to making use of a likelihood principle on 3 Dirac laws. For a target α at 0.25, the limits proposed in (Ivanova et al., 2007) are: $\eta_L = 0.16$ and $\eta_U = 0.34$; and this is exactly equivalent maximising the likelihood with respect to the laws δ_α , δ_{α_1} and δ_{α_2} , with $\alpha_1 \approx 0.0912$ and $\alpha_2 \approx 0.438$.

The CCD, and the BOIN, in common with mTPI are methods for local decision making. The root of the regression function that models (X, Y) is determined by a step by step search over the dose range. These methods do not appeal to a model describing the probability of toxicity as a function of the dose at which the patient is being treated and, for this reason, they can be considered to be non-parametric. Such a lack of structure has the advantage of simplicity but also some disadvantages. These disadvantages can of course be problematic and, in order to overcome some of the drawbacks, there is a need for some ad-hoc rules and further additions such as isotonic regression which is implemented at the end of the trial in order to determine the MTD. This lack of structure can also be mitigated by replacing the selection function φ_0 by φ_1 . The changes appear quite minor and yet they transform quite profoundly the nature of the methods which now can be expressed via a sliding model over 3 doses centered on the current doses. The new methods provided by these changes are called SP-mTPI, SP-CCD and SP-BOIN. Figure illustrates evolving topological support of the model. The following section which focuses on simulations allows us to summarize in a table the parameters of these several methods within the context of SPM.

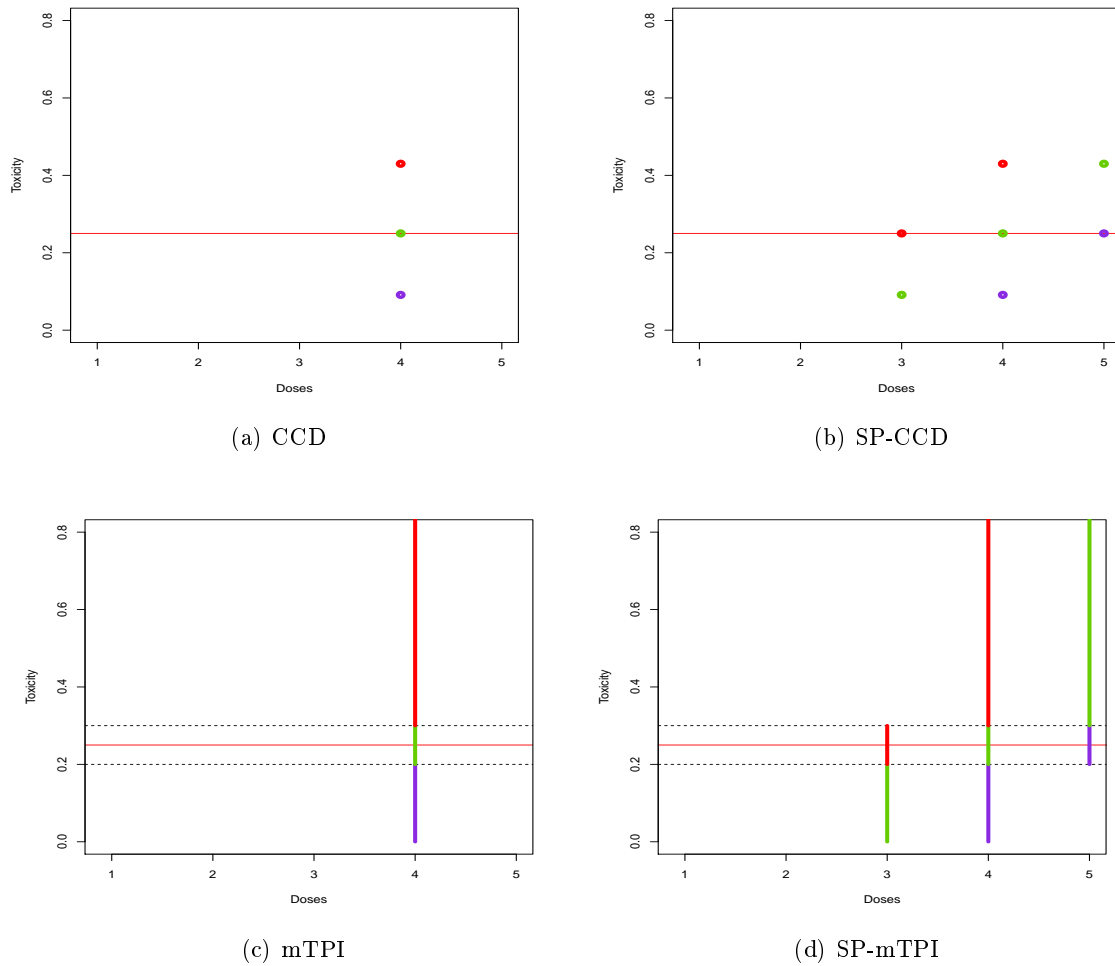


Figure 4.2: Topological support of Λ_θ . One colour for each $\theta \in D$. The current dose is the fourth and $\epsilon = 0.05$.

4.5 Simulations

In the following simulations, there are six doses and the threshold is fixed at 25%. There are 25 patients in the study except in Table 4.2 where, for the purposes of illustration, we show patient by patient inclusions upto the first 20 patients.

Definition 4.5.1. *A method works by memory-cohort of k patients if we estimate the next dose after each patient when a first cohort of k patients has been tested at the current dose.*

This offers a simple compromise between one-by-one inclusion and cohort inclusion. We would expect that it would behave similarly regarding performance as one-by-one inclusion and improve on that of simple cohort inclusion. We use it here as it helps better read the results of Table 2 and gives some insight into the behaviour of different methods. Five criteria are tested, one of them making use of the optimal method introduced by (O’Quigley et al., 2002):

- PCS: percentage of correct selection.

- TR: percentage of patients treated at the dose.
- TR2: percentage of patients treated at the two consecutive doses either side of the threshold.
- Δ : mean of the difference between the toxic rates at which patients are treated and the toxicity at the MTD.
- R- Δ : relative difference of Δ between the method and the optimal method. For a method \mathcal{M} , we have: $R-\Delta(\mathcal{M}) = (\Delta(\mathcal{M}) - \Delta(OPT))/(\Delta(\mathcal{M}^*) - \Delta(\mathcal{M}))$, \mathcal{M}^* is the better method according to the criterion Δ .

Table 4.1: SPM's parameters according to methods.

Parameters	Assumptions	Π	Λ_θ	φ	$\hat{\theta}$
CCD	4.2.3, 4.2.4 (i), 4.4.1	\mathcal{U}_D	$\alpha_1 = 0.912, \alpha_2 = 0.$	φ_0	$\hat{\theta}_n^l$
SP-CCD	4.2.3, 4.2.4 (i), 4.4.1	\mathcal{U}_D	$\alpha_1 = 0.15, \alpha_2 = 0.438$	φ_1	$\hat{\theta}_n^l$
mTPI	4.2.3, 4.2.4 (ii), 4.4.1	\mathcal{U}_D	$c_j^\theta = 0, q_j^\theta = 0, \epsilon = 0.05$	φ_0	$\hat{\theta}_n^l$
SP-mTPI	4.2.3, 4.2.4 (ii), 4.4.1	\mathcal{U}_D	$c_j^\theta = 0, q_j^\theta = 0, \epsilon = 0.05$	φ_1	$\hat{\theta}_n^l$
BOIN	4.2.3, 4.2.4 (i), 4.4.1	\mathcal{U}_D	$\alpha_1 = 0.15, \alpha_2 = 0.35$	φ_0	$\hat{\theta}_n^l$
SP-BOIN	4.2.3, 4.2.4 (i), 4.4.1	\mathcal{U}_D	$\alpha_1 = 0.15, \alpha_2 = 0.35$	φ_1	$\hat{\theta}_n^l$
SPM ₁	4.2.3, 4.2.4 (ii)	\mathcal{U}_D	$\epsilon = 0$ and see below	None	$\hat{\theta}_n$
SP-CRM	4.2.3, 4.2.4 (ii)	Π_0	$\epsilon = 0$ and see below	None	$\hat{\theta}_n$
CRM	as usual	$(\Pi_{H_j})_{j \in D}$	$\Lambda_\theta(a) = \Pi_{ H_\theta}(a)$	None	$\tilde{\theta}_n$

In the case of the CRM, we make use of a two stage version (O'Quigley and Shen, 1996) based on some some lead-in rule until we observe the first toxicity and then we use maximum likelihood. The chosen skeleton is $u = (0.07, 0.13, 0.25, 0.40, 0.55, 0.74)$ with Normal law, $\mathcal{N}(0, 1.34^2)$, for the prior. The CRM is based on the simple power model. The calibration of others methods are summarized in table 4.1. In that table, \mathcal{U}_D is the uniform law on the range of dose and Π_0 is proportional to $(1, 0.99, 0.95, 0.9, 0.84, 0.7)$. The SP-CRM has been calibrated on the basis of the skeleton of CRM. The method is anticipated to behave similarly to the CRM. The dispersion's parameter is always the same $c = 48$ and the modes satisfy: $q^\theta = (f(x|a_\theta))_{x \in D}$ with $f(\theta, a_\theta) = \alpha$. For SPM₁, the dispersion's parameter is $c = 40$ and we have: $\forall (j, \theta) \in D^2 : q_j^\theta = 0.13$ when $j < \theta - 1$; $q_j^\theta = 0.15$ when $j = \theta - 1$; $q_j^\theta = 0.25$ when $j = \theta$; $q_j^\theta = 0.35$ when $j = \theta + 1$; $q_j^\theta = 0.37$ when $j > \theta + 1$. Table 4.2 shows the effects of the *Conservative Rules*. For the observed toxicities, all the local methods (mTPI, BOIN, CCD) give the same sequence of dose. Scenario A lay out the necessity of the Conservative Rules in order to dampen the oscillation of local methods. For scenario B, we present sequences

where two toxicities are observed during the first cohort of patients treated at the MTD. Such an outcome for a trial is not the most likely, but is not rare. In such a circumstance the *Conservative Rules* remove the MTD from the possible doses as the SP-mTPI will have treated more than 50% of the patients enrolled at the MTD. The methods that model $Y \approx f(X)$ can work without the *Conservative Rules* and such additional restraints only have a very limited role to play.

Table 4.2: Example of trial under specific scenarios

Patients	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
A: Scenario= (0.05, 0.9, 1, 1, 1, 1)																				
Local methods without <i>Conservative Rules</i>																				
Doses	1	1	1	2	2	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2
Toxicity	0	0	0	1	1	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1
Local methods with <i>Conservative Rules</i>																				
Doses	1	1	1	2	2	2	1	2	1	1	1	1	1	1	1	1	1	1	1	1
Toxicity	0	0	0	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0
SP-mTPI (local model) without <i>Conservative Rules</i>																				
Doses	1	1	1	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	2
Toxicity	0	0	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1
B: Scenario= (0.1, 0.1, 0.25, 0.9, 1, 1)																				
Local methods without <i>Conservative Rules</i>																				
Doses	1	1	1	2	2	2	3	3	3	2	3	2	3	2	3	2	3	2	3	2
Toxicity	0	0	0	0	0	0	1	1	0	0	0	0	1	0	1	0	0	0	0	0
Local methods with <i>Conservative Rules</i>																				
Doses	1	1	1	2	2	2	3	3	3	2	3	2	3	2	2	2	2	2	2	2
Toxicity	0	0	0	0	0	0	1	0	1	0	0	0	1	0	0	1	0	0	0	0
SP-mTPI (local model) without <i>Conservative Rules</i>																				
Doses	1	1	1	2	2	2	3	3	3	2	2	3	3	3	3	3	3	3	3	3
Toxicity	0	0	0	0	0	0	1	0	1	0	0	0	1	0	0	1	0	0	0	0

In Table 4.3, the local methods and their semi-parametric alternative have been tested on 100 000 scenarios generated by the very general pseudo-uniform algorithm (chapter 3, simulations).

For all criterion, with or without *Conservative Rules*, the semi-parametric versions improve the performances of the original methods. Looking at the columns with or without *Conservative Rules* is interesting. The treatment allocation performance of mTPI, CCD and BOIN are worsened without *Conservative Rules* whereas they are improved for their semi-parametric alternatives. The analysis on scenario A and B of table 4.2 explains at least partially this result.

In Table 4.4, the methods have been tested on the same sample as the previous. These models cover the full range of doses. The CRM and its strong model perform better than the local methods, in particular for the ethical criterion. Changing from a strong model to the weaker semi-parametric complement, results in very slightly improved performance. These simulations

Table 4.3: Local methods and semi-parametric complements.

	With <i>Conservative Rules</i>		Without <i>Conservative Rules</i>	
	mTPI	SP-mTPI	mTPI	SP-mTPI
PCS	46.42	46.39	47.9	48.08
TR	32.46	32.74	32.13	33.22
TR2	50.76	50.83	50.8	51.33
Δ	11.82	11.76	11.87	11.64
R- Δ	1.16	1.11	1.21	1.0
	BOIN	SP-BOIN	BOIN	SP-BOIN
PCS	46.84	48.45	48.67	50.32
TR	31.39	33.11	30.84	33.56
TR2	49.78	51.75	49.83	52.30
Δ	12.11	11.81	12.26	11.72
R- Δ	1.43	1.15	1.56	1.07
	CCD	SP-CCD	CCD	SP-CCD
PCS	46.23	48.63	48.29	49.39
TR	32.18	32.71	31.96	33.18
TR2	50.47	50.94	50.71	51.34
Δ	11.88	11.76	11.97	11.64
R- Δ	1.22	1.11	1.30	1.0

Table 4.4: Model versus prior-model (without *Conservative Rules*).

Criterion	PCS	TR	TR2	Δ	R- Δ
CRM	47.43	33.14	52.02	11.64	1.0
SP-CRM	49.37	33.87	52.72	11.64	1.0
SPM ₁	50.63	34.02	52.65	11.72	1.07

shows the usefulness of local modelling of the regression function. Moreover, as presented in figure 4.5, the direct modelling of Π and its a posteriori distribution enable us to see for each dose its associated estimated probability of being the MTD. This is something that is dynamic and can be monitored as the trial progresses. From the point of view of further development, dose expansion cohorts or Phase I/II trials, we can use these features to select two or more doses with an agreed high probability of containing the true MTD. Table 4.5 shows what

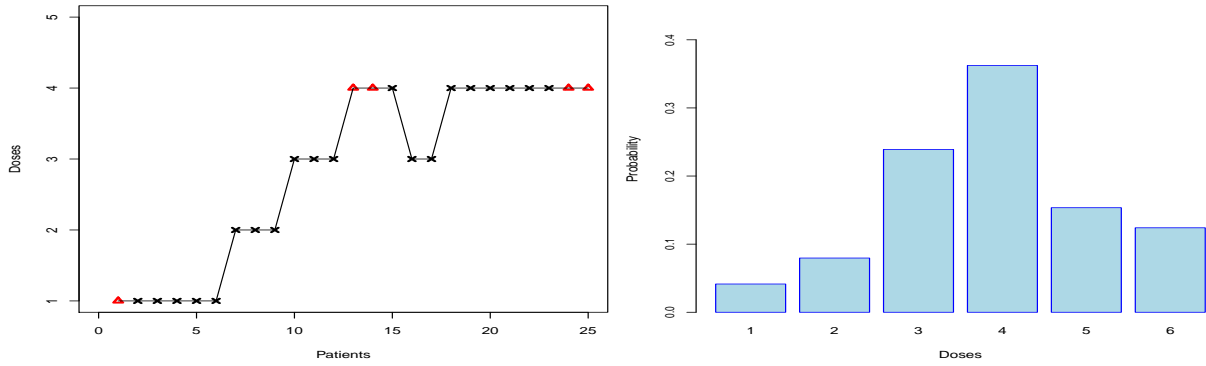


Figure 4.3: A trial for SPM_1 : scenario = $(0.05, 0.12, 0.21, 0.28, 0.45, 0.6)$, Δ =toxic, \times =non-toxic; and the final posterior Π .

such monitoring could look like based on a situation where the data do not allow a sharp distinction between levels 3 and 4. For all $r \in [0, 1]$, criterion $PCS(r)$ is the percentage of correct selection when we recommend the set $\{\theta \in D : \Pi(\theta) > r\}$ as containing the MTD and $SI(r)$ is the average size of this set. Thus, for SP-CRM, the selection of the dose associated with a probability higher than 0.25 of being the MTD is correct in 70.6% of cases and the average size of the recommended set is 1.70; for SPM_1 , the selection of dose associated with a probability higher than 0.1 of being the MTD is correct in 93% of cases and the average size of the recommended set is 3.30.

Table 4.5: Model versus prior-model.

Criterion	RE(0.1)	SI(0.1)	RE(0.15)	SI(0.15)	RE(0.2)	SI(0.2)	RE(0.25)	SI(0.25)
SP-CRM	83.7	2.58	79.4	2.23	74.1	1.93	70.6	1.70
SPM_1	93.3	3.30	84.5	2.43	74.6	1.85	65.7	1.48

Note that these results could be refined since they do not take into account the number of patients treated at each dose. In particular, figure 4.5 show a probability higher than 0.1 at dose 5 and 6 even if we do not observe any patients at these doses.

4.6 Conclusion and prospective work

The impossibility theorem of Azriel et al. (2011) shows that no method exists that would, for all situations, allow the current estimator to be strongly consistent for the MTD. Only particular configurations with respect to the employed method result in strong consistency. This argument has led us to consider new definitions concerning asymptotic behaviour. A method is said to be balanced when the current estimator converges to the whole set $M^* = \{a, b\}$ where a (above) and b below are the two consecutive doses on the both side of α (chapter

3, Large sample theory). Let $d(., .)$ be the euclidean distance. A sequence $(X_n)_{n \in \mathbb{N}}$ converges to a set B , denoted by $X_n \xrightarrow{S} B$, if :

$$\sup_{x \in B} \left(\liminf_{n \rightarrow +\infty} d(X_n, x) \right) = 0. \quad (4.11)$$

Definition 4.6.1. *Let D be a range of doses. A method, \mathcal{M} , is called balanced, if for all β we have,*

$$\mathcal{M}(X_1^n, Y_1^n) \xrightarrow{S} \{a, b\}, \text{ a.s.}$$

Among the methods tested previously, the SPM_1 and the SP-CRM are the only ones to have this property. This is important from a theoretical point of view and according to the usual assumption 4.2.2 of monotonicity, this property characterizes the most efficient manner to find with certainty the MTD. It is helpful to see these several different methods coming under a single umbrella. On the one hand we can see that overall performances between the methods are close. Choosing one method over another is to some degree a matter of taste since for most situations the actual results of a study will only depend very weakly on the chosen design, as long as it belongs to the semi-parametric class. This umbrella is also helpful when it comes to tackling more involved questions using any of these designs. The theoretical results apply across the board and can be made use of regardless of design choice. The structure in place can also be made use of when, in any given particular situation, we wish to make a decision about which of two or more competing designs we would like to use.

Chapter 5

The Role of Minimal Sets in Early Phase Dose Finding Studies

5.1 Introduction

Drug development, in particular in chronic diseases such as cancer, typically proceeds in 3 phases, the so-called Phase I, Phase 2 and Phase 3 clinical trials. Bayesian methods have had a very great impact on all three of these phases but, arguably, nowhere greater than in the earliest phase, the initial dose-finding study. With the relatively recent development of including dose expansion cohorts focused on efficacy, between the Phase I and the Phase 2 studies, the distinction among these phases is becoming blurred. Classically, the Phase I study has as its objective the locating of the maximum tolerated dose (MTD), the highest dose with a tolerable rate of toxicity that can be taken forward for further studies focused on efficacy.

In Phase 2 and Phase 3 studies the endpoint of interest, in many cases more than a single endpoint, can always be very sharply defined. As a result the task of crafting appropriate statistical designs to meet the requirements of efficient inference relating to such endpoints is at least in principle relatively straightforward. Traditionally, this has not been the case for Phase I studies where concepts such as the MTD would mean different things to different people with no common definition in place. Following (Storer, 1989), the idea of the MTD being a population quantity that we ought try to estimate based on the result of a study, gained ground and formed the basis for model based designs such as the continual reassessment method (O’Quigley et al., 1990). Model based methods have continued to develop with recent designs such as CCD (Ivanova et al., 2007), mTPI (Ji et al., 2010) and BOIN (Liu and Yuan, 2015) being successfully used in recent years.

All model based designs have as their objective the identification of the MTD where the MTD is defined as the dose being the closest to some target. The target dose is described as some ideal dose, conceptualized on a continuous spectrum, at which the rate of toxicity is equal to some, acceptable, target rate. This target rate is chosen by the clinical investigators and usually corresponds to 1 patient out of 5, 1 out of 4 or 1 out of 3. In principle any rate lying

between zero and one works but those three rates are by far the ones most commonly chosen. The number of doses under study is usually quite small, often around 6, but sometimes less and rarely much greater than 10.

The problem here is that the above definition of the MTD, while a much sharper one than has often been used, in the past and even the present, is not unambiguous. There is enough lack of precision surrounding the definition for us to be unable to state unequivocally whether some proposed method will, with increasing sample size, be able to identify the MTD with probability one. Usually, we would like to have this as a minimal requirement. The problem occurs because the concept “closest to some target” cannot be made precise without reference to the topological structure in which the problem is set. The common Kullback-Leibler distance (or divergence) depends on the chosen model. Different models could then imply different MTDs. Even without a model, the dose closest to the target based on Euclidean distance may not be the same dose as one based on the logarithm of that distance. It is common practice to work with Euclidean distance but there is no compelling reason for this and, indeed, more linear type measures such as $\log(-\log)$ can be seen as more natural.

The statistical challenge of efficient estimation of the MTD, while accommodating ethical requirements of experimentation, becomes impossible once we recognize that the definition of the MTD is not precise. Deeper thought tends to suggest that it is not really the MTD that we wish to estimate. It is more likely a cut between levels whereby the level above the cut has an associated rate of toxicity higher than some acceptable amount and the level below the cut would have a rate below the target. Such a change in emphasis may appear to be subtle but it has quite deep ramifications. More importantly it leads us in the direction of direct modelling of something that we can sharply define, the cut or the couple of dose associated, and that can be very well expressed within a hierarchical Bayesian framework. Such hierarchical modelling has at the one end, clear inference relating to an interpretable parameter of direct interest (such as a cut), while at the other end we can incorporate all of the features that we would hope for a method to benefit from. Great generality and great flexibility become immediately available. The purpose of this article is to investigate such hierarchical Bayesian modelling in this setting and to show, via several examples, that we obtain methods that are in all cases, as good, if not better, than methods that are currently available.

5.2 Semi-Parametric Modelling of the Maximum Tolerated Dose

The observations are the sequences: $(X_n, Y_n)_{n \in \mathbb{N}}$ where the dose X_n can be viewed as the predictor for Y_n , the observed toxicity and dependent variable taking values in $\{0, 1\}$ (1 for a DLT and 0 otherwise). The conditional distribution of Y_n given $X_n = d$ is Bernoulli with parameter β_d which means that we assume that the toxic probability at dose d does not depend of the way in which patients are included into the study. The doses of range D have been chosen by clinical expertise so that the Bernoulli distributions are stochastically ordered.

Assumption 5.2.1. $\forall n \in \mathbb{N}, \beta_d = \mathbb{P}(Y_n = 1 | X_n = d)$, with: $\beta_1 < \beta_2 < \dots < \beta_m$.

The true situation is then fully described by a scenario of increasing toxicities $(\beta_1, \dots, \beta_m)$. The goal of the study is to find the MTD from the doses available in the range D , having a probability of toxicity the closest to some chosen amount α fixed by the investigators.

Definition 5.2.1. *According to a certain metric $\Delta(\cdot, \cdot)$, the MTD is the dose d^* in D such that:*

$$\forall d \in D, \Delta(\beta_d, \alpha) \geq \Delta(\beta_{d^*}, \alpha).$$

The most frequently chosen metric for this situation, although not necessarily the most natural, is one based on Euclidean distance. The aim of identifying the MTD with increasing precision is hampered by the ethical constraint that we must, at each step, treat the patients at what is believed to be the best level for them, i.e., the current estimator of the MTD itself. In chapter 3, a wide class of semi-parametric methods (SPM), is introduced. Under the SPM heading we can include several well-known model based methods ranging from the CRM (O'Quigley et al., 1990) to the mTPI (Ji et al., 2010) or the BOIN method (Liu and Yuan, 2015)(chapter 4). The hierarchical approach employed leads to very ready improvements to those methods in terms of accuracy of final recommendation while rigorously respecting the ethical criterion. Let us briefly recall the notation and main ideas behind SPM. The key point is to directly model the MTD itself. The letter F denotes the vector space of Bernoulli parameters needed to cover the space of possible scenarios: $F = [0, 1]^m$. 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 term of the main parameter of interest, θ , which leads to the creation of m distinct classes, each class containing an infinite set of members all of which share the same MTD.

$$F = \bigcup_{\theta \in D} F_\theta, \text{ and } F_\theta = \{q \in F : \forall j \in D, |q_\theta - \alpha| \leq |q_j - \alpha|\}. \quad (5.1)$$

For all β in F_θ , the MTD is indexed by θ . Let n_j^1 be the number of toxicities at dose j and $n_j = n_j^1 + n_j^0$, the number of patients treated at dose j . At step n , the likelihood function is $L_n(q) = \prod_{1 \leq j \leq m} q_j^{n_j^1} (1 - q_j)^{n_j^0}$. The set F is endowed with a probability measure $\Lambda \otimes \Pi$ such that Π is a measure on D and the topological support of the measure $\Lambda(\cdot | \theta) = \Lambda_\theta(\cdot)$ is included in the class F_θ . The posterior distribution of θ and $(q | \theta)$ given (X_1^n, Y_1^n) are:

$$\Pi_n(\theta) \propto \left[\int L_n(q) \Lambda_\theta(dq) \right] \times \Pi(\theta) \text{ and } \Lambda_{\theta, n}(A) \propto \int_A L_n(q) \Lambda_\theta(dq), \quad (5.2)$$

with A an element of the sigma-algebra associated with Λ_θ . 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 . The posterior of Π is obtained by weighting according to the expected value of the likelihood with respect to $(\Lambda_\theta)_{\theta \in D}$. Finally,

the dose for which Π_n attains its maximum value is the natural estimator of the MTD:

$$\hat{\theta}_n = \arg \max_{\theta \in D} \Pi_n(\theta). \quad (5.3)$$

5.3 Minimal sets, cuts and large sample behavior

Definition 5.3.1. *The minimal set M^* is the smaller set of doses d for which the MTD may be expressed as a function of the exact dose/toxicity relation (d, β_d) :*

- $\exists f$, a function such that: $d^* = f[(d, \beta_d)_{d \in M_D}]$.
- $\forall M \subsetneq M_D$, $\nexists f$, a function such that: $d^* = f[(d, \beta_d)_{d \in M}]$.

In the canonical case (assumption 5.2.1) that we study, the minimal is simple: $M^* = \{b, a\}$ where b (below) and a (above) are the two consecutive doses associated with probability of toxicity on each side of the threshold α . there exist other settings, as drug combinations, where the minimal set is less obvious. We now explain why the minimal set is a central notion of dose finding studies. A basic definition and a simple assumption on the prior model helps our development.

Definition 5.3.2. *Let $d(\cdot, \cdot)$ be the euclidean distance. A sequence $(X_n)_{n \in \mathbb{N}}$ converges to a set B , denoted by $X_n \xrightarrow{S} B$, if: $\sup_{x \in B} \left(\liminf_{n \rightarrow +\infty} d(X_n, x) \right) = 0$.*

Assumption 5.3.1. Λ_θ is a product of unidimensional distributions: $\Lambda_\theta(dq) = \Lambda_\theta^1(dq_1) \times \dots \times \Lambda_\theta^m(dq_m)$; and its topological support S_θ satisfies: $S_\theta = B^{\theta-1} \times I_\epsilon \times A^{m-\theta}$, with $I_\epsilon = [\alpha - \epsilon, \alpha + \epsilon]$, $A = [\alpha + \epsilon, 1]$ and $B = [0, \alpha - \epsilon]$, with $\epsilon \geq 0$.

Note that the central interval I_ϵ can be reduced as much as we wish and, ultimately, to a single point α . In that case, for all θ , the distribution Λ_θ^θ is a Dirac law in α . For SPM and under certain regularity assumptions, two types of asymptotic behaviour occur. The first is almost sure convergence to a single dose having a true rate of toxicity contained within I_ϵ ; the second is almost sure convergence to the minimal set when the two doses are associated with toxicity rates lying outside the interval I_ϵ .

Definition 5.3.3. *Let us 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) *A method is ϵ -sensitive, if for all β such that $\mathcal{E}(I_\epsilon, \beta) \neq \emptyset$, we have: $\mathbb{P}_\beta[\exists N, \forall n > N : X_n \in \mathcal{E}(I_\epsilon, \beta)] = 1$.*

(b) *A method is balanced, if for all β such that $\mathcal{E}(I_\epsilon, \beta) = \emptyset$, we have: $X_n \xrightarrow{S} M^*$, a.s.*

At the first sight, the property ϵ -sensitive may appear more compelling than balanced behaviour. However, this convergence depends on the scenario and the parametrization of the prior model. If two or more doses are associated with toxicities belonging to I_ϵ , ϵ -sensitivity

will not ensure the convergence to the MTD. Conversely, if neither dose is close enough to the threshold, the SPM is balanced which means that the current estimator shall recommend alternatively the two doses of the minimal set. This couple of consecutive doses are the set on which we need to have observations when the goal is that of determining almost surely the MTD. In fact, in the absence of a strong parametric model, experimentation outside of this minimal set brings less and less, asymptotically negligible, information regarding estimation of the MTD. This intuitive result is a corollary of the impossibility theorem of Azriel et al. (2011) showing that no method exists that would, for all situations, allow the current estimator to be strongly consistent at the MTD. Thus, the balanced behaviour is necessary to attain asymptotically the practical goal of the study. When ϵ is equal to 0, and under some assumptions on the prior model (chapter 3, Large sample theory), the SPM is always balanced, whatever the scenario, and this provides us with the ability to build consistent estimators of the MTD on the basis of the empirical distribution of the toxicities or by using an isotonic regression at the end of the trial. For such a parametrization ($\epsilon = 0$), performance according to both the practical and ethical criteria is good, matching that of SPM in which $\epsilon > 0$.

The ethical criteria are based on the quality of treatment for all patients entered into the actual trial; our focus is the average number of patients treated at the MTD or adjacent levels in term of Euclidean distance. Having as our objective the accurate identification of the MTD does not result in any deterioration of performance with respect to the ethical constraints. This may seem to contradict what is suggested by the title of the article: *The treatment versus experimentation dilemma in dose finding studies* (Azriel et al., 2011). This dilemma does exist but is mainly a theoretical one and is stated by the impossibility theorem. This theorem expresses the complementary and antagonistic asymptotic properties of any method: the ϵ -sensitivity and the balanced behaviour. When the recommendations are concentrated on the minimal set $\{b, a\}$, which can be seen as minimal concerning the goal of the study, this helps improve the whole performance of the method. An important and difficult question arises out of this analysis: how should this concentration be best distributed between the two members of the set. Some answers are given by the quantification of the asymptotic ratio of observations on doses b and a . In chapter 3, two convergence results are obtained: $n_a/n_b \rightarrow D_{KL}(\beta_b, \alpha)/D_{KL}(\beta_a, \alpha)$ *a.s.*, for classical SPM with $\epsilon = 0$, or $n_a/n_b \rightarrow 0$ *a.s.*, if b is the MTD and vice versa, for SPM with a varying ϵ .

Minimal sets can be expressed in term of cuts between levels whereby whereby the dose above the cut has an associated rate of toxicity higher than α and the dose below the cut a rate lower than α . For the remainder of this paper, our focus is on estimating the Maximum Tolerated Cut, or equivalently the minimal set. Our first objective is to demonstrate the good properties of this approach. Not only do we not lose any precision regarding estimation of the MTD when compared with efficient methods whose purpose is only MTD estimation, but also we do not suffer any losses regarding allocation precision (the effective treatment of patients). Putting emphasis on the MTC leads us to consider sampling (allocation) strategies on the

minimal set. Little attention has been directly focused on this aspect and it helps shine a spotlight on the precise objectives and sampling constraints of the study. The idea of a cut is implicit in the work of optimal design (O’Quigley et al., 2002). Focusing on the cut itself, viewed as a contour in the setting of drug combinations, has been studied by Mander and Sweeting (2015). We complete the picture by a direct modelling of the cut in the canonical case of dose finding studies.

5.4 Modelling of the Maximum Tolerated Cut (SPMc)

Recall from section 5.2 that D is a discrete range of m doses ordered in terms of increasing probability of toxic response. According to the assumption 5.3.1, D can be partitioned into two sets of consecutive doses: the doses associated with probabilities of toxicity lower and those with probabilities of toxicity greater than α so that:

$$D = D^- \sqcup D^+ = \{j \in D : \beta_j < \alpha\} \sqcup \{j \in D : \beta_j > \alpha\}.$$

From a large sample perspective, we only have to approximate the toxicity of the two doses, $b = \max D^-$ and $a = \min D^+$, or a single dose when D^- or D^+ is empty. Thus, the couple $M^* = \{a, b\}$ is the minimal set on which we need to have observations. There are $m+1$ possible ways to cut the set D into D^- and D^+ . Our model is then structured around the set of all possible locations of a cut: before dose 1, in $]1, 2[$, $]2, 3[$, \dots , $]m-1, m[$ or after dose m . We want to find the MTC, noted c^* , from the range $C = \{1, m+1\}$, or equivalently the minimal set M^* . Cut 1 corresponds to the case where no dose has a toxicity lower than the threshold and cut $m+1$ is the case where all the toxicities are greater than α ; in these particular cases the MTC corresponds to $\{1\}$ and $\{m\}$ respectively. Otherwise, when the cut j lies within $\llbracket 2, m \rrbracket$, the MTC is the set $\{j-1, j\}$. Note that estimating the MTC and concentrating experimentation at those two doses allows us also to estimate the MTD regardless of the metric chosen to define this dose (definition 5.2.1). The semi-parametric modelling of the MTC is noted SPMc. The space F (see (5.1)) can be written as:

$$F = \bigcup_{\theta \in C} F'_\theta, \text{ with } F'_\theta = \{q \in F : j < \theta \Rightarrow q_j \leq \alpha \text{ and } j \geq \theta \Rightarrow q_j \geq \alpha\}, \quad (5.4)$$

with θ indexing the cut. The set F is now endowed with a probability measure $\tilde{\Lambda} \otimes \tilde{\Pi}$ such that $\tilde{\Pi}$ is a measure on C and the topological support of the measure $\tilde{\Lambda}_\theta(\cdot)$ is included in the class F'_θ . As for modelling the MTD, the posterior distribution of θ and $(q|\theta)$ given (X_1^n, Y_1^n) are:

$$\tilde{\Pi}_n(\theta) \propto \left[\int L_n(q) \tilde{\Lambda}_\theta(dq) \right] \times \tilde{\Pi}(\theta) \text{ and } \tilde{\Lambda}_{\theta,n}(A) \propto \int_A L_n(q) \tilde{\Lambda}_\theta(dq), \quad (5.5)$$

with A an element of the sigma-algebra associated with $\tilde{\Lambda}_\theta$. A simple parametrization of the prior-model $(\tilde{\Lambda}_\theta)_{\theta \in C}$ makes the following assumption:

Assumption 5.4.1. $\tilde{\Lambda}_\theta$ is a product of unidimensional distributions:

$$\tilde{\Lambda}_\theta(dq) = \tilde{\Lambda}_\theta^1(dq_1) \times \dots \times \tilde{\Lambda}_\theta^m(dq_m),$$

and its topological support \tilde{S}_θ is equal to $\tilde{B}^{\theta-1} \times \tilde{A}^{m-\theta+1}$, with $\tilde{B} = [0, \alpha]$ and $\tilde{A} = [\alpha, 1]$.

Note that the union of all the topological supports of $\tilde{\Lambda}_\theta$ is equal to F , the set of all admissible scenarios according to the monotonicity assumption. Each possible scenario is included in a particular \tilde{S}_θ . A natural estimator of the MTC is $\tilde{\theta}_n = \arg \max_{\theta \in C} \tilde{\Pi}_n(\theta)$. However, this estimator does not of itself provide an unequivocal choice for the next dose. It does provide us with the current estimate of the minimal set, $\tilde{M}_n = \{\tilde{\theta}_n - 1, \tilde{\theta}_n\}$. In order to choose the next dose in \tilde{M}_n , we require some allocation strategy. The allocation strategy can be framed within the context of 2-armed bandit problem in which we need to determinate the treatment arm closest to some threshold with respect to some metric.

Allocation strategy 1 One possibility is to use the SPMc-model itself on \tilde{M}_n . At step n , the next dose is the one with the highest average probability $\tilde{\Pi}_n$ on either side of the interval:

$$\hat{d}_n^1 = \arg \max_{\theta \in \tilde{M}_n} \tilde{\Pi}_n(\theta) + \tilde{\Pi}_n(\theta + 1).$$

- (1) Update $\tilde{\Pi}_n$ and determine the most likely minimal set \tilde{M}_n .
- (2) Use an allocation strategy \hat{d}_n on \tilde{M}_n for selecting the next dose.

All of the selection strategies proposed in that article satisfy the following assumption.

Assumption 5.4.2. For all couples M , we have:

$$\left\{ \sum_{i=1}^n \mathbb{1}_{\{\tilde{M}_n=M\}} \xrightarrow{n \rightarrow +\infty} +\infty \right\} \Leftrightarrow \left\{ n_j \xrightarrow{n \rightarrow +\infty} +\infty, \forall j \in M \right\}.$$

This means that the doses of the minimal set will be fully explored asymptotically. If the SPMc converges almost surely to the true cut, the balanced behavior is a straightforward consequence of this assumption. The consistency of $(\tilde{\theta}_n)_{n \in \mathbb{N}}$ can be obtained under the following regularity constraints.

Assumption 5.4.3. For all $d \in D$ and $c \in C$, the marginal distribution Λ_c^d is absolutely continuous with respect to Lebesgue measure and λ_c^d denotes its density function. The prior-model respects one of the following properties:

- (i) There exist two numbers s and S in \mathbb{R}_+^* , such that, for all c and d , we have:

$$\forall q_d \in S_c^d, s < \lambda_c^d(q_d) < S.$$

(ii) $\forall (c, c') \in C^2$, with $c < c'$, $\forall d : d < c$ or $d \geq c'$, we have $\Lambda_c^d = \Lambda_{c'}^d$.

These two points describe different parametrizations of SPM. As seen in the section *Large sample theory* of chapter 3, the first one can be easily obtained, and the second assumption covers the case of local models such as SP-mTPI.

Theorem 5.4.1. *Under assumptions 5.4.1, 5.4.2 and 5.4.3, the SPMc is balanced.*

Proof. We first prove the balanced property. The assumption of independance allows us to focus on the marginal ratio:

$$\frac{M_{n,r}^k}{M_{n,t}^k} = \frac{\int g(q_j, n_k^1, n_k^0) \Lambda_r^k(dq_k)}{\int g(q_k, n_k^1, n_k^0) \Lambda_t^k(dq_k)} \quad \text{and} \quad \frac{I_{n,r}}{I_{n,t}} = \prod_{k \in D} \frac{M_{n,r}^k}{M_{n,t}^k} \quad \text{with} \quad (t, r) \in C, \quad k \in D.$$

Note that the assumptions 5.4.1 (or alternative assumption 5.4.1 (ii)) and Lemma 3.6.1 (ii) involves:

$$d(\beta_k, S_r^k) = d(\beta_k, S_t^k) \implies 0 < \liminf_{n \rightarrow +\infty} \frac{M_{n,r}^k}{M_{n,t}^k} \leq \limsup_{n \rightarrow +\infty} \frac{M_{n,r}^k}{M_{n,t}^k} < +\infty, \quad a.s. \quad (5.6)$$

We will show that:

$$M = \{c-1, c\} \neq M^* \implies \mathbb{P}(\{n_M \rightarrow +\infty\}) = 0, \quad (5.7)$$

with $n_M = \sum_{i=1}^n \mathbb{1}_{\{\tilde{M}_n = M\}}$. As $M \neq M^*$, there exists $k \in M$ such that β_k is included in $S_{c^*}^k$ and not in S_c^k . By using Lemma 5.8.1, we have:

$$\mathbb{P}\left(\lim_{n \rightarrow +\infty} \frac{M_{n,c}^k}{M_{n,c^*}^k} = 0 \mid \{n_M \rightarrow +\infty\}\right) = \mathbb{P}\left(\lim_{n \rightarrow +\infty} \frac{M_{n,c}^k}{M_{n,c^*}^k} = 0 \mid \{n_k \rightarrow +\infty\}\right) = 1,$$

where the first equality arises from assumption 5.4.2. For all doses $d \in D$, Λ_{c^*} consistently models the rate of toxicity which means: $\beta_d \in S_{c^*}$. We then have:

$$\mathbb{P}\left(\lim_{n \rightarrow +\infty} \frac{I_{n,c}}{I_{n,c^*}} = 0 \mid \{n_M \rightarrow +\infty\}\right) = 1 \quad \text{and} \quad \mathbb{P}\left(\left\{\lim_{n \rightarrow +\infty} \frac{I_{n,c}}{I_{n,c^*}} = 0\right\} \cap \{n_M \rightarrow +\infty\}\right) = 0,$$

which proves the balanced property. \square

The main implication of this theorem is the ability to construct estimators that are almost surely convergent. Such a property was already obtain for the classical SPM (section 5.2). In the remainder of this section, we look more closely at the relative relations of SPM modelling in terms of the MTC or in term of the MTD. In order to do this we need to make the following minor modification to SPM: the range of dose D on which is built the prior Π is extended to the set $\bar{D} = \{0, \dots, m+1\}$. The new doses 0 and $m+1$ are fictive doses and must be considered as additional parametrization corresponding respectively to dose 1 and m , which means that

the new natural estimator will be:

$$\hat{\theta}_n = \arg \max_{\theta \in D} \bar{\Pi}_n(\theta), \quad (5.8)$$

$\bar{\Pi}_n(1) = \Pi_n(0) + \Pi_n(1)$, $\bar{\Pi}_n(m) = \Pi_n(m) + \Pi_n(m+1)$ and equality otherwise: $1 < j < m$, $\bar{\Pi}_n(j) = \Pi_n(j)$. The distributions Λ_0 and Λ_{m+1} model the case where all the toxicities are, respectively, in A and B . The prior-model satisfies the following assumption.

Assumption 5.4.4. *The distributions Λ_θ are products of unidimensional distributions. For $\theta \in \llbracket 1, m \rrbracket$, their topological support satisfies assumption 5.3.1 with $\epsilon = 0$, and we have: $S_0 = A^m$ and $S_{m+1} = B^m$.*

Modelling the MTD this way leads naturally to a way to model the MTC by using a mean operator on the doses belonging to the potential candidates for being the minimal set:

$$\forall \theta \in C, \tilde{\Pi}(\theta) \propto \Pi(\theta-1) + \Pi(\theta) \text{ and } \tilde{\Lambda}_\theta = r_\theta \Lambda_{\theta-1} + (1-r_\theta) \Lambda_\theta \text{ with } r_\theta = \frac{\Pi(\theta-1)}{\Pi(\theta-1) + \Pi(\theta)}. \quad (5.9)$$

When the SPM works under assumption 5.4.4, the SPMc defined by (5.9) satisfies assumptions 5.4.1. In particular, the marginals $\tilde{\Lambda}_\theta^{\theta-1}$ and $\tilde{\Lambda}_\theta^\theta$ have an atom in α and their supports are $[0, \alpha]$ and $[\alpha, 1]$, respectively. The relation between the distribution Π and $\tilde{\Pi}$ remains unchanged.

Property 5.4.2. *If $\tilde{\Lambda} \otimes \tilde{\Pi}$ satisfies (5.9), we have: $\tilde{\Pi}_n(\theta) \propto \Pi_n(\theta-1) + \Pi_n(\theta)$.*

Proof. We have:

$$\begin{aligned} \Pi_n(\theta) \propto \left[\int L_n(q) \tilde{\Lambda}_\theta(dq) \right] \times \tilde{\Pi}(\theta) &= \left[\int L_n(q) (r_\theta \Lambda_{\theta-1}(dq) + (1-r_\theta) \Lambda_\theta(dq)) \right] \\ &\times (\Pi(\theta-1) + \Pi(\theta)) = F(\theta), \end{aligned}$$

and

$$\begin{aligned} F(\theta) &= \left[\int L_n(q) \Lambda_{\theta-1}(dq) \right] \times r_\theta (\Pi(\theta-1) + \Pi(\theta)) \\ &+ \left[\int L_n(q) \Lambda_\theta(dq) \right] \times (1-r_\theta) (\Pi(\theta-1) + \Pi(\theta)) \\ &= \Pi_n(\theta-1) + \Pi_n(\theta). \end{aligned}$$

□

This results shows that an MTD-model leads naturally to an MTC-model. All MTC-models however are not the result of an MTD model. The class of MTD models is all the greater given that a wide variety of allocation strategies can be used. Constructing an MTD-model from an MTC-model (5.9) is also possible by taking the mean on the cuts for which the dose

lies in the minimal set:

$$\forall \theta \in D, \Pi(\theta) \propto \tilde{\Pi}(\theta) + \tilde{\Pi}(\theta + 1) \text{ and } \Lambda_\theta = \tilde{r}_\theta \tilde{\Lambda}_{\theta-1} + (1 - \tilde{r}_\theta) \tilde{\Lambda}_\theta, \quad (5.10)$$

with $\tilde{r}_\theta = \tilde{\Pi}(\theta)/[\tilde{\Pi}(\theta) + \tilde{\Pi}(\theta + 1)]$. This leads to the same theoretical result as Property 5.4.2. When the SPMc satisfies assumption 5.4.1, the topological support of Λ_θ^θ is then the entire parameter space $[0, 1]$, for all θ . Moreover, when the law $\tilde{\Pi}_n$ remains unimodal throughout the trial, the SPM resulting from the SPMc is precisely equivalent to SPMc working with selection strategy 1. The reason for this is that, in the particular case when $(S_\theta^\theta = [0, 1], \forall \theta \in D)$, the SPM is ϵ -balanced. Unimodality at each step of the trial is a desirable property, which is very easy to obtain in practice for any reasonable model; it depends on the parametrization of $\tilde{\Lambda} \otimes \tilde{\Pi}$ and can be checked empirically.

We can conclude that the MTD and the MTC point of view are very similar. However, the MTC viewpoint is richer. It provides us with information not just on our current best estimate of the MTD and on our next best estimate the other side of the threshold, but also we focus, through the allocation strategy on the relative probabilities either dose has ultimately of being the closest to the threshold. Obtaining information at both doses will not only help best identify the MTD but will provide valuable information on the second dose that can be of great value for example when we terminate the initial dose-escalation study and proceed to investigating dose expansion cohorts.

5.5 A family of allocation strategies

All methods iterate between updating the current estimates of the MTD, or, here, the MTC, and then allocating the next patient or patient cohort to some level, most often the MTD itself. Almost all discussion in the field focuses on sequential estimation and very little on allocation. And yet the problem of allocation is a very important one in its own right and remains a relatively uncharted area. Different allocation strategies will result in different types of behavior and the formulation of an allocation strategy can also be used in a formal way to define certain methods such as the mTPI, BOIN and CCD methods. Any sensible allocation strategy will want to make use of some metric describing the distance between levels and the MTD, as well as the total amount of information at the relevant levels. For example, suppose the target is 0.2 and the estimated rates on the estimated minimal set are 0.18 and 0.90. Clearly we would want to very heavily favour the lower of these two levels as regards allocation, and, this, all the more so as the precision at these levels increases. Two kinds of information are available on the doses that make up the estimated minimal set: (a) the distance between the rates of toxicity and the threshold and (b) the numbers of patients treated at each dose. The different strategies described here take care of either or both of these. Let R be a function of the threshold α , an estimator $\hat{\beta}_{k,n}$ of the toxicity at dose k and the observation's number n_k at this dose. Let $s \in \mathbb{R}^+$ be a cutoff point. The current estimator \hat{d}_n selects a dose in the

estimated minimal set $\tilde{M}_n = \{k, k+1\}$:

$$\hat{d}_n(s) = \begin{cases} k & , \text{ if } R(\alpha, \hat{\beta}_{k,n}, n_k)/R(\alpha, \hat{\beta}_{k+1,n}, n_{k+1}) \leq s \\ k+1 & , \text{ otherwise.} \end{cases} \quad (5.11)$$

R is as an increasing function of n_k and a distance Δ between α and $\hat{\beta}_{k,n}$. In particular, R can take the following form, with $w \in [0, 1]$:

$$R(\alpha, \hat{\beta}_{k,n}, n_k) = (n_k)^w \times \Delta(\alpha, \hat{\beta}_{k,n})^{(1-w)}. \quad (5.12)$$

Note that such a strategy with $w = 0$ does not meet the requirement of Assumption 5.4.2. Azriel et al. (2011) proposed such an estimator for the MTD and use a randomized allocation design in order to meet the requirement of making observations either side of the threshold. Randomization can be avoided if we so wish by replacing w by a sequence $(w_{n'_M})_{n_M \in \mathbb{N}}$ converging to 0 with n'_M the number of observations on the minimal set $M = \{k, k+1\}$: $n'_M = n_k + n_{k+1}$. Under the following assumption, the estimator \hat{d}_n will have good asymptotic properties.

Assumption 5.5.1. *The allocation strategy fulfills (5.12) and, when $n'_M \rightarrow +\infty$, the decreasing sequence $(w_{n'_M})_{n \in \mathbb{N}}$ converges to 0 with $w_{n'_M} = o(1/\log(n'_M))$. The family of estimators $(\hat{\beta}_j)_{j \in D}$ is almost surely convergent: $\mathbb{P}(\hat{\beta}_{j,n} \rightarrow \beta_j | \{n_j \rightarrow +\infty\}) = 1$. The function $\Delta(\alpha, \cdot)$ is continuous on the whole parameter space $[0, 1]$, with $\Delta(\alpha, \hat{\beta}_{k,n}) = 0$ if and only if $\alpha = \hat{\beta}_{k,n}$.*

Allocation strategy 2: This simple allocation strategy is based on the number of patients observed at each dose of the minimal set, which correspond to $w = 1$ in (5.12).

$$\hat{d}_n^2(s) = \begin{cases} k & , \text{ if } n_k/n_{k+1} \leq s \\ k+1 & , \text{ otherwise.} \end{cases} \quad (5.13)$$

Setting $p = 1$, corresponds to the case in which the investigator would like to have the same number of observation on the doses either side of the threshold.

Allocation strategy 3: This strategy use the formula (5.12) with the same weight w on the two parts of (5.12). $\Delta(\alpha, \cdot)$ is continuous on the whole parameter space. We can describe this strategy in the particular case where Δ is the divergence of Kullback-Leibler. 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: $D_{KL}(p||q) = H(q|p) - H(p) = p \log(\frac{p}{q}) + (1-p) \log(\frac{1-p}{1-q})$. We provide the following estimator:

$$\hat{d}_n^3 = \arg \min_{j \in \tilde{M}_n} n_j \times D_{KL}(\hat{\beta}_{j,n} || \alpha). \quad (5.14)$$

When the estimator of the rate of toxicity is the frequentist estimator, $\widehat{\beta}_{j,n} = n_j^1/n_j$, the allocation strategy can be written down in a more intuitive manner. The dose selected by $\widehat{d}_n^{\mathfrak{B}}$ is the one in \widetilde{M}_n for which the relative likelihood function at α is the highest:

$$\widehat{d}_n^{\mathfrak{B}} = \arg \max_{j \in \widetilde{M}_n} \left(\frac{\alpha^{n_j^1} (1 - \alpha)^{n_j^0}}{\alpha^{n_j \alpha} (1 - \alpha)^{n_j (1 - \alpha)}} \right).$$

However, an estimator of the toxicity based on an isotonic regression should obtain better results by using the available information on the whole range of doses.

Allocation strategy 4: We propose to use the semi-parametric model $\Lambda \otimes \Pi$ satisfying assumption 5.3.1 on the doses of \widetilde{M}_n together with a probability's cutoff p over which the more toxic dose, $\tilde{\theta}_n$, will be selected.

$$\widehat{d}_n^A(s) = \begin{cases} k & , \text{ if } \Pi_n(k) > s \\ k + 1 & , \text{ otherwise .} \end{cases} \quad (5.15)$$

We have then a family of estimators with differing degrees of conservativeness according to the value of s . Note that the cutoff s can be replaced by a sequence $(s_n)_{n \in \mathbb{N}}$. With the aim of using all the available information, the traditional likelihood function can be replaced by the isotonic likelihood function:

$$L_n^*(q) = q^{n_j \hat{t}_{j,n}} (1 - q)^{n_j (1 - \hat{t}_{j,n})} \quad (5.16)$$

where $\hat{t}_{j,n}$ is the estimator of the isotonic regression of the toxicity at dose j .

We now set out some asymptotic properties of these selection strategies.

Theorem 5.5.1. *In all cases, we suppose that SPMc works under assumption 5.4.1 and 5.4.3.*

1) *If the SPMc works with $\widehat{d}_n^{\mathfrak{B}}$ and the estimators of the rate of toxicity are consistent: for all $j \in D$, $\widehat{\beta}_{j,n}$ converges almost surely to β_j on $\{n_j \rightarrow +\infty\}$, then we have:*

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

2) *Under assumption 5.5.1, the method is balanced and if $\Delta(\alpha, \beta_a) > \Delta(\alpha, \beta_b) > 0$ or $\Delta(\alpha, \beta_b) > \Delta(\alpha, \beta_a) > 0$ we have, a.s.: $n_b/n \rightarrow 1$ or $n_a/n \rightarrow 1$, respectively.*

Proof. 1) These allocation strategies satisfy assumption 5.4.2 and the SPMc is balanced. We focus on the convergence of the ratio n_a/n_b . As the proof is almost the same in these two cases, only the case of relative entropy will be treated with $\widehat{\beta}_{j,n} = n_j^1/n_j$. Let W be the space of full measure on which the SPMc is balanced and the law of large numbers is valid for

the frequentist ratios n_b^1/n_b and n_a^1/n_a . We set:

$$J_n = \frac{\exp \left[-n_a \times \Delta(\alpha, \widehat{\beta}_{a,n}) \right]}{\exp \left[-n_b \times \Delta(\alpha, \widehat{\beta}_{b,n}) \right]}.$$

On W and for n large enough, we have:

$$\widehat{d}_n^1 = \begin{cases} a & , \text{ if } J_n > 1 \\ b & , \text{ otherwise .} \end{cases} \quad (5.17)$$

We argue by contradiction, in supposing that there exists a subsequence $(n')_{n \in \mathbb{N}}$ and real $t > 0$ such that:

$$n'_a > n'_b(t + \Delta(\alpha, \beta_b)/\Delta(\alpha, \beta_a)).$$

We then have:

$$J_{n'} < \frac{\exp \left[-n'_b \Delta(\alpha, \widehat{\beta}_{a,n'}) \left(\frac{\Delta(\alpha, \beta_b)}{\Delta(\alpha, \beta_a)} \right) \right]}{\exp \left[-n'_b \Delta(\alpha, \widehat{\beta}_{b,n}) \right]} \exp \left[-t n'_b \Delta(\alpha, \widehat{\beta}_{a,n}) \right].$$

On W , we have: $\Delta(\alpha, \widehat{\beta}_{b,n}) \xrightarrow{n_b \rightarrow +\infty} \Delta(\alpha, \beta_b)$ and $\Delta(\alpha | \widehat{\beta}_{a,n}) \xrightarrow{n_b \rightarrow +\infty} \Delta(\alpha | \beta_a)$. Therefore, there exists $\delta > 0$ such that, for n large enough, we have: $\Delta(\alpha | \widehat{\beta}_{a,n}) > \delta$. Thus, the sequence $(J_{n'})_{n \in \mathbb{N}}$ converges to 0. If we accept that the sequence $(J_n)_{n \in \mathbb{N}}$ can not have an accumulation point in 0, this then leads to a contradiction.

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

It remains to show that $(J_n)_{n \in \mathbb{N}}$ does not have a sub-sequence converging to zero. Let n' be the infinite subsequence of n such that $\widehat{\theta}_{n'} \neq \widehat{\theta}_{n'-1}$. The SPM being ϵ -balanced, the sequence $(J_{n'})_{n \in \mathbb{N}}$ falls an infinite number of times in each of the intervals $[0, 1]$ and $[1, +\infty]$. The variations $(J_{n+1})/J_n$ being bounded below by a strictly positive number K , we obtain the expected result for the subsequence $(J_{n'})_{n \in \mathbb{N}}$.

Let $n_{a,k} = n_a - \lfloor n_a/k \rfloor$, where $\lfloor \cdot \rfloor$ is the floor function. $n_{a,k}^1$ and $n_{a,k}^0$ are defined as usual: $n_{a,k}^1 = \sum_{i \in N(n,a,k)} \mathbb{1}_{\{Y_i=1\}}$ with $N(n,a,k) = \{i \in \mathbb{N}, X_i = a, Y_i = 1, n_{a,k} \leq i \leq n_a\}$ and $n_{a,k} = n_{a,k}^1 + n_{a,k}^0$. Moreover, for all $j \in D$, we have $\widehat{\beta}_{j,n_{a,k}} \rightarrow \beta_j$ almost surely on $\{n_{a,k} \rightarrow +\infty\}$ and thus:

$$\exp \left[-n_{a,k} \times \Delta(\alpha, \widehat{\beta}_{j,n_{a,k}}) \right] < 1/K \text{ a.s.} \quad (5.18)$$

We set $\varphi(n) = \max_{n' \in E_n} (n' - (n-1)')$, where $E_n = \{n' < n, \widehat{\theta}_{(n-1)'} = a\}$. We suppose that there exists a strictly positive number k such that for n large enough: $\varphi(n) - 1 > n/k$. Then, for all N , there exists $n' > N$ such that $\widehat{\theta}_{(n-1)'} = a$ and $n' - (n-1)' > n/k$. Using the inequality

(5.18), there exists N large enough, such that if $n' > N$, $\hat{\theta}_{(n-1)'} = a$ and $n' - (n-1)' > n/k$ then we can find $n \in \llbracket (n-1)'; n' \rrbracket$ with $\hat{\theta}_n = b$, which is a contradiction. Thus, we have: $n' - (n-1)' \in o(n)$. Hence the result:

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

2) As $\Delta(\alpha, \beta_a) > 0$ and $\Delta(\alpha, \beta_b) > 0$, no dose is associated with a rate of toxicity equal to the threshold. Suppose that $\Delta(\alpha, \beta_a) > \Delta(\alpha, \beta_b) > 0$. We show that the allocation strategy considered fulfills assumption 5.4.2. For a minimal set $M = \{k, k+1\}$, we fix: $n_M = \sum_{i=1}^n \mathbb{1}_{\{\tilde{M}_n = M\}}$. Recall that: $R_n(\alpha, \hat{\beta}_{k,n}, n_k) = (n_k)^{w_{n_M}} \times \Delta(\alpha, \hat{\beta}_{k,n})^{(1-w_{n_M})}$. According to the event $A_k = \{n_k \rightarrow +\infty\} \cap \{n_{k+1} \rightarrow +\infty\}^c$, we have: $\Delta(\alpha, \hat{\beta}_{k+1,n}) > 0$, almost surely. Conditionally on A_k , for all ϵ , the law of large number leads to $\Delta(\alpha, \hat{\beta}_{k+1,n}) \in [\Delta(\alpha, \beta_k) - \epsilon, \Delta(\alpha, \beta_k) + \epsilon]$, for n large enough; thus, there exists $m > 0$ and $N \in \mathbb{N}$ such that: $\forall n > N$, $\Delta(\alpha, \hat{\beta}_{k,n})/\Delta(\alpha, \hat{\beta}_{k+1,n}) > m$. As $w_{n_M} = o(\log(n_M))$, we have:

$$\mathbb{P} \left(B_k = \left\{ \frac{R_n(\alpha, \hat{\beta}_{k,n}, n_k)}{R_n(\alpha, \hat{\beta}_{k+1,n}, n_{k+1})} \xrightarrow{n \rightarrow +\infty} +\infty \right\} \middle| A_k \right) = 1 \text{ and } \mathbb{P}(B_k \cap A_k) = 0,$$

therefore $\mathbb{P}(A_k) = 0$ and the allocation strategy fulfill assumption 5.4.2 which implies the balanced behavior. In the case $\Delta(\alpha, \beta_a) > \Delta(\alpha, \beta_b) > 0$, $\Delta(\alpha, \hat{\beta}_{a,n})/\Delta(\alpha, \hat{\beta}_{b,n})$ converges almost surely to $\Delta(\alpha, \beta_a)/\Delta(\alpha, \beta_b) > 1$ and there exists $K > 1$ and $N \in \mathbb{N}$ such that, for all $n > N$, we have: $\Delta(\alpha, \hat{\beta}_{a,n})/\Delta(\alpha, \hat{\beta}_{b,n}) > K$. When $X_n = a$ and $n > N$, we then have $n_a < n_b(1/K)^{(1-w_n)/w_n}$. Since the method is balanced, there exists a non-decreasing sequence $(\varphi(n))_{n \in \mathbb{N}}$, such that for n large enough, we have: $\varphi(n) < n$ and $X_{\varphi(n)} = a$. Therefore, for n large enough,

$$n_a < n_b(1/K)^{(w_{\varphi(n)}^{-1}-1)} \text{ and } \frac{n_b}{n} \sim \frac{n_b}{n_a + n_b} \rightarrow 1 \text{ a.s.}$$

□

The further is the rate of toxicity at dose level a from α relative to that for b from α with respect to Δ , then the less the dose level a will be selected. Thus, the more the toxicity of dose a is far from α relatively to the toxicity of dose b , with respect to the divergence measure Δ , the less the dose a will be chosen asymptotically. The number of observations at doses outside the minimal set M^* is almost surely finite. Moreover, when the weight can vary, we are able to build a sequence $(w_n)_{n \in \mathbb{N}}$ such that the number of patients treated at the MTD is asymptotically equivalent to the number of patients enrolled in the trial. In the following section we outline some simulations with different allocation strategies.

5.6 Simulations

We evaluate the performance of an approach based on modelling the MTC, the SPMc, and an isotonic regression in conjunction with the allocation strategies described above. We compare these designs with other methods which have already been shown to have good performance, in particular the usual SPM, the CRM and the local BOIN methods. Cohorts are taken to be of size one meaning that we estimate the dose after each patient. We include the classical restriction that no skipping is allowed. The number of subject enrolled in the trial is 25. Six criteria are used for evaluation:

- PCS: percentage of correct selection of the MTD. For all methods, an isotonic regression is used on the data collected during the trial.
- PCSc: percentage of correct selection for the MTC.
- TR: percentage of patients treated at the dose.
- TR2: percentage of patients treated at the two consecutive doses either side of the threshold.
- Δ : mean of the difference between the toxic rates at which patients are treated and the rate of toxicity at the MTD.
- TM: empirical rate of toxicity observed during the trial.

We consider different structures in order to evaluate the interval or the minimal set. For SPMc, the prior $\tilde{\Pi}$ is uniform and the whole prior model can be summarized by the number c of pseudo-observations for each of the marginal distributions of the laws Λ_θ and the ratios of pseudo-toxicities observed for each marginal, the vector's family $(q^\theta)_{\theta \in D}$. This leads to the following formula:

$$\tilde{\lambda}_\theta(q) = \prod_{j=1}^m \tilde{\lambda}_\theta^j \propto \prod_{j=1}^{\theta-1} h(q_j, c, q_j^\theta) \mathbb{1}_{\tilde{B}}(q_j) \times \prod_{j=\theta}^m h(q_j, c, q_j^\theta) \mathbb{1}_{\tilde{A}}(q_j), \quad (5.19)$$

with $h(q, c, r) = q^{cr} \times (1-q)^{c(1-r)}$. When the support of the marginal distribution is \tilde{B} , the ratio is $1/7$ and when the support is \tilde{A} , the ratio is $3/8$: $j < \theta \Rightarrow q_j^\theta = 1/7$ and $j \geq \theta \Rightarrow q_j^\theta = 3/8$. Two cases are considered. The SPMc with a number of pseudo-observations equal to 40 and the SPMc-u where there are no pseudo-observations ($c = 0$). For the SPMc-u, the distributions on the intervals \tilde{B} and \tilde{A} are uniform. A simple alternative in order to evaluate the cut is to make use of an isotonic regression at each step of the trial (noted RI in table 5.1). In that method, a dose with no observation is supposed fully toxic. This does not result in being over conservative, as we might guess, because the next dose is selected by an allocation strategy. Nat (for natural) is strategy 1, Num (for number) is strategy 2 and KL (for Kullback-Leibler)

is allocation strategy 3. Strategy 4 (with $p=1/2$) is a usual SPM on the two doses of the estimated minimal set.

Local BOIN is parametrized as proposed in the original article (Liu and Yuan, 2015). In the case of the CRM, we make use of a two stage version (O'Quigley and Shen, 1996) based on some some lead-in rule until we observe the first toxicity and then we use maximum likelihood on the simple power model. The chosen skeleton is $u = (0.07, 0.13, 0.25, 0.40, 0.55, 0.74)$ with Normal law, $\mathcal{N}(0, 1.34^2)$, for the prior. The usual SPM fulfills assumption 5.3.1 with $\epsilon = 0$. The central ditributions Λ_θ^θ are then Dirac laws in α and this method is always balanced (see Definition 5.3.3(b)). The number of pseudo-observations for each marginal of the laws Λ_θ is 40 and the ratio of pseudo-toxicities for the interval B an A are $1/7$ and $3/8$, respectively, as for the SPMc.

Table 5.1: Model versus prior-model (without *Conservative Rules*).

Criteriaons		PCS	PCSc	TR	TR2	Δ	TM
BOIN		50.51	/	37.55	57.85	11.55	25.11
CRM		51.06	/	39.55	59.87	11.09	24.88
SPM		52.86	/	40.05	60.18	11.08	24.95
SPMc	KL	51.23	49.09	39.51	59.50	11.20	25.08
SPMc	Num	52.08	45.78	33.61	55.20	13.06	27.42
SPMc	Nat	53.91	50.26	40.09	60.78	11.31	25.44
SPMc-u	KL	51.68	49.03	39.40	59.44	11.24	25.24
SPMc-u	Nat	53.61	47.07	39.48	59.81	11.96	28.38
RI	KL	51.54	49.49	40.08	60.51	11.89	25.32
RI	spm	51.98	49.72	39.04	60.4	10.78	24.84

In table 5.1, all methods have been tested over 100 000 scenarios randomly generated by the very general pseudo-uniform algorithm (chapter 3, Simulations). Poor results are higlighted in red. The strategy Num, which only takes account of the numbers of observations at each dose on the estimated minimal set, obtains good results concerning the pragmatcal criterion, PCS. But this strategy has poorer results from the point of view of treatment allocation during the trial (the ethical criterion) which is due to its variability on the minimal set. This argument explains also the ethical results of the BOIN. The methods in blue obtain the best results. The first is the classical SPM and the second the SPM modelling on the minimal set with strategy 1. This strategy is natural because it uses only the model on the cut by selecting the dose with the higher mean of probability for either member of the minimal set. In section 5.4, we have shown the link between these two way of modelling and we recall

that, under minimal conditions, the SPMc with strategy 1 is exactly equivalent to a usual SPM (see 5.10). The SPM and the SPMc (Nat) converge in many generated trials. If the rate of toxicity associated with the MTD coincides with the value α , the SPM is then almost surely consistent, which means that the method will not experiment on other doses when the observations at the current doses are sufficiently close to the threshold. The SPMc-u (Nat) has a high percentage of observed toxicities compared to other methods which appears to be a consequence of its lack of structure. Isotonic regression working with SPM on the two sliding doses of the minimal set obtains good results. The results are also encouraging when the KL strategy is used, although the high value of criterion Δ when compared with other methods suggests higher variability so that, at the cost of being less rigid in allocation, we make some modest gain in recommendation accuracy. Such choices need to be discussed in advance by the investigating team and will depend on the priorities that the team feels should guide the study.

5.7 Conclusion

The central feature of the SPM is the direct modelling of the key parameter of interest, the MTD. For the SPMc models, the focus shifts and is on the minimal set, or equivalently the maximum tolerated cut, corresponding to the two consecutive doses on each side of the regression function's root. This might be viewed as a more statistically meaningful goal in the light of the impossibility theorem of Azriel et al. (2011). Furthermore, such a goal coincides with the majority of current trials where more than a single dose may be considered for closer study, in expansion cohorts for example or in preliminary Phase II studies. We show formally in this work that the two approaches to modelling are very similar and, in certain cases, will be fully equivalent. At the same time, the SPMc introduces a much more general perspective and one that can be readily calibrated to function with a wide variety of allocation strategies in order to obtain particular behavior as well as particular asymptotic properties, properties that will manifest themselves in some respect in the realistic small samples available to us. If we so wish we can use the methods immediately to consider more than two doses; we might also describe the interval in different ways, for example the lower level being greater than one threshold while the upper level remains less than a second threshold. Many combinations are possible and raise no additional methodological concerns.

5.8 Appendix: Comparison of two discrete priors

The couple (Ω, \mathcal{A}) denotes an abstract space endowed with its σ -field. The letter I denotes a finite set and $(X_k)_{k \in \mathbb{N}}$ is a sequence of I -valued and independent random variables. Let F be the space of function from I to $[0, 1]$. For an element $q \in F$, q_i denotes its value at $i \in I$. Let the probability space be $S = \{q \in F : \sum_{i \in I} q_i = 1\}$. A random variable X following the distribution q satisfies $\mathbb{P}\{X = i\} = q_i$, for $i \in I$. Let Λ_1 and Λ_2 be two probabilities on the Borel σ -field \mathcal{B} of S . Let S_1 and S_2 be the respective topological support of Λ_1 and Λ_2 . We would like to know the asymptotic behaviour of the ratio of the expected likelihood under Λ_1 to the one under Λ_2 . We define the operator r as follows:

$$r(\Lambda_1, \Lambda_2, n) = \frac{\int \prod_{k=1}^n q_{X_k} \Lambda_1(dq)}{\int \prod_{k=1}^n q_{X_k} \Lambda_2(dq)} = \frac{\int \prod_{i \in I} q_i^{n_i} \Lambda_1(dq)}{\int \prod_{i \in I} q_i^{n_i} \Lambda_2(dq)},$$

where $n_i = \sum_{k=1}^n \mathbf{1}_{\{X_k=i\}}$, for $i \in I$. We assume that $(X_k)_{k \in \mathbb{N}}$ is identically distributed under the true probability β . The convergence of $r(\Lambda_1, \Lambda_2, n)$ depends mainly on the localization of β relatively to the support S_1 and S_2 . For this we make use of the usual concept of entropy. The entropy of q relative to p is $H(q|p) = -\sum_{i \in I} p_i \log q_i$, with $\log 0 = -\infty$ and $0 \times -\infty = 0$. We suppose that β is closer to S_1 than S_2 in term of entropy.

Assumption 5.8.1. *Let V be a subspace of S_2 satisfying $\Lambda_2(V) > 0$. There exists $\delta > 0$ such that:*

$$\inf_{q \in S_1} H(q|\beta) - \sup_{q \in V} H(q|\beta) > 4\delta.$$

This leads to a simple characterisation of the behavior of $r(\Lambda_1, \Lambda_2, n)$.

Lemma 5.8.1. *Under Assumption 5.8.1, we have: $r(\Lambda_1, \Lambda_2, n) \xrightarrow{n \rightarrow +\infty} 0$.*

Proof. Let the empirical probability be $\hat{\beta}_n = (\hat{\beta}_{n,i})_{i \in I}$ where $\hat{\beta}_{n,i} = n_i/n$. By discarding a \mathbb{P}_β -null set, the law of large number implies that $\lim_{n \rightarrow +\infty} \hat{\beta}_{n,i}(\omega) = \beta_i$, for all $i \in I$ and for all $\omega \in \Omega$. The main fact to establish is

$$\inf_{q \in S_1} H(q|\hat{\beta}_n) - \sup_{q \in V} H(q|\hat{\beta}_n) > \delta, \quad (5.20)$$

for n large enough. Let $I_+ = \{i \in I : \beta_i > 0\}$, and $S_+ = \{(q_i)_{i \in I_+} : \sum_{i \in I_+} q_i \geq 1\}$, the projection of S on the vector space indexed by I_+ . In the same way, we defined $S_{1,+}$ and V_+ such that

$$S_{1,+} = \{(q'_i)_{i \in I_+} : \exists q \in S_1, \forall i \in I_+, q'_i = q_i\},$$

and

$$V_+ = \{(q'_i)_{i \in I_+} : \exists q \in V, \forall i \in I_+, q'_i = q_i\}.$$

We begin by proving the following inequality for n large enough

$$\left| \sup_{q \in V} H(q|\beta) - \sup_{q \in V} H(q|\hat{\beta}_n) \right| < \delta. \quad (5.21)$$

Assumption 5.8.1 implies $\sup_{q \in V} H(q|\beta) < \infty$, hence there exists $\kappa > 0$ such that $V_+ \subset T = [\kappa, 1]^{\#I_+}$. $\overline{V_+}$ denotes the closure of V_+ . Let $H_{I_+}(q|p) = -\sum_{i \in I_+} p_i \log q_i$. For all $p \in S_+$, the function $H_{I_+}(\cdot|p)$ defined over T is Lipschitz continuous for the constant $K_p = (\sup_{i \in I_+} p_i) / \log(\kappa)$. As the sequence $\hat{\beta}_n$ converge to β , there exists a Lipschitz constant K_0 available for the whole sequence $(H_{I_+}(\cdot|\hat{\beta}_n))_{n \in \mathbb{N}}$ and its limit $H_{I_+}(\cdot|\beta)$ over $\overline{V_+}$. As $\overline{V_+}$ is a compact set, for all $\varepsilon > 0$, there exists a finite family $(p_j)_{1 \leq j \leq J}$ satisfying the following property: $\forall q \in V_+, \exists j \in \llbracket 1, J \rrbracket$ such that $d(q, p_j) < \varepsilon$. For all $q \in S_+$, we can find $k \in \llbracket 1, J \rrbracket$, such that

$$\begin{aligned} \left| H_{I_+}(q|\beta) - H_{I_+}(q|\hat{\beta}_n) \right| &\leq \left| H_{I_+}(q|\beta) - H_{I_+}(p_j|\beta) \right| + \left| H_{I_+}(p_j|\beta) - H_{I_+}(p_j|\hat{\beta}_n) \right| \\ &\quad + \left| H_{I_+}(p_j|\hat{\beta}_n) - H_{I_+}(q|\hat{\beta}_n) \right|, \end{aligned}$$

where the second term is bounded by $(2K + 1)\varepsilon$, for n large enough. Thus, $(H(\cdot|\hat{\beta}_n))_{n \in \mathbb{N}}$ converges uniformly to $H(\cdot|\beta)$ and, for n large enough, we obtain the inequality (5.21). We continue by establishing

$$\left| \inf_{q \in S_1} H(q|\beta) - \inf_{q \in S_1} H(q|\hat{\beta}_n) \right| < \delta, \quad (5.22)$$

for n large enough and when $(\inf_{q \in S_1} H(q|\hat{\beta}_n))_{n \in \mathbb{N}}$ is bounded. Otherwise the inequality (5.20) is immediate. As $(\hat{\beta}_n)_{n \in \mathbb{N}}$ converge to β , there exists κ' such that $S_{1,+} \subset [\kappa', 1]^{\#I_+}$. Thus, we prove (5.22) by the same argument used for (5.21). Assumption 5.8.1 together with (5.21) and (5.22) lead to inequality (5.20). From (5.20) it follows that

$$\int_V \prod_{i \in I} q_i^{n_i} \Lambda_2(dq) / \Lambda_2(V) > \int_V \prod_{i \in I} q_i^{n_i} \Lambda_1(dq) \times \exp(n\delta),$$

and

$$r(\Lambda_1, \Lambda_2, n) \leq \frac{\int \prod_{i \in I} q_i^{n_i} \Lambda_1(dq)}{\int_V \prod_{i \in I} q_i^{n_i} \Lambda_2(dq)} < \exp(-n\delta) \Lambda_2(V),$$

which completes the proof. □

Chapter 6

The Case of Partial Ordering

6.1 Introduction

This study puts itself in the context of dual-agent Phase I clinical trials. The place of these trials in drugs development is more and more important because of the increasing use of drugs combination in therapies. The practical benefits of drugs combination are numerous : several modes of action can be combined or side effects of a drug can be lightened by another one for example. The aim of Phase I trials, typically in oncology, is to find the one or several drug combinations having the toxicity probability the closest to a threshold α , fixed in advance by clinicians. In general, this threshold is around 25%. We name these combinations the maximum tolerated dose (MTD).

There exist some publications about algorithmic designs used to identify the MTD within a discrete set of doses for Phase I trials Storer (1989); Skolnik et al. (2008) for single-agent trials, Huang et al. (2007) for dual-agent trials. These designs have no modeling, and the escalation, de-escalation rules are determined as a function of some set of the most recent observations, that is why they are used so often. Furthermore they have a Markov property, sometimes referred to in this context as a lack-of-memory property. However the sequence of dose selected by such models has a high variance and does not converge. Multiple model-based designs have been suggested, but they are hardly used in practice. This is even more the case with dual-agent trials because of the complexity and lack of interpretability of the proposed models, but also because of the difficulty to implement these designs with software. Existing model-based designs can be sorted in two categories. The first one is the class of parametric models. For dual-agent trials, we can cite Wages et al. (2011) with the partial ordering continual reassessment method (CRM), Wang and Ivanova (2005) and Braun and Jia (2013) with extensions of the CRM, Yin and Yuan (2009a) and Yin and Yuan (2009b) with copula models, Thall et al. (2003) with a six-parameter logistic-type model and Braun and Wang (2010) with the use of beta distributions and log-linear models. The second one is the class of non-parametric models. For dual-agent trials, we can cite Mander and Sweeting (2015) with the product of independent beta probabilities escalation design (PIPE).

In Phase I trials, we are interested in several criteria. Here we consider three of them. Two of them are well explained by Azriel et al. (2011). The first one is called the treatment principle : we want to treat patients at the MTD as often as possible. The second one is called the experimentation principle : we want to obtain a good estimate for the MTD at the end of the study. We know (Azriel et al., 2011) that if the first principle is verified, then the second cannot be, in the sense that we do not have almost sure convergence to the MTD. So there is a trade-off to make, but ideally, we want a design who outperforms the others for these two criteria. The last one is called the coherence, was introduced by (Cheung, 2005) and can be explain that way : if we observe a dose-limiting toxicity (DLT) for the n -th patient, we do not want the design to recommend a higher dose to the $(n + 1)$ -th patient ; in the same way, if we do not observe a DLT for the n -th patient, we do not want the design to recommend a lower dose to the $(n + 1)$ -th patient. This last criterion is essential to be confident in the design.

In this chapter, we introduce a new model-based design for dual-agent trials. This is an extension of SPM for single-agent trials described in the preceding chapters.

6.2 Context and Notations

In this paper, we are interested in estimating the root of a dose-toxicity regression function as observations are accumulated sequentially. Let (i, j) represents the combination of the i -th dose of a drug A1 and the j -th dose of a drug A2, where $i \in \{1, \dots, I\}$ and $j \in \{1, \dots, J\}$. Let D be the set of all the combinations of doses: $D = \{1, \dots, I\} \times \{1, \dots, J\}$. We denote the observations by the sequence $(X_n, Y_n)_{n \in \mathbb{N}}$. At step n , corresponding to the n -th patient enrolled in the trial, the variable X_n is the dose selected among the IJ available doses; and the variable Y_n is the observed binary response at this dose taking values $\{0, 1\}$ (1 for a DLT and 0 otherwise). The conditional distribution of Y_n given $X_n = (i, j)$ follows a Bernoulli distribution with parameter $\beta_{(i,j)}$. This can be written as follow.

Assumption 6.2.1. $\forall n \in \mathbb{N}, \quad \mathbb{P}(Y_n = 1 | X_n = (i, j)) = \beta_{(i,j)}$.

For a threshold α , fixed by clinicians, we want to find the dose with a probability of toxicity the closest to α . This is achieved by estimating the root of the regression function. This dose is called the maximum tolerated dose (MTD) and we denote it d^* . It is assumed that the ranges $\{1, \dots, I\}$ and $\{1, \dots, J\}$ are ordered in terms of the probability of toxic response. This leads naturally to an order on D . 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 6.2.2. $(i, j) < (r, s) \Rightarrow \beta_{(i,j)} < \beta_{(r,s)}$.

For all $X \in D$, this partial ordering partitions D . We denote $A_X = \{(k, l) \in D : (k, l) > X\}$ the set of doses associated with toxicity higher than β_X and $B_X = \{(k, l) \in D : (k, l) < X\}$ the set of doses associated with toxicity lower than β_X . The set $C_X = D \setminus (A_X \cup B_X)$ contains the doses which are not ranked with X . The above sets are illustrated in Figure 6.1. We suppose

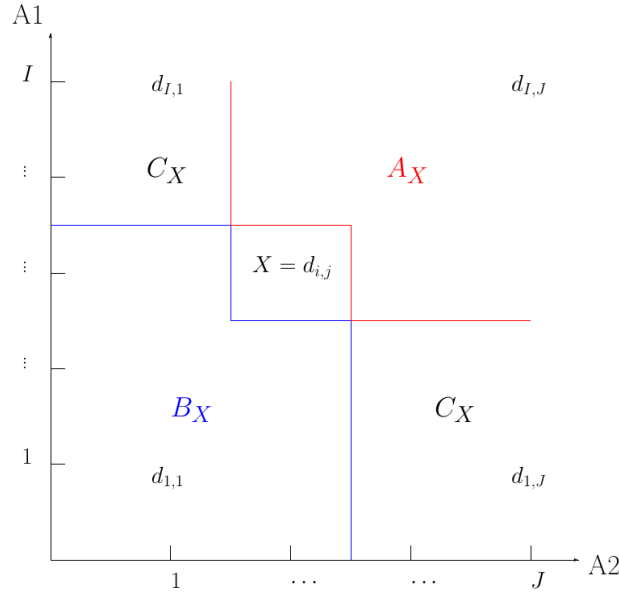


Figure 6.1: Illustration of the sets A_X , B_X and C_X .

that all the sample $(X_1^n, Y_1^n) = ((X_1, Y_1), \dots, (X_n, Y_n))$ can be used for the recommendation of the dose X_{n+1} . More precisely, for all the methods exposed in that article, X_{n+1} is completely determined by the history (X_1^n, Y_1^n) .

Assumption 6.2.3. *The current estimator of the method \mathcal{M} 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.*

We now introduce a generalization, for the case of partial ordering, of the coherence principle introduced by Cheung (2005) for the case of total ordering.

Definition 6.2.1 (Partial ordering coherence). *A method \mathcal{M} is coherent if its current estimator satisfies:*

$$(X_n, Y_n) = (d, 0) \Rightarrow \mathcal{M}(X_1^n, Y_1^n) \in D \setminus B_d \text{ and } (X_n, Y_n) = (d, 1) \Rightarrow \mathcal{M}(X_1^n, Y_1^n) \in D \setminus A_d.$$

First condition is the coherence in de-escalation and second the coherence in escalation.

This latter definition matches the one of Cheung (2005) in case of a single-agent trial, and so is a reasonable extension of this criteria in case of partial ordering.

6.3 Modelling from an MTD point of view

6.3.1 General setting

The po-SPM is the generalisation of SPM introduced in the preceding chapters. Let $F = [0, 1]^D$ be the set of curves (or functions) which go from D to $[0, 1]$. An element of F is denoted by

$q = (q_{(i,j)})$ where $i \in \{1, \dots, I\}$ and $j \in \{1, \dots, J\}$. Then, $q_{(i,j)}$ is a probability of toxicity for the dose (i, j) . In particular, $\beta \in F$. In the context of Bayesian hierarchical framework, let θ be the random variable of interest taking values in D . We now partition F according to θ :

$$F = \bigcup_{\theta \in D} F_\theta, \text{ where } F_\theta = \{q \in F / \forall (i, j) \in D, |q_\theta - \alpha| \leq |q_{(i,j)} - \alpha|\}.$$

We can see that F is partitioned in IJ classes F_θ , where every curve of a same class have the same MTD, by definition. For a fixed $q \in F$, we can make explicit the probability of the history (X_1^n, Y_1^n) :

$$\begin{aligned} \mathbb{P}_q(X_1^n, Y_1^n) &= \mathbb{P}_q(Y_1|X_1) \times \prod_{p=1}^n \left[\mathbb{P}_q(Y_p|X_p) \mathbb{P}_q(X_p|X_1^{p-1}, Y_1^{p-1}) \right] \\ &= \prod_{p=1}^n (q_{X_p})^{Y_p} (1 - q_{X_p})^{1-Y_p} = \prod_{(i,j) \in D} q_{(i,j)}^{n_{(i,j)}^1} (1 - q_{(i,j)})^{n_{(i,j)}^0}, \end{aligned}$$

where $n_{(i,j)}^1$ represents the number of DLTs occurred at dose (i, j) , and $n_{(i,j)} = n_{(i,j)}^1 + n_{(i,j)}^0$ the number of patients treated at dose (i, j) . We endow F with a probability measure $\Lambda \otimes \Pi$, where Π is a measure on D , and the topological support of the measure $\Lambda(\cdot|\theta) = \Lambda_\theta(\cdot)$ is included in F_θ . We can now express the posterior distribution of θ given the history (X_1^n, Y_1^n) :

$$\Pi_n(\theta) = \Pi(\theta|X_1^n, Y_1^n) = \frac{\int \prod_{(i,j) \in D} q_{(i,j)}^{n_{(i,j)}^1} (1 - q_{(i,j)})^{n_{(i,j)}^0} \Lambda_\theta(dq) \Pi(\theta)}{\sum_{\theta \in D} \int \prod_{(i,j) \in D} q_{(i,j)}^{n_{(i,j)}^1} (1 - q_{(i,j)})^{n_{(i,j)}^0} \Lambda_\theta(dq) \Pi(\theta)}. \quad (6.1)$$

Let us assume that Λ_θ is absolutely continuous with respect to a measure ν with density λ_θ . Then the posterior distribution of F_θ given (X_1^n, Y_1^n) is

$$\lambda_{\theta,n}(q) = \frac{\prod_{(i,j) \in D} q_{(i,j)}^{n_{(i,j)}^1} (1 - q_{(i,j)})^{n_{(i,j)}^0} \lambda_\theta(q)}{\int \prod_{(i,j) \in D} q_{(i,j)}^{n_{(i,j)}^1} (1 - q_{(i,j)})^{n_{(i,j)}^0} \lambda_\theta(q) \nu(dq)}. \quad (6.2)$$

Using Equation (6.1) and Equation (6.2) together, we obtain:

$$\Pi_n(\theta) = \frac{\left[\int q_{X_n}^{Y_n} (1 - q_{X_n})^{1-Y_n} \lambda_{\theta, n-1}(q) \nu(dq) \right] \Pi_{n-1}(\theta)}{\sum_{\theta \in D} \left[\int q_{X_n}^{Y_n} (1 - q_{X_n})^{1-Y_n} \lambda_{\theta, n-1}(q) \nu(dq) \right] \Pi_{n-1}(\theta)}.$$

The distribution Π is updated by weighting according to the expected likelihood. At step n , we integrate the likelihood according to the posterior $\Lambda_{\theta, n-1}$. Thus, the amount of information of a fixed observation (X_n, Y_n) varies according to the whole history (X_1^{n-1}, Y_1^{n-1}) . The family of distributions $(\Lambda_\theta)_{\theta \in D}$ plays a predictive and adaptive model-like role and will be called prior model. We then choose the estimator of the MTD (or current estimator) such as

$$\hat{\theta}_n = \operatorname{argmax}_{\theta \in D} \Pi_n(\theta). \quad (6.3)$$

Note that we can also obtain estimators for every probability $\beta_{(i,j)}$:

$$\hat{\beta}_{(i,j)}^{(n)} = \mathbb{E}_{(\Lambda \otimes \Pi)_n} [q_{(i,j)}] = \sum_{q \in D} \left[\int q_{(i,j)} \Lambda_{\theta, n}(dq_{(i,j)}) \right] \Pi_n(\theta).$$

Now let S_θ and S_θ^d be the topological support of Λ_θ and its marginal Λ_θ^d , with $d \in D$. The interval of toxicity $[0, 1]$ is partitioned into three sets: $I = [\alpha - \varepsilon_1, \alpha - \varepsilon_2]$, $B = [0, \alpha - \varepsilon_1]$ and $A = [\alpha + \varepsilon_2, 1]$, with $\varepsilon_1, \varepsilon_2 \in [0, \max(\alpha, 1 - \alpha)]$. For all doses θ and d , the marginal support S_θ^d satisfies:

$$\left\{ \begin{array}{l} d \in A_\theta \Rightarrow S_\theta^d \subset A \\ d \in B_\theta \Rightarrow S_\theta^d \subset B \\ d \in C_\theta \Rightarrow S_\theta^d \subset [0, 1] \\ d = \theta \Rightarrow S_\theta^d \subset I \end{array} \right.$$

6.3.2 A simple and coherent prior model

We are now interested in the prior distribution inside the class (Λ) and call the family $(\Lambda_\theta)_{\theta \in D}$ the prior-model because we use it to make inference. In a single class F_θ , the Bernoulli's parameters at each dose $q_{(i,j)}$ are considered independent, which is summarised by the following assumption.

Assumption 6.3.1. Λ_θ is a product of unidimensionnal distribution, i.e.

$$\Lambda_\theta(dq) = \prod_{i=1}^I \prod_{j=1}^J \Lambda_\theta^{(i,j)}(dq_{(i,j)}).$$

When there is no confusion, we write $\Lambda_\theta(dq_{(i,j)})$ for $\Lambda_\theta^{(i,j)}(dq_{(i,j)})$. We then state a stochastic partial ordering assumption on the prior-model, which involves the coherence principle. An example of such a prior model is given below.

Assumption 6.3.2. *Let d and d' be two doses such that $d < d'$. For all marginal $(i, j) \in D$, the posterior $\Lambda_{d,n}^{(i,j)}$ is stochastically greater than $\Lambda_{d',n}^{(i,j)}$:*

$$\Lambda_{d,n}^{(i,j)}([0, x]) \leq \Lambda_{d',n}^{(i,j)}([0, x]) \quad \forall x \in [0, 1].$$

Theorem 6.3.1. *Under Assumptions 6.3.1 and 6.3.2, if $(\Lambda_\theta)_{\theta \in D}$ is the single prior model used, then the po-SPM is coherent in the sense of Definition 6.2.1.*

Proof. We only need to deal with the following cases:

- Progression in drug A1: Let $s \in \{1, \dots, J\}$ and $(r, t) \in \{1, \dots, I\}^2$ such as $r < t$,
- Progression in drug A2: Let $s \in \{1, \dots, I\}$ and $(r, t) \in \{1, \dots, J\}^2$ such as $r < t$.

Suppose that $Y_{n+1} = 1$. The case $Y_{n+1} = 0$ can be solved identically. By construction, we have

$$\Pi_{n+1}(\theta) \propto \left[\int q \Lambda_{\theta,n}(dq) \right] \Pi_n(\theta).$$

Furthermore,

$$\begin{aligned} \int q_{(r,s)} \Lambda_{\theta,n}(dq) &= \int \left[\int \mathbf{1}_{\{0 \leq x \leq q_{(r,s)}\}} \mu(dx) \right] \Lambda_{\theta,n}(dq_{(r,s)}) \\ &= \int \Lambda_{\theta,n}^{(r,s)}(\cdot | x, 1) \mu(dx). \end{aligned}$$

If $\hat{\theta}_n = (r, s)$, then for all $\theta \in D$, $\Pi_n(\theta) \leq \Pi_n(r, s)$. Let $t > r$. According to Assumption 6.3.2, we know that $\Lambda_{(r,s),n}^{(r,s)}(\cdot)$ is stochastically greater than $\Lambda_{(t,s),n}^{(r,s)}(\cdot)$, *i.e.*

$$\int \Lambda_{(r,s),n}^{(r,s)}(\cdot | x, 1) \mu(dx) \geq \int \Lambda_{(t,s),n}^{(r,s)}(\cdot | x, 1) \mu(dx),$$

that is

$$\int q_{(r,s)} \Lambda_{(r,s),n}(dq) \geq \int q_{(r,s)} \Lambda_{(t,s),n}(dq).$$

Finally $\Pi_{n+1}(r, s) \geq \Pi_{n+1}(t, s)$. With the same arguments, we can show that $\Pi_{n+1}(s, r) \geq \Pi_{n+1}(t, r)$ for $t > s$. So the SPM is coherent by definition of $\hat{\theta}_n$ (Equation (6.3)). \square

In the following example, the prior model satisfies the assumptions 6.3.1 and 6.3.2 and the conjugacy for the likelihood.

Example 6.3.1. The prior model is defined by a triplet $((S_\theta)_{\theta \in D}, (q^\theta)_{\theta \in D}, c)$. The sets S_θ are the topological supports of the distributions Λ_θ and for all $(i, j) \in D$, the marginal supports, $S_\theta^{(i,j)}$, fulfill the following assumption.

Assumption 6.3.3. The topological supports satisfy:

$$\left\{ \begin{array}{ll} (i, j) \in A_\theta & \Rightarrow S_\theta^{(i,j)} = A \\ (i, j) \in B_\theta & \Rightarrow S_\theta^{(i,j)} = B \\ (i, j) \in C_\theta & \Rightarrow S_\theta^{(i,j)} = [0, 1] \\ (i, j) = \theta & \Rightarrow S_\theta^{(i,j)} = I, \end{array} \right.$$

with θ and (i, j) in D .

The vectors $q^\theta \in [0, 1]^m$ are the modes of the distributions Λ_θ . Thus $q_{(i,j)}^\theta$ takes the maximum value of the density function $\lambda_\theta^{(i,j)}$:

$$q^\theta = \arg \max_{q \in F_\theta} \Lambda_\theta(q) .$$

These modes are ranked according to the product order: if $(i, j) < (r, s)$ then $q_{(i,j)}^\theta < q_{(r,s)}^\theta$. The positive vector $c = (c_1, c_2)$ is a dispersion parameter of the distributions ; c_1 corresponds to a number of pseudo-patients observed at the doses ranked with θ , i.e. the doses in the set $D \setminus C_\theta$; and c_2 corresponds to a number of pseudo-patients observed at the doses non-ranked with θ , i.e. the doses in the set C_θ . Uniform priors on the topological supports are updated. We denote l the real function such as for all $x, y, z \in [0, 1]$, $l(q, c, r) = q^{cr}(1 - q)^{c(1-r)}$. Then, for all $\theta \in D$ and all $q \in F_\theta$, we have the following result.

$$\lambda_\theta^{(i,j)}(q_{(i,j)}) \propto \left\{ \begin{array}{ll} l(q_{(i,j)}, c_1, q_{(i,j)}^\theta) \mathbb{1}_{\{q_{(i,j)} \in A\}} & \text{if } (i, j) \in A_\theta, \\ l(q_{(i,j)}, c_1, q_{(i,j)}^\theta) \mathbb{1}_{\{q_{(i,j)} \in B\}} & \text{if } (i, j) \in B_\theta, \\ l(q_{(i,j)}, c_2, q_{(i,j)}^\theta) \mathbb{1}_{\{q_{(i,j)} \in [0,1]\}} & \text{if } (i, j) \in C_\theta, \\ l(q_{(i,j)}, c_1, q_{(i,j)}^\theta) \mathbb{1}_{\{q_{(i,j)} \in I\}} & \text{if } (i, j) = \theta. \end{array} \right.$$

Thus, Λ_θ is a product of beta priors on $[0, 1]$ and incomplete beta priors on A , B and I .

In that case, the marginals of the prior model are absolutely continuous according to the Lebesgue measure. This last assumption is based on the regularity of the prior model.

Assumption 6.3.4. The following conditions are valid except when Λ_θ^θ is a Dirac measure.

(a) For all $(i, j) \in D$, the marginal distribution $\Lambda_\theta^{(i,j)}$ is absolutely continuous with respect to the Lebesgue measure and $\lambda_\theta^{(i,j)}$ denotes its density function.

(b) There exist two numbers s and S in \mathbb{R}_+ , such that, for all θ and (i, j) in D , we have:

$$\forall q_{(i,j)} \in S_\theta^j, s < \lambda_\theta^{(i,j)}(q_{(i,j)}) < S .$$

The second point is only useful for the sake of the demonstration when some toxicities $\beta_{(i,j)}$ are equal to 0 or 1.

6.3.3 Asymptotical results and perspective

In the Bayesian paradigm, the choice of a topological support for the prior model determines consistency. It is therefore a central step for po-SPM and we will design it according to the goal of our study. In the previous example, the support of marginal Λ_θ^θ is $I = [\alpha - \varepsilon_1, \alpha + \varepsilon_2]$. In that case, observing at dose θ leads eventually to recommend θ if and only if the toxicity β_θ is included in I . Such an assumption on the scenario is necessary for the almost sure convergence to the MTD (Azriel et al., 2011, Theorem 1). In case of a range of dose totally ordered, this leads to consider the ε -sensitivity behaviour of a method (Cheung, 2011). The definition does not change in case of partial ordering.

Definition 6.3.1. *Let $\varepsilon \geq 0$ and $I = [\alpha - \varepsilon; \alpha + \varepsilon]$. We consider the set $\mathcal{E}(I, \beta)$ of the collection of doses associated with a toxicity belonging to I_ε , i.e. $\mathcal{E}(I, \beta) = \{j \in D : \beta_j \in I\}$. A method, \mathcal{M} , is called ε -sensitive, if for all β such that $\mathcal{E}(I, \beta) \neq \emptyset$, we have:*

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

The ε -sensitivity corresponds to a strong consistency to one dose associated with toxicity close to the threshold α . The almost sure convergence to the MTD is obtained if the MTD is the single dose in $\mathcal{E}(I, \beta)$. This involves to choose a little interval I which could contain no toxicity associated with a dose in our range. For this reason, the complementary behaviour of a method, called ε -balanced, is introduced in case of total ordering (see chapter 3). We extend this definition to the case of partial ordering. Let $\delta(., .)$ be the euclidean distance. A sequence $(X_n)_{n \in \mathbb{N}}$ converges to a set B , denoted by $X_n \xrightarrow{S} B$, if :

$$\sup_{x \in B} \left(\liminf_{n \rightarrow +\infty} \delta(X_n, x) \right) = 0.$$

The key notion is the minimal set of doses on which we need to have observations in order to determine almost surely the MTD. Let L and U be the partition of D into the sets of doses associated with toxicities respectively lower and upper than α , i.e. $L = \{d \in D : \beta_d \leq \alpha\}$ and $U = \{d \in D : \beta_d \geq \alpha\}$. The minimal set M_D satisfies $M_D = M_L \cup M_U$, where $M_L = \{d \in L : A_d \cap L = \emptyset\}$ and $M_U = \{d \in L : B_d \cap U = \emptyset\}$. In case of total ordering, M_D is equal to $\{a, b\}$, where b (below) and a (above) are the two consecutive doses associated to toxicities either side of the target α . This allow us to introduce the following general definition.

Definition 6.3.2. *Let D be a range of doses. A method, \mathcal{M} , is called ε -balanced, if for all β such that $\mathcal{E}(I, \beta) = \emptyset$, we have:*

$$\mathcal{M}(X_1^n, Y_1^n) \xrightarrow{S} M_D, \text{ a.s.}$$

Note that if $I = \{\alpha\}$, that is if $\varepsilon = 0$, this behaviour is referred as *balanced*.

This means that if no doses are associated with a toxicity in I , the doses recommended infinitely often by the current estimator are all the doses in M_D and only these doses.

According to the prior model parametrization, the po-SPM has the following asymptotic properties.

Theorem 6.3.2. *Under assumptions 6.3.1, 6.3.3 and 6.3.4 the po-SPM is ε -sensitive and ε -balanced.*

The ε -balanced property allows the po-SPM to concentrate the observations on the minimal set M_D which, in a asymptotical point of view, leads to determine almost surely the MTD. Indeed, the MTD is included in the minimal set for which each dose is infinitely observed. However, this behaviour is only asymptotical and, at a finished rank, the simulations seem often to account a convergence to a single dose identified as the MTD. In the usual case of one dimensional space of drugs, some questions are recently raised about the pragmatism aim of an early phase dose finding study in relation with the expansion cohort. As seen in the preceding chapter, maintaining more than one dose in the study does not lead to poorer results. From this point of view and under certain conditions of proximity between the toxicities and the threshold, the minimal set $\{a, b\}$ has the property to be the best candidate. This new paradigm is particularly relevant in the two dimensional context in which the goal of finding the MTD is less obvious. Indeed, the dose/toxicity and dose/efficiency relations can both verify the partial ordering assumption. In this case, there exist some configuration where a dose less toxic than the MTD is more efficient. Such a scenario is impossible in case of total ordering for the two relation considered previously. Moreover, in the partial order case, and from the point of view of extrapolation, the observations are usually less informative than in the total order case. Thus, in our situation, the MTD which is not necessarily the most adapted response for the clinicians is even more difficult to find. That is why it seems important to have the capacity to recommend to the phase II a sensitive set of doses which contains the MTD with high probability and other doses close enough to it.

6.4 Modelling from an MTC point of view

In the partial ordering case, the minimal set is less intuitive than in the case of total ordering. However, it is possible to consider the minimal set on each line $l_j = \{(i, j), i \in \{1, \dots, I\}\}$ and column $c_i = \{(i, j), j \in \{1, \dots, J\}\}$. This leads to the notion of Maximum Tolerated Contour introduced by Mander and Sweeting (2015) and for which we give the following formal definition:

$$MTC = \bigcup_{(i,j) \in [1,I] \times [1,J]} (M_{c_i} \cup M_{l_j}). \quad (6.4)$$

From a treatment efficiency perspective, clinicians could prefer to find the MTC than the MTD. Note that M_D is also the minimal set of doses on which we need to have observations in

order to determine with certainty the MTC. This provides a particular interest to the balanced property as it unifies different points of view.

The class of Partial Ordering and Semi-Parametric Methods on the Contour (po-SPMc) are inspired by the PIPE model introduced by Mander and Sweeting (2015) and is an extension of unidimensional semi-parametric methods on cut (chapter 5) to the partial ordering case. The modelization includes the PIPE model as a particular case. Now let \mathcal{K} be the set of the available contours which are not in contradiction with the Assumption 6.2.2. The letter \mathcal{C} denotes an element of \mathcal{K} and \mathcal{C}^* is the true MTC. We first define the MTC, using the general notion of contour, as the set of couple of doses which are the minimal set on each lines and columns of D , see Equation (6.4). Each contour can also be seen as a polygonal chain going

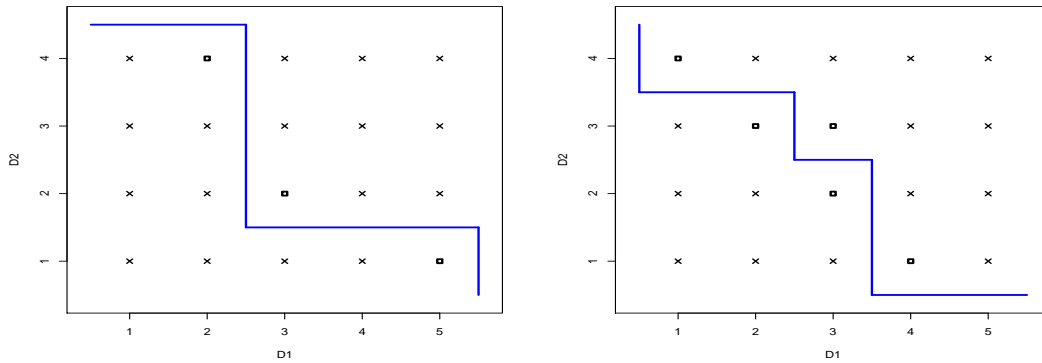


Figure 6.2: **Two contours**, the symbol \square represents the doses of minimal set and \times the other doses.

from $(0.5, J + 0.5)$ to $(I + 0.5, 0.5)$ by step of size 1 in abscissa and -1 in ordinate. Examples of two different contours can be found in Figure 6.2. This second definition allows us to find easily the total number of contours: $\#\mathcal{K} = \binom{I+J}{I}$. For each contour \mathcal{C} , we introduce three sets of dose. $L(\mathcal{C}) = \{d \in D : \exists d' \in \mathcal{C}, d' > d\}$ and $U(\mathcal{C}) = \{d \in D : \exists d' \in \mathcal{C}, d' < d\}$ are respectively the set of doses lower and upper than \mathcal{C} , and $M(\mathcal{C}) = \{d \in L(\mathcal{C}) : A_d \cap L(\mathcal{C}) = \emptyset\} \cup \{d \in U(\mathcal{C}) : B_d \cap U(\mathcal{C}) = \emptyset\}$ denotes the minimal set corresponding to the contour \mathcal{C} . In particular, $M(\mathcal{C}^*) = M_D$. Note that finding the contour is equivalent to find the minimal set. We endow the set of all scenario, F , with a probability measure $\tilde{\Lambda} \otimes \tilde{\Pi}$, where $\tilde{\Pi}$ is a measure on \mathcal{K} . The distribution $\tilde{\Lambda}_{\mathcal{C}} = \tilde{\Lambda}(\cdot | \mathcal{C})$ and its support $S_{\mathcal{C}}$ satisfy the following hypotheses.

Assumption 6.4.1. (a) $\tilde{\Lambda}_{\mathcal{C}}$ is a product of unidimensionnal distribution :

$$\tilde{\Lambda}_{\mathcal{C}}(dq) = \prod_{i=1}^I \prod_{j=1}^J \tilde{\Lambda}_{\mathcal{C}}^{(i,j)}(dq_{(i,j)}).$$

Let $S_{\mathcal{C}}^d$ be the marginal support for dose d . We have

$$\begin{cases} d \in L(\mathcal{C}) & \Rightarrow S_{\mathcal{C}}^d = [0, \alpha] \\ d \in U(\mathcal{C}) & \Rightarrow S_{\mathcal{C}}^d = [\alpha, 1]. \end{cases}$$

(b) For all $d \in D$, the marginal distribution $\tilde{\Lambda}_{\mathcal{C}}^d$ can have an atom in α and is absolutely continuous with respect to the Lebesgue measure on the rest of its support. $\tilde{\lambda}_{\mathcal{C}}^d$ denotes its density function. There exist two numbers s and S in \mathbb{R}_+^* , such that, for all \mathcal{C} and d ,

$$\forall q_d \in S_{\mathcal{C}}^d, s < \tilde{\lambda}_{\mathcal{C}}^d(q_d) < S.$$

We obtain the posterior distribution $\tilde{\Pi}_n$ of \mathcal{C} in the same way as Π (Equation (6.1)) and a natural estimator of the MTC is $\tilde{\mathcal{C}}_n = \arg \max_{\mathcal{C} \in \mathcal{K}} \tilde{\Pi}_n(\mathcal{C})$. From an asymptotical point of view, Assumption 6.4.1 (b) may be deleted when the marginals of $\tilde{\Lambda}$ satisfy

$$\forall \mathcal{C}, \mathcal{C}' \in \mathcal{K}, \forall d \in (U(\mathcal{C}) \cap U(\mathcal{C}')) \cup (L(\mathcal{C}) \cap L(\mathcal{C}')), \tilde{\Lambda}_{\mathcal{C}}^d = \tilde{\Lambda}_{\mathcal{C}'}^d. \quad (6.5)$$

This assumption is fulfilled by the PIPE method. It is not sure that such a restriction on the prior model and the localization of the weights leads to weaker performances. Note that, in the modelization of PIPE, the prior model $\tilde{\Lambda}$ and the distribution $\tilde{\Pi}$ are never explicitly mentioned and are, in some way, indifferentiated. The whole parameter of the po-SPMc is summarized by the family $(a_d, b_d)_{d \in D}$, where the couple (a_d, b_d) is a fictive amount of observations on the dose d . We integrate this piece of information into a uniform prior on the parametric set of toxicity at dose d . When we include the PIPE method in the po-SPMc settings, this design fulfils Assumptions 6.4.1 (a) and Equation (6.5), and its probability distribution $\tilde{\Lambda} \otimes \tilde{\Pi}$ satisfies, for all \mathcal{C} and d , $\tilde{\Lambda}_{\mathcal{C}}^d \propto B(\alpha; a_d, b_d) \mathbb{1}_{d \in L(\mathcal{C})} \times (1 - B(\alpha; a_d, b_d)) \mathbb{1}_{d \in U(\mathcal{C})}$ and

$$\tilde{\Pi}(\mathcal{C}) \propto \prod_{d \in L(\mathcal{C})} B(\alpha; a_d, b_d) \times \prod_{d \in U(\mathcal{C})} (1 - B(\alpha; a_d, b_d)),$$

where $B(\alpha; a_d, b_d)$ is the incomplete beta function. There are some practical problems with the indifferentiation of $\tilde{\Pi}$ and $\tilde{\Lambda}$. The modelization of PIPE is forced to use a very weak prior-model because the weight of information contained in $(a_d, b_d)_{d \in D}$ has a strong impact on the distribution $\tilde{\Pi}$. When this model is parametrized for a threshold at 0.25 with $\sum_{d \in D} a_d + b_d = 1$, as recommended in the original article, and the prior strength is constant across the doses (with $\#D = 12$), $B(\alpha, a_d, b_d)$ varies in the interval $]0.2227, 0.2677[$. The distribution Π is then almost proportional to a product of interval's length ($\simeq 0.25$ and $\simeq 0.75$). In this configuration, when the aim of the trial is to find a lower threshold as 0.1, the strength of the prior $\tilde{\Pi}$ is then increased. From our point of view, it is more convenient to choose directly the prior on \mathcal{K} and the prior model without hidden effect due to the threshold considered or the local strength

of prior model. This is the main theoretical benefit of the po-SPMc which corresponds to a broader class of methods more easily calibrated through the distinction of $\tilde{\Pi}$ and $\tilde{\Lambda}$.

The po-SPMc is a class of sequential methods for which the law $\tilde{\Pi}_n$ does not provide an obvious estimator of the next dose. The estimation of the next dose is made in two steps: 1) update the estimator of the contour \tilde{C}_n , and 2) according to the result of step 1, use a specific selection strategy to choose the next dose. We expose different strategies in the previous section, most of them are easily adaptable to the partial ordering case: Strategies 1, 3, 4. The strategy 2 was already introduced in Mander and Sweeting (2015, 2.1). They all satisfy the following assumption.

Assumption 6.4.2. (a) *The dose X_{n+1} is selected in $M(\tilde{C}_n)$.*

(b) *If the contour \mathcal{C} is infinitely recommended then all the doses in $M(\mathcal{C})$ are infinitely selected. In other terms,*

$$\left\{ n_{\mathcal{C}} = \sum_{i=1}^n \mathbb{1}_{\{\tilde{c}_i = \mathcal{C}\}} \xrightarrow{n \rightarrow \infty} \infty \right\} \implies \forall d \in M(\mathcal{C}), \left\{ n_d \xrightarrow{n \rightarrow \infty} \infty \right\}, \text{ a.s.}$$

Assumption 6.4.2 (a) is natural while Assumption 6.4.2 (b) is often satisfied and allows us to lay out an asymptotical property. This result is also valid for the PIPE method as particular case of the po-SPMc.

Theorem 6.4.1. *Under Assumptions 6.4.1 and 6.4.2, the po-SPMc leads to a strongly consistent estimator of the MTC and its minimal set, i.e. $\tilde{C}_n \xrightarrow{n \rightarrow \infty} \mathcal{C}^*$, a.s.*

Proof. We note, as in the proof of ε -balanced behaviour, that the regularity assumption 6.4.1 (or assumption alternative (6.5)) involves:

$$\delta(\beta_k, S_r^k) = \delta(\beta_k, S_t^k) \implies 0 < \liminf_{n \rightarrow \infty} \frac{M_{n,r}^k}{M_{n,t}^k} \leq \limsup_{n \rightarrow \infty} \frac{M_{n,r}^k}{M_{n,t}^k} < \infty, \text{ a.s.}$$

We will show the following assertion

$$\mathcal{C} \in \mathcal{K} \setminus \mathcal{C}^* \implies \mathbb{P}(n_{\mathcal{C}} \rightarrow \infty) = 0 \tag{6.6}$$

As $\mathcal{C} \neq \mathcal{C}^*$, we have $M(\mathcal{C}) \neq M(\mathcal{C}^*)$ and there exists $k \in M(\mathcal{C})$ such that β_k is included in $S_{\mathcal{C}}^k$ and not in $S_{\mathcal{C}^*}^k$. By using Lemma 5.8.1, we have

$$\mathbb{P} \left(\lim_{n \rightarrow \infty} \frac{M_{n,\mathcal{C}}^k}{M_{n,\mathcal{C}^*}^k} = 0 \mid n_{\mathcal{C}} \rightarrow \infty \right) = \mathbb{P} \left(\lim_{n \rightarrow \infty} \frac{M_{n,\mathcal{C}}^k}{M_{n,\mathcal{C}^*}^k} = 0 \mid n_k \rightarrow \infty \right) = 1,$$

where the first equality arises from Assumption 6.4.2(b). For all doses $d \in D$, the distribution

$\Lambda_{\mathcal{C}^*}$ models properly the toxicity, in other words $\beta_d \in S_{\mathcal{C}^*}$. We then have

$$\mathbb{P} \left(\lim_{n \rightarrow \infty} \frac{I_{n,\mathcal{C}}}{I_{n,\mathcal{C}^*}} = 0 \mid n_{\mathcal{C}} \rightarrow \infty \right) = 1 \text{ and } \mathbb{P} \left(\left\{ \lim_{n \rightarrow \infty} \frac{I_{n,\mathcal{C}}}{I_{n,\mathcal{C}^*}} = 0 \right\} \cap \{n_{\mathcal{C}} \rightarrow \infty\} \right) = 0,$$

which proves implication (6.6). \square

The convergence to \mathcal{C}^* is provided by the convergence to the minimal set in the same way as the balanced behaviour of the po-SPM. Indeed, there exist some relations between the modelling from the MTC and the MTD point of view. For illustrate these links, we introduce some modifications of the po-SPM. The range of doses on which is build the prior Π is extended to the set $\bar{D} = D \cup \{u, l\}$. The doses u and l are fictive. They can be seen as additional parametrization of doses $(1, 1)$ and (I, J) , respectively. From the point of view of l and u , all the toxicities are respectively lower and upper than the threshold α . The topological support of Λ satisfies now the following assumption.

Assumption 6.4.3. *All the distributions Λ_{θ} are product of unidimensional distributions. Their topological supports satisfy Assumption 6.3.3(ii) with $\varepsilon = 0$, and we have $S_u = A^m$ and $S_l = B^m$.*

Let $\bar{\Pi}_n$ be the distribution provided by Π on the range of doses D such as $\bar{\Pi}_n(1, 1) = \Pi_n(u) + \Pi_n(1, 1)$, $\bar{\Pi}_n(I, J) = \Pi_n(I, J) + \Pi_n(l)$ and $\bar{\Pi}_n(d) = \Pi_n(d)$ otherwise. The natural estimator of the MTD is then: $\hat{\theta}_n = \arg \max_{\theta \in D} \bar{\Pi}_n(\theta)$. Finally, for modification, we do not take into account the update of Π for the non-ordered dose. Such an assumption is reasonable having regard to the willingness to explore the whole contour. The posterior of Π according to the observations is then

$$\Pi_n(\theta) \propto \int \prod_{(i,j) \in D \setminus \mathcal{C}_{\theta}} q_{(i,j)}^{n_{(i,j)}^1} (1 - q_{(i,j)})^{n_{(i,j)}^0} \Lambda_{\theta}(dq) \Pi(\theta). \quad (6.7)$$

This update fosters doses poorly observed, that is doses d with only a few observations on d , A_d and B_d . This latest modification ensures the almost sure convergence of the po-SPM to the minimal set. The modelling from the MTD provides naturally a modelling from the MTC:

$$\forall \mathcal{C} \in \mathcal{K}, \tilde{\Pi}(\mathcal{C}) \propto \sum_{\theta \in M(\mathcal{C})} \Pi(\theta) \text{ and } \tilde{\Lambda}_{\mathcal{C}} = \sum_{\theta \in M(\mathcal{C})} r_{\mathcal{C}}^{\theta} \times \Lambda_{\theta}, \quad (6.8)$$

where $r_{\mathcal{C}}^{\theta} = \Pi(\theta) / \sum_{\theta \in M(\mathcal{C})} \Pi(\theta)$. Note that in Equation (6.8), the minimal set $M(\mathcal{C})$ might be replaced by the whole contour. When the po-SPM works under Relation (6.7) and Assumption 6.4.3, the po-SPMc defined by Equation (6.8) satisfies Assumption 6.4.1 (a). Furthermore, the relation between the distribution Π and $\tilde{\Pi}$ carries on.

Property 6.4.2. *If $\tilde{\Lambda} \otimes \tilde{\Pi}$ satisfies Equation (6.8), we have $\tilde{\Pi}_n(\mathcal{C}) \propto \sum_{\theta \in M(\mathcal{C})} \Pi_n(\theta)$.*

Proof. We have the following equalities.

$$\begin{aligned}\tilde{\Pi}_n(\mathcal{C}) &\propto \left[\int L_n(q) \tilde{\Lambda}_{\mathcal{C}}(dq) \right] \times \tilde{\Pi}(\mathcal{C}) \\ &= \left[\int L_n(q) \sum_{\theta \in M(\mathcal{C})} r_{\mathcal{C}}^{\theta} \times \Lambda_{\theta}(dq) \right] \times \left(\sum_{\theta \in M(\mathcal{C})} \Pi(\theta) \right) \\ &= F(\theta),\end{aligned}$$

where

$$\begin{aligned}F(\theta) &= \sum_{\theta \in M(\mathcal{C})} \left[\int L_n(q) \Lambda_{\theta}(dq) \right] \times r_{\mathcal{C}}^{\theta} \Pi(\theta) \\ &= \sum_{\theta \in M(\mathcal{C})} \Pi_n(\theta).\end{aligned}$$

□

We can notice that this result is still valid when we replace the minimal set $M(\mathcal{C})$ by the whole contour in Equation (6.8). Conversely, modelling from the MTC provides a distribution on D . The probability of toxicity of a dose θ is then proportional to the expected probability of all the contour \mathcal{C} for which θ lies in the minimal set $M(\mathcal{C})$. Let $N(\theta)$ be the set of all contour \mathcal{C} such that $\theta \in M(\mathcal{C})$. We set:

$$\forall \theta \in D, \Pi(\theta) \propto \sum_{\mathcal{C} \in N(\theta)} \tilde{\Pi}(\mathcal{C}) \text{ and } \Lambda_{\theta} = \sum_{\mathcal{C} \in N(\theta)} \tilde{r}_{\mathcal{C}}^{\theta} \times \tilde{\Lambda}_{\mathcal{C}},$$

with $\tilde{r}_{\mathcal{C}}^{\theta} = \tilde{\Pi}(\theta) / \sum_{\mathcal{C} \in N(\theta)} \tilde{\Pi}(\mathcal{C})$. This definition of Π on D , coming from a distribution on the contour, leads to a natural strategy selection: maximising Π on the minimal set of the selected contour. When we use such a strategy selection, the methods stemmed from a MTD or MTC modelling are exactly the same. However, every MTC model is not issued of an MTD modelling. The class of MTC methods is all the greater given that a wide variety of strategy selection can be used (see the previous chapter).

6.5 Appendix: Proof of asymptotical results for po-SPM

6.5.1 A lemma

In order to state asymptotical properties of the po-SPM, we introduce the set \tilde{D} of doses infinitely observed:

$$j \in \tilde{D} \Leftrightarrow n_j \xrightarrow[n \rightarrow \infty]{} \infty.$$

Lemma 6.5.1. *Let r be a dose in $\mathcal{E}(I, \beta) \cap \tilde{D}$. Then, there exists no dose d ordered with r in $\mathcal{E}(I, \beta) \cap \tilde{D}$.*

Proof. The proof is based on reductio ad absurdum. Let $r \in \mathcal{E}(I, \beta) \cap \tilde{D}$. Assume that there exists $d \in \mathcal{E}(I, \beta) \cap \tilde{D}$ ordered with r . As $d \in \tilde{D}$, the dose sequence $(x_n)_{n \in \mathbb{N}}$ has the following form

$$(x_n) = (\cdots, \underbrace{d, d, \cdots, d}_{s_1}, x_{j_0}, \cdots, x_{j_1}, \underbrace{d, \cdots, d}_{s_2}, x_{j_2}, \cdots, x_{j_3}, \underbrace{d, \cdots, d}_{s_3}, x_{j_4}, \cdots),$$

where $j_0 < j_1 < j_2 < \cdots$ and, if we denote S the number of sequences s_i , for all $i \in S$, s_i is a sequence including only the dose d . In words, the dose sequence $(x_n)_{n \in \mathbb{N}}$ switches between sequences including only the dose d and sequences which never include it. s_i is then the i -th sequence of dose d . We denote by A_i the random variable defined by

$$A_i = \begin{cases} \{\text{size}(s_i) < \infty\}, & \text{if } s_i \text{ exists,} \\ \emptyset, & \text{otherwise,} \end{cases}$$

where $\text{size}(s_i)$ denotes the number of terms in the sequence s_i . We are now interested in the evaluation of $\mathbb{P}(S = \infty)$. First we note B_i the random variable $\{\tilde{s}_i < \infty\}$, where $\tilde{s}_1, \tilde{s}_2, \cdots$ are independent sequences of dose d , such as $\Pi_{\tilde{n}_i}(d)/\Pi_{\tilde{n}_i}(r) = 1$ with \tilde{n}_i the start index of \tilde{s}_i . In other terms, we place us in the most negative case, *i.e.* the sequence switches from a dose $x_{\tilde{n}_i-1} \neq d$ to the dose d with a ratio of probabilities equals to 1. Then,

$$\begin{aligned} \mathbb{P}(S = \infty) &= \mathbb{P}\left(\bigcap_{i \in \mathbb{N}} A_i\right) = \mathbb{P}(A_1)\mathbb{P}(A_2|A_1) \cdots \mathbb{P}(A_n|A_{n-1}, \cdots, A_1) \cdots \\ &\leq \mathbb{P}(B_1)\mathbb{P}(B_2) \cdots \mathbb{P}(B_n) \cdots \end{aligned} \tag{6.9}$$

$$= \mathbb{P}\left(\bigcap_{i \in \mathbb{N}} B_i\right), \tag{6.10}$$

where Inequality (6.9) is due to the definition of the random variables B_i and Equation (6.10) to their independence.

Moreover, we have that for all $i \in \mathbb{N}$, $\mathbb{P}(B_i) < 1$. Indeed, let us define the sequence of doses $(x_n)_{n \in \mathbb{N}}$ such that $x_n = d$ for all $n \in \mathbb{N}$. By the law of the iterated logarithm, there exists $N \in \mathbb{N}$ such that $\mathbb{P}(B_0) = 1$, where $B_0 = \{x_N, x_{N+1}, \cdots\}$. We denote by A_0 the set $\{x_1, \cdots, x_{N-1}\}$. As the length of A_0 is finite, we know that $\mathbb{P}(A_0) > 0$. With the inequality $\mathbb{P}(\bar{B}_i) \geq \mathbb{P}(A_0 \cap B_0)$, due to the definition of B_i , we have that $\mathbb{P}(B_i) < 1$, for all $i \in \mathbb{N}$. Thus $\mathbb{P}(\bigcap_{i \in \mathbb{N}} B_i) < 1$. By the Kolmogorov's zero-one law, it implies that $\mathbb{P}(\bigcap_{i \in \mathbb{N}} B_i) = 0$ and then $\mathbb{P}(S = \infty) = 0$. In words, it means that there exists a last infinite sequence of dose d . This is in contradiction with the fact that $r \in \tilde{D}$. \square

6.5.2 Proof of theorem 6.3.2

Proof. Let us start with the proof of the ε -sensitivity. In this proof, we are interested in an asymptotic behaviour of the po-SPM, that is why we ignore the doses tested only a finite number of times and we reason as if they were never allocated. In other words, the doses in this proof are always considered in \tilde{D} . Assume now that $\mathcal{E}(I, \beta)$ is not empty. Let $r \in \tilde{D} \setminus \mathcal{E}(I, \beta)$. We can distinguish two cases.

The first one is the existence of a dose $d \in \mathcal{E}(I, \beta)$ such that d is ordered with r . We are then reduced to the SPM in case of total ordering and the proof can be found in chapter 3 (annexe 1).

The second one is the opposite, *i.e.* there exists no dose in $\mathcal{E}(I, \beta)$ ordered with r . So there exists a dose $d \in \mathcal{E}(I, \beta)$ not ordered with r , because we assume that $\mathcal{E}(I, \beta)$ is not empty. We want now to compare the integrals $I_{n,r}$ and $I_{n,d}$.

$$\frac{I_{n,r}}{I_{n,d}} = \prod_{k \in \tilde{D}} \frac{M_{n,r}^k}{M_{n,d}^k} = \prod_{k \in \tilde{D}} \frac{\int_{S_r^k} g(q_k, n_k^1, n_k^0) \Lambda_r(dq_k)}{\int_{S_d^k} g(q_k, n_k^1, n_k^0) \Lambda_d(dq_k)} \quad (6.11)$$

$$= \prod_{k \in \tilde{D}} \frac{\int_{S_r^k} g(q_k, n_k^1, n_k^0) \lambda_r(q_k) dq_k}{\int_{S_d^k} g(q_k, n_k^1, n_k^0) \lambda_d(q_k) dq_k} \quad (6.12)$$

$$\leq \frac{S}{s} \prod_{k \in \tilde{D}} \frac{\int_{S_r^k} g(q_k, n_k^1, n_k^0) dq_k}{\int_{S_d^k} g(q_k, n_k^1, n_k^0) dq_k}, \quad (6.13)$$

where Equation (6.11) follows from Assumption 6.3.1, Equation (6.12) from Assumption 6.3.4 (a) and Inequation (6.13) from Assumption 6.3.4 (b). We state then the following property. For all function f continuous on $[0, 1]$, we have

$$\int f(q_k) \frac{g(q_k, n_k^1, n_k^0)}{\text{Beta}(n_k^1 + 1, n_k^0 + 1)} dq_k \xrightarrow{n_k \rightarrow \infty} \int f(q_k) \text{pi}_{\{\beta_k\}}(q_k) \gamma(dq_k) = f(\beta_k), \quad (6.14)$$

where $\text{Beta}(\cdot)$ denotes the Beta function and γ the counting measure. Let then $k \in \tilde{D}$. We are looking for the behaviour, when n_k goes to infinity, of

$$Q_k = \frac{\int_{S_r^k} g(q_k, n_k^1, n_k^0) dq_k}{\int_{S_d^k} g(q_k, n_k^1, n_k^0) dq_k}.$$

The case $S_r^k = S_d^k$ is obvious. For the other cases, we use the convergence expressed in Equation (6.14). As d is not ordered with r , with Assumption 6.3.3, we have six cases left to study.

- $S_r^k = A$ or B or I and $S_d^k = [0, 1]$. Then $Q_k \xrightarrow{n_k \rightarrow \infty} \text{pi}_A(\beta_k)$ or $\text{pi}_B(\beta_k)$ or 0. This last result is due to the fact that $r \notin \mathcal{E}(I, \beta)$.
- $S_r^k = [0, 1]$ and $S_d^k = I$. Then $Q_k \xrightarrow{n_k \rightarrow \infty} 1$, because $d \in \mathcal{E}(I, \beta)$.

- $S_r^k = [0, 1]$ and $S_d^k = A$. Then $Q_k \xrightarrow[n_k \rightarrow \infty]{} 1/\mathbb{1}_A(\beta_k) = 1$. Indeed, with Lemma 6.5.1, as $d \in \mathcal{E}(I, \beta) \cap \tilde{D}$, $k \in \tilde{D}$ cannot be in $\mathcal{E}(I, \beta)$ and so $\beta_k \in A$.
- $S_r^k = [0, 1]$ and $S_d^k = B$. With the same argument as the previous case, we have $Q_k \xrightarrow[n_k \rightarrow \infty]{} 1$.

Finally, going back to Inequality (6.13), we conclude that $I_{n,r}/I_{n,d}$ tends to 0 when n goes to infinity. This leads to a contradiction because, as $r \in \tilde{D}$, this ratio is greater than 1 infinitely often. So $\tilde{D} \subset \mathcal{E}(I, \beta)$, *i.e.* the po-SPM is ε -sensitive.

We prove now that the po-SPM is also ε -balanced. Assumption 6.3.1 allows us to focus on the marginal ratio

$$\frac{M_{n,r}^k}{M_{n,t}^k} = \frac{\int g(q_j, n_k^1, n_k^0) \Lambda_r^k(dq_k)}{\int g(q_k, n_k^1, n_k^0) \Lambda_t^k(dq_k)} \quad \text{and} \quad \frac{I_{n,r}}{I_{n,t}} = \prod_{k \in D} \frac{M_{n,r}^k}{M_{n,t}^k}.$$

We note, as in the lemma 3.6.1, that Assumption 6.3.4 involves:

$$d(\beta_k, S_r^k) = d(\beta_k, S_t^k) \implies 0 < \liminf_{n \rightarrow \infty} \frac{M_{n,r}^k}{M_{n,t}^k} \leq \limsup_{n \rightarrow \infty} \frac{M_{n,r}^k}{M_{n,t}^k} < +\infty, \quad a.s. \quad (6.15)$$

The ε -balanced behaviour corresponds to the case where, for all $k \in D$, $\beta_k \notin I$. We show that

$$r \notin M_D \implies \mathbb{P}(\{n_r \rightarrow \infty\}) = 0 \quad (6.16)$$

By symmetry we can choose $r \in L$. The set $M_D \cap L \cap A_r$ is not empty. Let t be a dose in $M_D \cap L \cap A_r$. By using Equation (6.15), as $A_t \subset A_r$ and $S_t^k = [0, 1]$ when $k \in C_t$, we have

$$\forall (i, j) \in A_t \cup C_t, \quad \mathbb{P} \left(\limsup_{n \rightarrow \infty} \frac{M_{n,r}^k}{M_{n,t}^k} < \infty \mid n_k \rightarrow \infty \right) = 1.$$

If $k \in (B_r \cup C_r) \cap B_t$ then $\beta_{ji} \in B$ and we obtain the same result

$$\mathbb{P} \left(\limsup_{n \rightarrow \infty} \frac{M_{n,r}^k}{M_{n,t}^k} < \infty \mid n_k \rightarrow \infty \right) = 1.$$

We consider now the ratios $M_{n,r}^k/M_{n,t}^k$ when $k \in (B_t \cap A_r) \cup \{t, r\}$. In that case, we have $\beta_k \in B$, $S_t^k = B$ and S_r^k equals to I or A . By using Lemma 5.8.1, we have

$$\forall k \in (B_t \cap A_r) \cup \{t, r\}, \quad \mathbb{P} \left(\lim_{n \rightarrow \infty} \frac{M_{n,r}^k}{M_{n,t}^k} = 0 \mid n_k \rightarrow \infty \right) = 1.$$

We then have

$$\mathbb{P}\left(\lim_{n \rightarrow \infty} \frac{I_{n,r}}{I_{n,t}} = 0 \mid n_r \rightarrow \infty\right) = 1 \text{ and } \mathbb{P}\left(\left\{\lim_{n \rightarrow \infty} \frac{I_{n,r}}{I_{n,t}} = 0\right\} \cap \{n_r \rightarrow \infty\}\right) = 0,$$

which proves Equation (6.16). We achieve the proof of the ε -balanced property by showing that

$$t \in M_D \implies \mathbb{P}(n_t \rightarrow \infty) = 1 \tag{6.17}$$

By using Equation (6.16), we have, for all r and t in M_D

$$\forall (i, j) \in D, k \neq r \text{ and } k \neq t, \mathbb{P}\left(\limsup_{n \rightarrow \infty} \frac{M_{n,r}^k}{M_{n,t}^k} < \infty\right) = 1.$$

Moreover, for all $r \in M_D$, S_t^r is equal to the segment $[0, 1]$. As $\beta_r \notin I = S_r^r$, we have

$$\forall r \in M_D, \mathbb{P}\left(\lim_{n \rightarrow \infty} \frac{M_{n,r}^r}{M_{n,t}^r} = 0 \mid n_r \rightarrow \infty\right) = 1.$$

Let E_r be the event $\{n_r \rightarrow \infty\} \cap \{n_t \rightarrow \infty\}^c$. Then,

$$\mathbb{P}\left(\lim_{n \rightarrow \infty} \frac{I_{n,r}}{I_{n,t}} = 0 \mid E_r\right) = 1 \text{ and } \mathbb{P}\left(\left\{\lim_{n \rightarrow \infty} \frac{I_{n,r}}{I_{n,t}} = 0\right\} \cap E_r\right) = 0,$$

As $\mathbb{P}(\sum_{M_D} n_k \rightarrow +\infty) = 1$, we have $\mathbb{P}(\cup_{M_D} E_r) = 1$, which proves Equation (6.17) and ends the demonstration of the po-SPM ε -balanced behaviour. \square

Bibliography

- Philippe Amiel and François Vialla. La vérité perdue du "code de nuremberg": réception et déformations du "code de nuremberg" en france. *Revue de droit sanitaire et social*, (4): 673–687, 2009.
- 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 Shelemياهو Zacks. Cancer phase i clinical trials: efficient dose escalation with overdose control. *Statistics in medicine*, 17(10):1103–1120, 1998.
- T. M. Braun and N. Jia. A generalized continual reassessment method for two-agent phase i trials. *Statistics in Biopharmaceutical Research*, 5(2):105–115, 2013.
- T. M. Braun and S. Wang. A hierarchical bayesian design for phase i trials of novel combinations of cancer therapeutic agents. *Biometrics*, 66(3):805–812, 2010.
- Leo Breiman. *Probability, volume 7 of Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992.
- Michel Broniatowski and Amor Keziou. Divergences and duality for estimation and test under moment condition models. *Journal of Statistical Planning and Inference*, 142(9):2554–2573, 2012.
- K. Cheung. Coherence principles in dose-finding studies. *Biometrika*, 92(4):863–873, 2005.
- Y. K. Cheung. *Dose finding by the continual reassessment method*. CRC Press, 2011.
- Ying Kuen Cheung and Rick Chappell. A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics*, 58(3):671–674, 2002.
- Joseph L Doob. Application of the theory of martingales. *Le calcul des probabilités et ses applications*, pages 23–27, 1949.

- Stephen D Durham and Nancy Flournoy. Random walks for quantile estimation. In *Statistical Decision Theory and Related Topics V*, pages 467–476. Springer, 1994.
- Stephen D Durham, Nancy Flournoy, and William F Rosenberger. A random walk rule for phase i clinical trials. *Biometrics*, pages 745–760, 1997.
- David A Freedman. On the asymptotic behavior of bayes' estimates in the discrete case. *The Annals of Mathematical Statistics*, pages 1386–1403, 1963.
- Mauro Gasparini and Jeffrey Eisele. A curve-free method for phase i clinical trials. *Biometrics*, 56(2):609–615, 2000.
- Constantine Gatsonis and Joel B Greenhouse. Bayesian methods for phase i clinical trials. *Statistics in medicine*, 11(10):1377–1389, 1992.
- Irving John Good. Some history of the hierarchical bayesian methodology. *Trabajos de estadística y de investigación operativa*, 31(1):489–519, 1980.
- X. Huang, S. Biswas, Y. Oki, J. Issa, and D. A. Berry. A parallel phase i/ii clinical trial design for combination therapies. *Biometrics*, 63(2):429–436, 2007.
- Alexia Iasonos and John O'Quigley. Design considerations for dose-expansion cohorts in phase i trials. *Journal of Clinical Oncology*, 31(31):4014–4021, 2013.
- Anastasia Ivanova, Nancy Flournoy, and Yeonseung Chung. Cumulative cohort design for dose-finding. *Journal of Statistical Planning and Inference*, 137(7):2316–2327, 2007.
- Y. Ji, Y. Li, and N. Bekele. A modified toxicity probability interval method for dose-finding trials. *Clinical Trials*, 7(6):653–663, 2010.
- Olav Kallenberg. *Foundations of modern probability*. Springer Science & Business Media, 2006.
- Edward L Korn, Douglas Midthune, T Timothy Chen, Lawrence V Rubinstein, Michael C Christian, and Richard M Simon. A comparison of two phase i trial designs. *Statistics in medicine*, 13(18):1799–1806, 1994.
- Friedrich Liese and Klaus-J Miescke. *Statistical decision theory: estimation, testing, and selection*. Springer Science & Business Media, 2008.
- 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.
- A. P. Mander and M. J. Sweeting. A product of independent beta probabilities dose escalation design for dual-agent phase i trials. *Statistics in Medicine*, 34(8):1261–1276, 2015. ISSN 1097-0258.

- John O'Quigley and Larry Z Shen. Continual reassessment method: a likelihood approach. *Biometrics*, pages 673–684, 1996.
- John O'Quigley and S Zohar. Experimental designs for phase i and phase i/ii dose-finding studies. *British journal of cancer*, 94(5):609–613, 2006.
- 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, Larry Z Shen, and Anthony Gamst. Two-sample continual reassessment method. *Journal of biopharmaceutical statistics*, 9(1):17–44, 1999.
- John O'Quigley, Xavier Paoletti, and Jean Maccario. Non-parametric optimal design in dose finding studies. *Biostatistics*, 3(1):51–56, 2002.
- Steven Piantadosi and Guanghan Liu. Improved designs for dose escalation studies using pharmacokinetic measurements. *Statistics in medicine*, 15(15):1605–1618, 1996.
- Giovanni C Porzio and Giancarlo Ragozini. On the stochastic ordering of folded binomials. *Statistics & Probability Letters*, 79(9):1299–1304, 2009.
- Ethan Reiner, Xavier Paoletti, and John O'Quigley. Operating characteristics of the standard phase i clinical trial design. *Computational Statistics & Data Analysis*, 30(3):303–315, 1999.
- William F Rosenberger and Linda M Haines. Competing designs for phase i clinical trials: a review. *Statistics in medicine*, 21(18):2757–2770, 2002.
- Stephen J Ruberg. Dose response studies i. some design considerations. *Journal of Biopharmaceutical Statistics*, 5(1):1–14, 1995.
- Mark J Schervish. *Theory of statistics*. Springer Science & Business Media, 2012.
- Larry Z Shen and John O'Quigley. Consistency of continual reassessment method under model misspecification. *Biometrika*, 83(2):395–405, 1996.
- J. M. Skolnik, J. S. Barrett, B. Jayaraman, D. Patel, and P. 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.
- William F Stout. A martingale analogue of kolmogorov's law of the iterated logarithm. *Probability Theory and Related Fields*, 15(4):279–290, 1970.
- P. F. Thall, R. E. Millikan, P. Mueller, and S. Lee. Dose-finding with two agents in phase i oncology trials. *Biometrics*, 59(3):487–496, 2003.

- N. A. Wages, M. R. Conaway, and J. O'Quigley. Continual reassessment method for partial ordering. *Biometrics*, 67(4):1555–1563, 2011.
- K. Wang and A. Ivanova. Two-dimensional dose finding in discrete dose space. *Biometrics*, 61(1):217–222, 2005. ISSN 1541-0420.
- J. Whitehead, H. Thygesen, and A. Whitehead. A bayesian dose-finding procedure for phase i clinical trials based only on the assumption of monotonicity. *Statistics in Medicine*, 29(17):1808–1824, 2010.
- John Whitehead and David Williamson. Bayesian decision procedures based on logistic regression models for dose-finding studies. *Journal of Biopharmaceutical Statistics*, 8(3):445–467, 1998.
- G. Yin and Y. Yuan. A latent contingency table approach to dose-finding for combinations of two agents. *Biometrics*, 65(3):866–875, 2009a.
- G. Yin and Y. Yuan. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 58(2):211–224, 2009b. ISSN 1467-9876.
- Sarah Zohar and John O'Quigley. Identifying the most successful dose (msd) in dose-finding studies in cancer. *Pharmaceutical statistics*, 5(3):187–199, 2006.