SEMIPARAMETRIC DOSE FINDING METHODS FOR PARTIALLY ORDERED DRUG COMBINATIONS

Matthieu Clertant

LAGA, LabEx Inflamex, Université Sorbonne Paris Nord, France

Nolan A. Wages

Translational Research and Applied Statistics, University of Virginia, U.S.A.

John O'Quigley

Department of Statistical Science, University College London, U.K.

Supplementary Material

The 8 sections contain: scenarios used in simulation, a notation table, numerical experiment in the MTD setting, the calibration methods and the proofs of all the properties and theorems.

S1. Scenarios for the MTC setting

Table 1: *Toxicity scenarios 1-4 for the two-drug combinations. Dose combinations in the minimal set are in bold. Doses associated to a toxicity probability equal to* 0.20 *belong to* D^- *and* D^+ *(see Section 4); they are always in the minimal set.*

	Dose								
		Drug A2							
	Level	1	2	3	4	1	2	3	4
		Scenario 1				Scenario 2			
	4	**22**	26	30	34	55	65	75	85
	3	16	**20**	24	28	40	50	60	70
	2	10	14	**18**	**22**	**25**	35	45	55
	1	4	8	12	**16**	10	**20**	30	40
Drug									
A1		Scenario 3				Scenario 4			
	4	11	21	31	41	**10**	**30**	50	80
	3	10	**20**	30	31	6	**15**	**30**	45
	2	9	19	29	30	4	10	**15**	**20**
	1	8	18	**28**	29	1	2	3	**4**

S2. Notations

Table 2: *Important notations*

Notation	**Description**
d or X	a dose combination
Y	response in term of toxicity
D	set of dose combinations
I and J	number of doses for agent 1 and 2
$\mathcal{A}_d$, $\mathcal{B}_d$ and $\mathcal{C}_d$	set of doses above, below and non-ordered with d
P	a scenario, i.e a $I \times J$ matrix of bernoulli parameters
c and $\mathcal{M}_c$	a contour and a minimal set associated to this contour
$\mathcal{A}_c$ and $\mathcal{B}_c$	set of dose combinations above and below c
C	set of possible contour in the range of dose combinations
$d^\star$	true Maximum Tolerated Dose (MTD)
$c^\star$	true Maximum Tolerated Contour (MTC)
Π	distribution on the MTD or the MTC parameter
θ	MTD parameter
γ	MTC parameter
$\Lambda = (\Lambda_\theta)_{\theta \in D}$	distributions on the scenarios conditioned by the MTD parameter
$\Lambda = (\Lambda_\gamma)_{\gamma \in C}$	distributions on the scenarios conditioned by the MTC parameter
S, S_θ, S_γ	support of Λ, Λ_θ and Λ_γ
Λ_θ^d (or Λ_γ^d)	marginal at dose combination d

S3. Allocation strategy

In a minimal set, many of the doses are not ordered and the question of allocation within this set needs to be addressed. This is not straightforward and we do not exhaust all possibilities here. Several allocation strategies can be made available and any particular choice will result in some particular

trial behaviour. The allocation strategy introduced hereafter is a compromise between one that spreads experimentation 'equally' against one that tests those doses that will, in expectation, bring more information into the study. We also need for experimentation to be restricted to an area of the minimal set considered to be safe.

While allocating patients within the estimated minimal set, we want to do our best to avoid the selection of dose combinations which are already associated with having an unacceptably high toxicity probability. We can address this issue by making use of a Bayesian test based on a uniform prior $\mathcal{U}$ on the space of Bernoulli parameters. The set H_n describes those dose combinations indicated as being too toxic following these local Bayesian tests at each of the doses.

$$H_n = \{d \in D : \mathbb{P}_{\mathcal{U}} \left[P_d > \theta_T | (n_d, n_d^1) \right] > \delta_T \} .$$

We are then in a position to use the partial ordering in D to extend these exclusions to other doses that are unfavourably ordered with the doses belonging to this set, e.g. a dose d' in $\mathcal{A}_d$ with d belonging to H_n.

Exclusion rule: A dose d_0 is excluded from the study if:

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

where T_n is the set of dose combinations considered as overly toxic. A

dose combination of the estimated minimal set $\mathcal{M}_{\hat{\gamma}_n}$ will not be tested if it belongs to T_n. Note that this exclusion rule corresponds also to a stopping rule when the evidence of unacceptably high toxicity is associated with the lowest dose. This rule is an extension in the context of partial ordering to rules introduced by Ji et al (2010).

The approach based on a compromise between spreading the observations along the contour as opposed to making observations on those doses that we estimate would bring more information into the study leans on two quantities: (i) the number of doses which will potentially benefit from the next observation through the partial ordering structure and (ii) the amount of information already collected. Each dose d, can be associated with a value k_d corresponding to the number of dose combinations ordered with d plus one: $k_d = \#\mathcal{A}_d + \#\mathcal{B}_d + 1$, where the cardinality of a set E is noted $\overline{E}$. The value k_d corresponds to the number of dose combinations on which we aim to learn something subsequent to an observation on d. Observing a DLT at dose d implies that the patient would have experienced a DLT for all the doses in $\mathcal{A}_d$, and conversely, all the doses in $\mathcal{B}_d$ would have been safe for a patient who does not experience a DLT at dose d. This structure is naturally taken into account by our model. When the objective is to optimize the quantity of information obtained during the trial, those doses associated

with the highest value k_d are, in expectation, those that provide the most information. Note that these doses are located along the diagonal going from the combination $(1,1)$ to (I,J). The notion of entropy is our second tool. It corresponds to a quantity of information per unit of data, in our case the observation related to a single patient. As the goal of the study is to determine doses that are close to α in term of toxicity probability, we will use the concept of relative entropy between each observation and the target α. If p and q denotes two Bernoulli distributions and their parameters, the entropy of p relative to q is: $H(p|q) = -q \times \log(p) - (1-q)\log(1-p)$, with the convention $\log(0) = -\infty$ and $0 \times (-\infty) = 0$. Thus, $n \times H(\alpha|1)$ is the quantity of information relative to the target alpha which is brought by n DLT ($H(\alpha|0)$ for the non toxic observations). It corresponds to the log-likelihood in α. After n observations, the next dose X_{n+1} is:

$$X_{n+1} = \min_{d \in \mathcal{M}_{\hat{\gamma}_n} \setminus T_n} \frac{1}{k_d} \left\{ \left(\sum_{j \in \mathcal{B}_d \cup \{d\}} n_j^0 \right) \times H(\alpha|0) + \left(\sum_{j \in \mathcal{A}_d \cup \{d\}} n_j^1 \right) \times H(\alpha|1) \right\}. \tag{S3.1}$$

This allocation strategy spreads the observations along the estimated contour and tests slightly more often the doses in the middle of the grid. Since a DLT corresponds to a greater quantity of information those doses with a smaller number of toxicities are tried more often (for $\alpha = 0.25$: $H(\alpha|1)/H(\alpha|0) = \log(\alpha)/\log(1-\alpha) \approx 4.82$).

S4. Numerical experiments: MTD setting

We evaluate the operating characteristics of poSPM and compared them to the 2d-BOIN method (Lin et al., 2016) and the poCRM (Wages, Conaway and O'Quigley, 2011). Performance evaluation was based on several metrics under 4 toxicity scenarios. Our goal is to evaluate: (1) how well each method provides a recommendation of sets of doses at and around the target rate (i.e. acceptable MTD's), and (2) how well each method allocates patients to acceptable MTD's. While traditional evaluation measures, such as the percentage of recommendation and allocation to the true MTD's are useful in assessing performance, it is also beneficial to consider the entire distribution of selected dose combination, as it provides more detailed information as to what combinations are being recommended. For evaluating recommendation, Cheung (2011) proposes to use the accuracy index so that, after n patients, he defines,

$$A_n = 1 - I \times J \times \frac{\sum_{i=1}^{I}\sum_{j=1}^{J}(P_d - \alpha)^2 \times \rho_d}{\sum_{i=1}^{I}\sum_{j=1}^{J}(P_d - \alpha)^2}, \tag{S4.1}$$

where P_d is the true toxicity probability at dose combination $d = (i, j)$, ρ_d is the percentage of trials in which combination $d = (i, j)$ was selected as the MTD, and n is the total sample size. For experimentation, the same

formula can be used with ρ_d representing the percentage of patients treated at combination $d = (i, j)$. The maximum value of A_n is 1 with larger values (close to 1) indicating that the method has high accuracy. For each method, we simulated 10,000 trials under the 4 different sets of assumed DLT probabilities in a 6×6 grid of combinations with varying positions and number of true MTD's, as shown in Table 3. The target toxicity rate is $\alpha = 0.25$ and the total sample size for each simulated study is 40 patients. For each method, a cohort of size 1 is used. The calibration of the three methods is fully described in Appendix S5.1 and S5.2.

Table 4 shows the operating characteristics of the 3 methods under 10 000 trials for the 4 scenarios considered. For each scenario, we report the percentage of patient allocation (experimentation %) and percentage of MTD selection (recommendation %) for doses contained within five different ranges of the true toxicity probabilities. Since $\alpha = 0.25$, the target interval containing the true MTD is $[0.20, 0.30]$. In Scenarios 1 and 2, the poSPM and poCRM have very comparable performance, with each method outperforming 2d-BOIN in terms of experimentation %, recommendation %, as well as accuracy of experimentation and recommendation. For both of these scenarios, 2d-BOIN induces a higher percentage of DLT's, on average. In Scenario 3, the 2d-BOIN exhibits the best performance when

Table 3: *True toxicity probabilities for the four scenarios of a* 6×6 *grid, with maximum tolerated doses shown in bold*

	Dose	Drug A1											
	Level	1	2	3	4	5	6	1	2	3	4	5	6
		Scenario 1						Scenario 2					
	6	20	29	31	43	47	50	37	45	51	54	55	58
	5	18	20	29	34	41	48	30	38	41	43	46	47
	4	16	19	21	32	36	42	23	**25**	35	36	40	42
	3	10	15	20	**25**	30	37	19	20	26	33	35	39
	2	3	9	16	19	21	32	15	17	21	24	31	33
	1	2	5	10	17	21	30	5	10	16	20	**25**	28
Drug													
A2		Scenario 3						Scenario 4					
	6	11	13	15	17	**25**	33	65	70	76	80	84	90
	5	9	11	13	15	16	**25**	55	63	69	77	80	85
	4	7	9	11	14	15	17	45	56	60	72	76	79
	3	5	8	10	11	13	15	35	42	52	65	70	73
	2	3	6	7	9	12	13	**25**	34	46	54	60	64
	1	1	3	5	7	9	11	15	**25**	36	43	49	55

considering experimentation and recommendation percentages in the target interval. The accuracy of recommendation of the poSPM is very close to that of the 2d-BOIN due to the higher recommendation percentage for 2d-BOIN on the toxicity probability interval $[0, 0.10)$. In the final scenario, poSPM is the best performing method according to all metrics.

We anticipate that there will always be certain scenarios in which some methods perform better than others. A useful tool in this regard for comparing dose-finding designs can be average performance over a range of scenarios. Across the 4 scenarios, the poSPM method, the 2d-BOIN design and poCRM methods demonstrated averages of 46.7%, 42.5%, and 43.3% recommendation percentages for combinations in the true target interval $[0.20, 0.30]$, respectively. The overall percentage of observed toxicities of the poSPM method, the 2d-BOIN design, and poCRM methods were on average; 24.1%, 27.3%, and 23.1%, respectively. It is desirable for the value to be as close as possible to the target rate α. This is achieved by the poSPM. The average percentage of patients allocated to a dose combination in the target interval of the poSPM method, the 2d-BOIN design, and poCRM methods were 34.13%, 30.75%, and 31.75%, respectively. Based on the accuracy index for recommendation, the poSPM yielded an average value of 0.77, the 2d-BOIN method produced an average value of 0.73, and

Table 4: *Experimentation and recommendation percentages for the poSPM, the BOIN design and the poCRM in four scenarios of a* 6×6 *grid of true toxicity probabilities. Column containing the MTD in bold.*

		Experimentation %						
Scenario	Method	$[0, 0.10)$	$[0.10, 0.20)$	$[\mathbf{0.20}, \mathbf{0.30}]$	$(0.30, 0.40]$	$(0.40, 1]$	A_{40}	%DLTs
	poSPM	11.1	18.7	**42.3**	22.3	6.6	0.35	23.9
1	BOIN	8.6	16.2	**32.0**	23.9	19.3	0.34	27.5
	poCRM	11.0	26.0	**43.0**	17.0	4.0	0.37	22.2
	poSPM	4.6	18.5	**44.0**	26.6	6.4	0.54	25.8
2	BOIN	4.3	15.3	**40.0**	31.2	18.2	0.26	29.3
	poCRM	4.0	22.0	**47.0**	24.0	4.0	0.64	25.1
	poSPM	14.6	52.6	**17.7**	15.1	0.0	0.38	17.5
3	BOIN	14.3	38.7	**25.1**	21.8	0.0	0.48	20.0
	poCRM	25.0	57.0	**10.0**	9.0	0.0	0.20	14.6
	poSPM	0.0	24.2	**32.5**	25.5	17.8	0.90	29.2
4	BOIN	0.0	18.4	**25.9**	24.4	28.9	0.83	32.3
	poCRM	0.0	18.0	**27.0**	29.0	26.0	0.86	32.3
		Recommendation %						
		$[0, 0.10)$	$[0.10, 0.20)$	$[\mathbf{0.20}, \mathbf{0.30}]$	$(0.30, 0.40]$	$(0.40, 1]$	A_{40}	
	poSPM	0.4	16.0	**54.2**	25.1	4.3	0.68	
1	BOIN	1.2	16.9	**46.6**	27.4	7.45	0.65	
	poCRM	1.0	24.0	**56.0**	16.0	3.0	0.68	
	poSPM	0.1	10.1	**56.5**	29.6	3.7	0.74	
2	BOIN	0.1	13.69	**44.5**	33.5	8.1	0.63	
	poCRM	0.0	15.0	**59.0**	25.0	3.0	0.76	
	poSPM	0.6	47.0	**30.6**	21.8	0.0	0.71	
3	BOIN	2.5	44.3	**37.4**	15.8	0.0	0.72	
	poCRM	12.0	63.0	**18.0**	8.0	0.0	0.46	
	poSPM	0.0	17.5	**45.5**	28.7	8.2	0.95	
4	BOIN	0.0	11.5	**41.4**	30.7	13.3	0.93	
	poCRM	0.0	12.0	**40.0**	35.0	11.0	0.94	

the poCRM design resulted in an average value of 0.71. For accuracy of experimentation, the average index values over the four scenarios were 0.54, 0.48, and 0.52 for the poSPM, 2d-BOIN, and poCRM, respectively.

S5. Calibration

S5.1 poCRM, 2d-BOIN and PIPE methods

For poCRM, we utilized six possible orderings in all scenarios, arranging the combinations across rows, up columns, and up or down any diagonal as suggested by Wages and Conaway (2013). A uniform prior was placed on the orderings. The skeleton values for poCRM were generated according to the algorithm of Lee and Cheung (2009) using the **getprior** function in `R` package **dfcrm**. Specifically, we used **getprior(0.035,0.25,12,36)**, and all simulation results were generated using the functions of the `R` package **poCRM**.

For the BOIN method, we used the default cutoff values $\alpha_1 = 0.6\alpha$ and $\alpha_2 = 1.4\alpha$ and a Beta$(0.5, 0.5)$ prior for the toxicity probability at each combination. The boundaries for the optimal interval for the BOIN method are $(0.197, 0.298)$.

The prior distributions for PIPE are set to be the true DLT probabilities taken from Scenario 1 in Table 1. PIPE uses a neighbourhood dose-skipping

constraint, as well as the closest doses chosen from the admissible dose set. The dose escalation algorithm employs the smallest sample size strategy, and a weak prior distribution (1/16) is specified.

S5.2 poSPM

The marginals of the prior model are beta distributions truncated according to their support (Assumption 3). Let $B_I(a, b)$ be the truncated beta distribution with parameters a and b and let $\mathsf{B}_I(m, T)$ be an alternative parametrization of this distribution such that: $\mathsf{B}_I(m, T) = B(m \times T + 1, (1 - m) \times T + 1)$. Then, m is the mode of the beta distribution and T is a dispersion parameter.

The family of $I \times J$ matrix $(m_\theta)_{\theta \in D}$ is associated with the prior model $(\Lambda_\theta)_{\theta \in D}$. The element of the matrix m_θ in position d, noted m_θ^d, is the mode of the marginal Λ_θ^d. The value T_1 and T_2 are the dispersion parameters used on the different marginals. In the Bayesian setting, they can be equated to a number of pseudo observations providing toxicities and non-toxicities. The rank of a dose combination $\theta = (i, j)$, noted $\overline{\theta}$, is equal to $i + j$. All the doses on the same diagonal have the same rank. Let r_1 and r_2 be two positive values used for the calibration of the prior Π. The poSPM model

used in the simulations can be summarized by:

$$\begin{cases} \Lambda_\theta & \sim \prod_{d\in\mathcal{B}_\theta} \mathsf{B}_B(m_\theta^d, T_1) \times \prod_{d\in\mathcal{A}_\theta} \mathsf{B}_A(m_\theta^d, T_1) \times \prod_{d\in\mathcal{C}_\theta} \mathsf{B}_{[0,1]}(m_\theta^d, T_2) \times \mathsf{B}_I(m_\theta^\theta, T_1) \\ \Pi(\theta) & \propto r_1^{\overline{\theta}-2} \times r_2^{\overline{\theta}-3}, \end{cases}$$

We set: $\mathcal{A}'_\theta = \{d \in \mathcal{A} : \overline{\theta} + 1 < \overline{d}\}$ and $\mathcal{B}'_\theta = \{d \in \mathcal{B} : \overline{\theta} > \overline{d} + 1\}$. The parameters of the prior model are : $m_\theta^d = \alpha(1 + 0.4\mathbb{1}_{\mathcal{A}_\theta}(d) + 0.2\mathbb{1}_{\mathcal{A}'_\theta}(d) - 0.4\mathbb{1}_{\mathcal{B}_\theta}(d) - 0.2\mathbb{1}_{\mathcal{B}'_\theta}(d))$, $T_1 = 40$ and $T_2 = 10$. The parameters r_1 and r_2 are chosen to produce a progressive allocation in the range of doses. The goal is to obtain a prior Π that is non-informative with the purpose of compensating the weight given to the highest dose combinations by the first non-toxic observations. For $\alpha = 0.25$, we choose: $r_1 = 0.942724$ and $r_2 = 0.95566$. In order to explore the diagonal going from $(1, 1)$ to (I, J) before a DLT is observed, a very small weight (10^{-5}) is added to these doses. Thus, the poSPM follows the sequence of dose combinations $(1, 1)$, $(1, 2)$, $(2, 2)$, $(2, 3)$, $(3, 3)$, ... until a first DLT is observed. After N observations, the final recommendation is made by using our current estimator $\hat{\theta}_N$, which corresponds to the most probable MTD according to the posterior Π_N.

S5.3 poSPMc

We use the same notation as for the poSPM calibration. The poSPMc model fulfills Assumption 4. The allocation strategy defined by Equation

(4.3) is used. The family of $I \times J$ matrix $(m_\gamma)_{\gamma \in C}$ is associated with the prior model $(\Lambda_\gamma)_{\gamma \in C}$. The rank of a contour γ, noted $\overline{\gamma}$, is the number of dose combination in $\mathcal{B}_\gamma$. T is the dispersion parameter of the truncated beta distributions, r_1 and r_2 are two positive values used for the calibration of the prior Π.

$$\begin{cases} \Lambda_\gamma & \sim \prod_{d \in \mathcal{B}_\gamma} \mathsf{B}_{[0,\alpha]}(m_\gamma^d, T) \times \prod_{d \in \mathcal{A}_\gamma} \mathsf{B}_{[\alpha,1]}(m_\gamma^d, T_1), \\ \Pi(\gamma) & \propto r_1^{\overline{\gamma}-2} \times r_2^{\overline{\gamma}-3}, \end{cases}$$

The parameters of the prior model are: $m_\gamma^d = \alpha(1+0.5\mathbb{1}_{\mathcal{A}_\gamma}(d)+0.25\mathbb{1}_{\mathcal{A}_\gamma \setminus \mathcal{M}_\gamma}(d)-0.4\mathbb{1}_{\mathcal{B}_\gamma}(d)-0.2\mathbb{1}_{\mathcal{B}_\gamma \setminus \mathcal{M}_\gamma}(d))$ and $T = 25$. The parameters r_1 and r_2 are chosen to produce a progressive allocation in the range of dose combinations. The goal is to obtain a prior Π that is non informative with the idea of compensating the weight given to the highest contour in term of rank by the the first non-toxic observations. For $\alpha = 0.2$, we choose $r_1 = 0.8739592$ and $r_2 = 0.9749345$ and for $\alpha = 0.3$, we choose $r_1 = 0.8117365$ and $r_2 = 0.950334$. For the exclusion rules (End of section 4.2), we choose to exclude a dose when the probability of our posterior on the upper interval attains $95\% : \delta = 0.95$. An $\epsilon = 10^{-5}$ is added to the part multiplying $1/k_d$ such that the value k_d drives the allocation between dose combinations for which we do not yet posses any observations. After N observations, we

recommend all the dose combinations in the estimated minimal set $\mathcal{M}_{\hat{\gamma}_N}$, except the doses for which we have less than 2 observations and the doses in H'_n which are considered as toxic:

$$H'_n = \{d \in D : \mathbb{P}_{\mathcal{U}}\left[P_d > \alpha + 0.05 | (n_d, n_d^1)\right] > 0.9\}.$$

S6. Computation of minimal set

In the two dimensional case, I and J are the number of doses for each of the agents. The set of contours could be described as all the polygonal chains tracing out a path from $(0.5, J + 0.5)$ to $(I + 0.5, 0.5)$ with steps of size 1 along the abscissae and of size -1 on the ordinate axis. Only rightward and downward steps are permitted. This set is also equivalent to all the combinations of I among $I + J$ which could be generated with the R function `combn` of package `combinat`. From there, it is easy to obtain matrices indicating which dose combinations are below or above the contour c that correponds to the sets $\mathcal{A}_c$ and $\mathcal{B}_c$.

Computing the minimal set is equivalent to computing $\min \mathcal{A}_c$ and $\max \mathcal{B}_c$. For the general problem "computing the maximum of a partially ordered set" the greatest calculation burden is $\mathcal{O}(n^2)$: in a set where no element are a priori ordered, all the elements have to be compared to solve

the problem. In the case of two drugs, the structure of $\mathbb{N}^2$ leads to a time complexity $\mathcal{O}(n)$. Let $\text{mat}_{\mathcal{B}_c}$ be a $I \times J$ matrix of 0 and 1 indicating the dose combinations in $\mathcal{B}_c$. From row N to 1:

1. Select the last element equal to 1 of the row in $\text{mat}_{\mathcal{B}_c}$, if it exists. Otherwise, go to next row.

2. Compare the selected element to the previous one. Keep it, if they are not ordered. Otherwise, delete it.

Finally, the selected elements are the maximum set of $\mathcal{B}_c$.

Note that in order to save time, all the arrays containing the minimal set associated to each contour can be saved for use in simulations. If we do so, the $\mathcal{O}(n^2)$ solution (exhaustive comparison two by two) could be used as it is only applied one time.

S7. Proofs of two properties and a general theorem

S7.1 Coherence

When there is no confusion, we write $\Lambda_\theta(dP_r)$ for $\Lambda^r_\theta(dP_r)$, with $r \in D$ We then state a stochastic partial ordering assumption on the prior-model.

Assumption 1. *Let d and d' be two doses such that $d < d'$. For all marginal*

$r \in D$, the posterior $\Lambda^r_{d,n}$ is stochastically greater than $\Lambda^r_{d',n}$:

$$\Lambda^r_{d,n}([0,x]) \leq \Lambda^r_{d',n}([0,x]) \, \forall x \in [0,1].$$

This assumption is satisfied by the calibration of poSPM presented in Section S5.2.

Proposition 1. *Under Assumptions 3 and 1 (supp. material), the poSPM is coherent.*

Proof. Suppose that $Y_{n+1} = 1$. The case $Y_{n+1} = 0$ can be solved in the same way. By construction, we have

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

Furthermore,

$$\begin{aligned} \int P_r \Lambda_{\theta,n}(dP) &= \int \left[\int \mathbb{1}_{\{0 \leq x \leq P_r\}} \mu(dx) \right] \Lambda_{\theta,n}(dP_r) \\ &= \int \Lambda^r_{\theta,n}(]x,1]) \mu(dx). \end{aligned}$$

If $\hat{\theta}_n = r$, then for all $\theta \in D$, $\Pi_n(\theta) \leq \Pi_n(r)$. Let $t > r$. According to Assumption 1, we know that $\Lambda^r_{r,n}(.)$ is stochastically greater than $\Lambda^t_{t,n}(.)$, *i.e.*

$$\int \Lambda^r_{r,n}(]x,1]) \mu(dx) \geq \int \Lambda^r_{t,n}(]x,1]) \mu(dx),$$

that is

$$\int P_r \Lambda_{r,n}(dP) \geq \int P_r \Lambda_{t,n}(dP).$$

Finally $\Pi_{n+1}(r) \geq \Pi_{n+1}(t)$, which ends the proof. □

S7.2 Conjugacy property

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

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

Proof. We have the following proportionnality relations:

$$\begin{aligned}\tilde{\Pi}_n(d) &\propto \frac{1}{\overline{\mathcal{V}}_d} \sum_{c \in \mathcal{V}_d} \Pi_n(c) \propto \frac{1}{\overline{\mathcal{V}}_d} \sum_{c \in \mathcal{V}_d} \left[\int L_n(P) \Lambda_c(dP) \right] \times \frac{\Pi(c)}{\sum\limits_{c \in \mathcal{V}_d} \Pi(c)} \times \sum_{c \in \mathcal{V}_d} \Pi(c) \\ &\propto \left[\int L_n(P) \sum_{c \in \mathcal{V}_d} r_c^d \times \Lambda_c(dP) \right] \times \left(\frac{1}{\overline{\mathcal{V}}_d} \sum_{c \in \mathcal{V}_d} \Pi(c) \right) \propto \left[\int L_n(P) \tilde{\Lambda}_d(dP) \right] \times \tilde{\Pi}(d).\end{aligned}$$

□

S7.3 Proof of Theorem 1

Proof. The proof is made by showing that the two first statements are equivalent to the third one. As the proofs of these two equivalences are very similar, we focus on: $(i) \Leftrightarrow (iii)$. The implication $(i) \Leftarrow (iii)$ is immediate

by using the Law of Large Number for the frequentist estimators n_d^1/n_d at each dose of the minimal set. We show $(i) \Rightarrow (iii)$ by using a reductio ad absurdum argument.

T is a scenario whose the maximum tolerated dose is MTD_T and the maximum tolerated contour is MTC_T. We choose this scenario T and a dose $d_0 \in MTC_T$ such that there exists a set of sequences of our sample, $A = \{a = (x_1^\infty, y_1^\infty)\}$ satisfying: (1) $\forall\, a = (x_1^\infty, y_1^\infty) \in A\,,\ F(x_1^n, y_1^n) \rightarrow MTD_T$, (2) $\mathbb{P}_T(A) > 0$, (3) $\forall\, a \in A\,,\ d_0 \notin \overset{\infty}{\underset{N=1}{\cap}} \overset{\infty}{\underset{n=N}{\cup}} \{x_n\}$. This is possible as the absence of (3) for any dose contradicts the negation of (iii) and the absence of (1) and (2) together contradicts (i). The point (3) implies that there exists $N_0 \in \mathbb{N}$ such that:

$$\mathbb{P}_T(B = \{s \in A : \nexists\, n > N_0\,,\ x_n = d_0\}) > 0.$$

We introduce the two following sets: $B' = \{\omega \in \Omega : \exists a = (x_1^\infty, y_1^\infty) \in A\,,\ (X_{N_0}^\infty, Y_{N_0}^\infty)(\omega) = (x_{N_0}^\infty, x_{N_0}^\infty)\}$ and $B'' = \{\omega \in \Omega : \exists a = (x_1^\infty, y_1^\infty) \in A\,,\ (X_1^{N_0}, Y_1^{N_0})(\omega) = (x_1^{N_0}, x_1^{N_0})\}$. We have: $\mathbb{P}_T(B) = \mathbb{P}_T(B'|B'') \times \mathbb{P}_T(B'')$ and then $\mathbb{P}_T(B'|B'') > 0$. Let T' be a scenario such that, for all $d \in D \backslash \{d_0\}$, $P_{T'}(Y = 1|X = d) = P_T(Y = 1|X = d)$ and $MTD_{T'} = d_0$. We then have: $\mathbb{P}_{T'}(B'|B'') > 0$, as the same method chooses the next dose being used for the two scenarios and the dose is not tried after N_0 for any sequence of B.

Moreover, $B'' = \underset{(x_1^\infty, y_1^\infty) \in A}{\cup} \{\omega \in \Omega : (X_1^{N_0}, Y_1^{N_0})(\omega) = (x_1^{N_0}, x_1^{N_0})\}$ with:

$$\mathbb{P}_{T'}(\{\omega \in \Omega : (X_1^{N_0}, Y_1^{N_0})(\omega) = (x_1^{N_0}, x_1^{N_0})\}),$$

as the constraints are on the first N_0 terms. Thus, $\mathbb{P}_{T'}(B) > 0$ and on this set of sequences of the sample the statistic F converges to $MTD_T \neq d_0 = MTD_{T'}$.

□

S8. Proof of asymptotic results for poSPM

S8.1 Proof of Theorem 2 (a)

This assumption leans upon the regularity of the prior model.

Assumption 2. *The following conditions are valid except when Λ_θ^θ is a Dirac measure.*

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

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

$$\forall\, P_d \in S_\theta^d\,,\; s < \lambda_\theta^d(P_d) < S\,.$$

Note that this assumption can be used for the poSPMc by replacing θ by γ, the parameter of the contour.

The second point is only useful for the sake of the demonstration when some $\beta_{(i,j)}$ are equal to 0 or 1. In order to establish asymptotic properties for the poSPM, we introduce the set $\tilde{D}$ to indicate those doses that are observed infinitely often:

$$d \in \tilde{D} \Leftrightarrow n_d \underset{n\to\infty}{\longrightarrow} \infty.$$

Proof. Let us start with the proof of ε-sensitivity. In this proof, we are interested in the asymptotic behavior of poSPM, that is why we are able to ignore those doses tested only a finite number of times and we can reason as though they had never been used. Specifically, the doses in this proof are always considered to be in $\tilde{D}$. We assume that $\mathcal{E}(I, P^T)$ is not empty (P^T denotes the true scenario). Let $r \in \tilde{D} \setminus \mathcal{E}(I, P^T)$. We can distinguish two cases. The first case is the existence of a dose $d \in \mathcal{E}(I, P^T)$ such that d is ordered with r. We are then reduced to the SPM in the case of total ordering and the proof can be found in Clertant and O'Quigley (2017).

The second case is where there exists no dose in $\mathcal{E}(I, P^T)$ ordered with r. So there exists a dose $d \in \mathcal{E}(I, P^T)$ not ordered with r, because we assume that $\mathcal{E}(I, P^T)$ is not empty. We want now to compare the integrals

$I_{n,r}$ and $I_{n,d}$, where $g(p,n,m) = p^n(1-p)^m$.

$$\frac{I_{n,r}}{I_{n,d}} = \prod_{k\in\tilde{D}} \frac{M^k_{n,r}}{M^k_{n,d}} = \prod_{k\in\tilde{D}} \frac{\int_{S^k_r} g\left(P_k, n^1_k, n^0_k\right) \Lambda_r\left(dP_k\right)}{\int_{S^k_d} g\left(P_k, n^1_k, n^0_k\right) \Lambda_d\left(dP_k\right)} \tag{S8.1}$$

$$= \prod_{k\in\tilde{D}} \frac{\int_{S^k_r} g\left(P_k, n^1_k, n^0_k\right) \lambda_r\left(P_k\right) dP_k}{\int_{S^k_d} g\left(P_k, n^1_k, n^0_k\right) \lambda_d\left(P_k\right) dP_k} \tag{S8.2}$$

$$\leq \frac{S}{s} \prod_{k\in\tilde{D}} \frac{\int_{S^k_r} g\left(P_k, n^1_k, n^0_k\right) dP_k}{\int_{S^k_d} g\left(P_k, n^1_k, n^0_k\right) dP_k}, \tag{S8.3}$$

where Equation (S8.1) follows from Assumption 3, Equation (S8.2) from Assumption 2 *(a)* and Inequation (S8.3) from Assumption 2 *(b)*. We are able to state the following property. For all functions f that are continuous on $[0,1]$, we have

$$\int f\left(P_k\right) \frac{g(P_k, n^1_k, n^0_k)}{\text{Beta}(n^1_k+1, n^0_k+1)} dP_k \underset{n_k\to\infty}{\longrightarrow} \int f\left(P_k\right) \mathbb{1}_{\{P_k\}}\left(P_k\right) \gamma\left(dP_k\right) = f\left(P_k\right), \tag{S8.4}$$

where Beta(.) denotes the Beta function and γ the counting measure. Let then $k \in \tilde{D}$. We are looking for the behavior, when n_k increases without bound, of

$$R_k = \frac{\int_{S^k_r} g\left(P_k, n^1_k, n^0_k\right) dP_k}{\int_{S^k_d} g\left(P_k, n^1_k, n^0_k\right) dP_k}.$$

The case $S^k_r = S^k_d$ is straightforward. For the other cases, we use the convergence expressed in Equation (S8.4). As d is not ordered with r, with Assumption **??**, we have six cases left to deal with.

- $S^k_r = A$ or B or I and $S^k_d = [0,1]$. Then $R_k \underset{n_k\to\infty}{\longrightarrow} \mathbb{1}_A\left(P^T_k\right)$ or $\mathbb{1}_B\left(P^T_k\right)$

or 0. This last result is due to the fact that $r \notin \mathcal{E}(I, P^T)$.

- $S_r^k = [0,1]$ and $S_d^k = I$. Then $R_k \underset{n_k \to \infty}{\longrightarrow} 1$, because $d \in \mathcal{E}(I, P^T)$.

- $S_r^k = [0,1]$ and $S_d^k = A$. Then $R_k \underset{n_k \to \infty}{\longrightarrow} 1/\mathbb{1}_A\left(P_k^T\right) = 1$. As there is no couple of ordered dose in $\mathcal{E}(I, P^T)$, P_k^T belongs to A.

- $S_r^k = [0,1]$ and $S_d^k = B$. With the same argument as the previous case, we have $R_k \underset{n_k \to \infty}{\longrightarrow} 1$.

Finally, going back to Inequality (S8.3), we conclude that $I_{n,r}/I_{n,d}$ tends to 0 when n increases without bound. This leads to a contradiction because, as $r \in \tilde{D}$, this ratio is greater than 1 infinitely often. So $\tilde{D} \subset \mathcal{E}(I, P)$, that is, the poSPM is ε-sensitive.

We prove now that poSPM is also balanced. Assumption 3 allows us to focus on the marginal ratio

$$\frac{M_{n,r}^k}{M_{n,t}^k} = \frac{\int g(P_j, n_k^1, n_k^0)\ \Lambda_r^k(dP_k)}{\int g(P_k, n_k^1, n_k^0)\ \Lambda_t^k(dP_k)} \quad \text{and} \quad \frac{I_{n,r}}{I_{n,t}} = \prod_{k \in D} \frac{M_{n,r}^k}{M_{n,t^k}}.$$

Assumption 2 involves:

$$d(P_k^T, S_r^k) = d(P_k^T, S_t^k) \Longrightarrow 0 < \liminf_{n \to \infty} \frac{M_{n,r}^k}{M_{n,t}^k} \leq \limsup_{n \to \infty} \frac{M_{n,r}^k}{M_{n,t}^k} < +\infty,\ a.s. \tag{S8.5}$$

The ε-balanced behavior corresponds to the case where, forall $k \in D$, $P_k^T \notin$

I. We show that

$$r \notin \mathcal{M}_{c^\star} \Longrightarrow \mathbb{P}(\{n_r \to \infty\}) = 0 \tag{S8.6}$$

By symmetry we can choose $r \in D^-$. The set $\mathcal{M}_{c^\star} \cap D^- \cap \mathcal{A}_r$ is not empty. Let t be a dose in $\mathcal{M}_{c^\star} \cap D^- \cap \mathcal{A}_r$. By using Equation (S8.5), as $\mathcal{A}_t \subset \mathcal{A}_r$ and $S_t^k = [0,1]$ when $k \in \mathcal{C}_t$, we have

$$\forall k \in A_t \cup C_t,\ \mathbb{P}\left(\limsup_{n\to\infty} \frac{M_{n,r}^k}{M_{n,t}^k} < \infty \ \middle|\ n_k \to \infty \right) = 1.$$

If $k \in (\mathcal{B}_r \cup \mathcal{C}_r) \cap \mathcal{B}_t$ then $P_k^T \in B$ and we obtain the same result

$$\mathbb{P}\left(\limsup_{n\to\infty} \frac{M_{n,r}^k}{M_{n,t}^k} < \infty \ \middle|\ n_k \to \infty \right) = 1.$$

We consider now the ratios $M_{n,r}^k / M_{n,t}^k$ when $k \in (\mathcal{B}_t \cap \mathcal{A}_r) \cup \{t, r\}$. In that case, we have $P_k^T \in B$, $S_t^k = B$ and S_r^k equals to I or A. By using Proposition 1, we have

$$\forall k \in (\mathcal{B}_t \cap \mathcal{A}_r) \cup \{t, r\},\ \mathbb{P}\left(\lim_{n\to\infty} \frac{M_{n,r}^k}{M_{n,t}^k} = 0 \ \middle|\ n_k \to \infty \right) = 1.$$

We then have

$$\mathbb{P}\left(\lim_{n\to\infty} \frac{I_{n,r}}{I_{n,t}} = 0 \ \middle|\ n_r \to \infty \right) = 1 \text{ and } \mathbb{P}\left(\left\{ \lim_{n\to\infty} \frac{I_{n,r}}{I_{n,t}} = 0 \right\} \cap \{n_r \to \infty\} \right) = 0,$$

which proves Equation (S8.6). We achieve the proof of the ε-balanced property by showing that

$$t \in \mathcal{M}_{c^\star} \Longrightarrow \mathbb{P}(n_t \to \infty) = 1 \tag{S8.7}$$

By using Equation (S8.6), we have, for all r and t in $\mathcal{M}_{c^\star}$

$$\forall k \in D,\ k \neq r \text{ and } k \neq t,\ \mathbb{P}\left(\limsup_{n\to\infty} \frac{M_{n,r}^k}{M_{n,t}^k} < \infty\right) = 1.$$

Moreover, for all $r \in \mathcal{M}_{c^\star}$, P_r^T is included in S_t^r. As $P_r^T \notin I = S_r^r$, we have

$$\forall r \in \mathcal{M}_{c^\star},\ \mathbb{P}\left(\lim_{n\to\infty} \frac{M_{n,r}^r}{M_{n,t}^r} = 0 \;\middle|\; n_r \to \infty\right) = 1.$$

Let E_r be the event $\{n_r \to \infty\} \cap \{n_t \to \infty\}^c$. Then,

$$\mathbb{P}\left(\lim_{n\to\infty} \frac{I_{n,r}}{I_{n,t}} = 0 \;\middle|\; E_r\right) = 1 \text{ and } \mathbb{P}\left(\left\{\lim_{n\to\infty} \frac{I_{n,r}}{I_{n,t}} = 0\right\} \cap E_r\right) = 0,$$

As $\mathbb{P}(\sum_{\mathcal{M}_{c^\star}} n_k \to +\infty) = 1$, we have $\mathbb{P}(\cup_{\mathcal{M}_{c^\star}} E_r) = 1$, which proves Equation ($S8.7$) and ends the demonstration of the poSPM ε-balanced behavior. □

S8.2 Proof of Theorem 2 (b)

Proof. We set:

$$I_{n,c} = \prod_{k\in\tilde{D}} M_{n,c}^k = \prod_{k\in\tilde{D}} \int_{S_c^k} g\left(P_k, n_k^1, n_k^0\right) \Lambda_c\left(dP_k\right)$$

We note that the regularity assumption 2 involves:

$$\delta(P_k^T, S_r^k) = \delta(P_k^T, S_t^k) \Longrightarrow 0 < \liminf_{n\to\infty} \frac{M_{n,r}^k}{M_{n,t}^k} \leq \limsup_{n\to\infty} \frac{M_{n,r}^k}{M_{n,t}^k} < \infty,\ a.s.$$

We will show the following assertion

$$c \in C \setminus c^* \Longrightarrow \mathbb{P}(n_c \to \infty) = 0 \tag{S8.8}$$

As $c \neq c^*$, we have $\mathcal{M}_c \neq \mathcal{M}_{c^\star}$ and there exists $k \in \mathcal{M}_c$ such that P_k^T is included in $S_{c^*}^k$ and not in S_c^k. By using Proposition 1, we have

$$\mathbb{P}\left(\lim_{n\to\infty} \frac{M_{n,c}^k}{M_{n,c^*}^k} = 0 \middle| n_c \to \infty\right) = \mathbb{P}\left(\lim_{n\to\infty} \frac{M_{n,c}^k}{M_{n,c^*}^k} = 0 \middle| n_k \to \infty\right) = 1,$$

where the first equality arises from Assumption 5(b). For all doses $d \in D$, the distribution Λ_{c^*} models correctly the probability of toxicity, in other words $P_d^T \in S_{\mathcal{C}^*}$. We then have

$$\mathbb{P}\left(\lim_{n\to\infty} \frac{I_{n,c}}{I_{n,c^*}} = 0 \middle| n_c \to \infty\right) = 1 \text{ and } \mathbb{P}\left(\left\{\lim_{n\to\infty} \frac{I_{n,c}}{I_{n,c^*}} = 0\right\} \cap \{n_c \to \infty\}\right) = 0,$$

which ends the proof. (S8.8). □

S9. A general bayesian property

The following property is used in the proof of asymptotic results. It has already been proved in (Clertant and O'Quigley, 2018). Here, for the convenience of the reader, we reproduce this property without its proof.

The couple $(\Omega, \mathcal{A})$ denotes an abstract space endowed with its σ-field. We denote by I a finite set and by $(X_k)_{k\in\mathbb{N}}$ a sequence of independent random variables taking their values in I. Let F be the set of functions from I to the segment $[0, 1]$. For any element $q \in F$, q_i denotes its value at $i \in I$. Let $S = \{q \in F : \sum_{i\in I} q_i = 1\}$ be the probability space on which we want to work. We say that a random variable X follows the distribution

q if $\mathbb{P}\{X = i\} = q_i$, for all $i \in I$. 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 topological supports of Λ_1 and Λ_2 respectively. We would like to know the asymptotic behavior of the ratio of the expected likelihood under Λ_1 on 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 $n_i = \sum_{k=1}^{n} \mathbb{1}_{\{X_k = i\}}$, for $i \in I$. We assume that, under the true probability β, the random variables $X_k,\ k \in \mathbb{N}$ are identically distributed. The convergence of $r(\Lambda_1, \Lambda_2, n)$ depends mainly on the localization of β compared to the supports S_1 and S_2. To deal with this problem, 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 the conventions $\log 0 = -\infty$ and $0 \times -\infty = 0$. We suppose that β is closer to S_1 than S_2 in terms of entropy.

Assumption 3. *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 $r(\Lambda_1, \Lambda_2, n)$.

Proposition 1. *Under Assumption 3, we have $r(\Lambda_1, \Lambda_2, n) \underset{n \to \infty}{\longrightarrow} 0$.*

Bibliography

D. Azriel, M. Mandel, and Y. Rinott (2011). The treatment versus experimentation dilemma in dose finding studies. *Journal of Statistical Planning and Inference*, 141(8):2759 – 2768.

T. M. Braun and S. Wang (2010). A hierarchical bayesian design for phase i trials of novel combinations of cancer therapeutic agents. *Biometrics*, 66(3):805–812.

K. Cheung (2005). Coherence principles in dose-finding studies. *Biometrika*, 92(4):863–873.

Y. K. Cheung (2011). *Dose finding by the continual reassessment method*. CRC Press.

M. Clertant and J. O'Quigley (2017). Semiparametric dose finding methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(5):1487–1508.

M. Clertant and J. O'Quigley (2018). The role of minimal sets in dose finding studies. *Journal of Nonparametric Statistics*, 30(4):1016–1031.

M. Clertant and J. O'Quigley (2018). Semiparametric dose finding methods:

special cases. *Journal of the Royal Statistical Society: Series C (Statistical Methodology)*, 68(2):271–288.

X. Huang, S. Biswas, Y. Oki, J. Issa, and D. A. Berry (2007). A parallel phase i/ii clinical trial design for combination therapies. *Biometrics*, 63(2):429–436.

Alexia Iasonos and John OQuigley (2016). Dose expansion cohorts in phase i trials. *Statistics in biopharmaceutical research*, 8(2):161–170.

Yuan Ji, Ping Liu, Yisheng Li, and B Nebiyou Bekele (2010). A modified toxicity probability interval method for dose-finding trials. *Clinical Trials*, pages 653–663.

A. P. Mander and M. J. Sweeting (2015). A product of independent beta probabilities dose escalation design for dual-agent phase i trials. *Statistics in Medicine*, 34(8):1261–1276.

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

P. F. Thall, R. E. Millikan, P. Mueller, and S. Lee (2003). Dose-finding with two agents in phase i oncology trials. *Biometrics*, 59(3):487–496.

N. A. Wages, M. R. Conaway, and J. O'Quigley (2011). Continual reassessment method for partial ordering. *Biometrics*, 67(4):1555–1563.

K. Wang and A. Ivanova (2005). Two-dimensional dose finding in discrete dose space. *Biometrics*, 61(1):217–222.