A Bayesian adaptive design in cancer phase I/II trials with drug combinations using escalation with overdose control (EWOC) and adaptive randomization

Sungjin Kim[1], José L. Jiménez[2], and Mourad Tighiouart[1]
[1]Biostatistics and Bioinformatics Research Center, Cedars-Sinai Medical Center, Los Angeles, California, U.S.A. [2]Politecnico di Torino, Turin, Italy

Abstract
The use of drug combinations in clinical trials is increasingly common during the last years since a more favorable therapeutic response may be obtained by combining drugs that, for instance, target multiple pathways or inhibit resistance mechanisms. However, most of the existing methodology in phase I trials recommends a single maximum tolerated dose (MTD), which may result in a failed phase II since other MTDs may present higher treatment efficacy for the same level of toxicity. We are motivated by a phase I/II trial that combines cisplatin with cabazitaxel for patients with prostate cancer with visceral metastasis. We present a Bayesian adaptive phase I/II design with drug combinations where a binary dose limiting toxicity (DLT) is used for dose escalation in stage 1 and a time to event endpoint is used for dose allocation in stage 2. The overall goal is to estimate the dose combination region associated with the highest median time to progression (TTP) among doses along the MTD curve. Conditional escalation with overdose control (EWOC) is used in stage 1 to allocate dose combinations to subsequent cohorts of patients and estimate the MTD. Stage 2 starts by allocating a first cohort of patients to dose combinations equally spaced along the MTD curve, and then allocates subsequent cohorts of patients to dose combinations likely to have high posterior median TTP using adaptive randomization. We perform extensive simulation studies to evaluate the operating characteristics of our method.

Keywords: Adaptive randomization; Continuous dose; Drug combination; Escalation with overdose control; Phase I/II trial

1. Introduction

In cancer phase I/II clinical trials, the main goal is to identify a safe dose that maximizes the treatment efficacy. In single-agent settings with binary or time to event endpoints where efficacy is observed relatively fast (e.g. one or two cycles of therapy), one-stage sequential designs where the joint probability of toxicity and efficacy is sequentially updated after each cohort of patients are usually employed (see e.g. [9, 14, 2, 6, 13, 3, 11] for binary endpoints, and [21] for time to event endpoints). This methodology has been extended to accommodate combination of drugs of any kind (see e.g. [5, 20] for binary endpoints and [21] for time to event endpoints), and proceed in a similar fashion as the methods referenced for single-agent.

In cases where efficacy is not ascertained in a short period of time, it is frequent to employ two-stage designs, where a maximum tolerated dose (MTD) set is first selected, and then tested for efficacy in a second stage with possibly a different population of patients than the one used in the first stage. This approach has been discussed by [10, 7, 4]. For drug combination trials, methodology for these type of two-stage designs have been proposed for binary efficacy endpoints (see e.g. [12, 15]).

One characteristic that most of these methods have in common is that they only recommend a single MTD either at the end of the phase I trial, or at the end of the first stage in a phase I/II trial. However, even if the recommended dose that will be tested for efficacy is indeed a valid MTD, there could be another MTD with higher efficacy, making the MTD recommended in the first place non-optimal.

In this article, we extend the work in [15] by proposing a two-stage design for drug combinations trials with binary DLT endpoint in the first stage, time to event efficacy endpoint in the second stage

and continuous dose levels in both stages. In the first stage, the dose finding method proposed by [16] is used to estimate the MTD curve. In the second stage, a Bayesian adaptive design that starts allocating a first cohort of patients to dose combinations equally spaced along the MTD curve, and then allocates subsequent cohorts of patients to dose combinations likely to have high posterior median TTP using adaptive randomization. To allow for different shapes in the median TTP curve, we employ a flexible family of cubic splines to model the dose - median TTP relationship. Adaptive randomization is sequentially used after a pre-defined time period to minimize the number of patients allocated at sub-therapeutic dose levels. At the end of the trial, the dose combination within the MTD with the highest *a posteriori* median TTP is selected and recommended for further phase II or III studies.

The manuscript is organized as follows. In section 2, we review the first stage of the proposed phase I/II trial previously described in [16, 15]. In section 3, we describe the second stage of the design. In section 4, we illustrate the methodology with the phase I/II drug combination trial of cisplatin and cabazitaxel in patients with prostate cancer with visceral metastasis where time to progression is a secondary endpoint. The goal in this trial is to find a tolerable dose combination with the highest TTP median. A discussion of the approach and final remarks are included in Section 5.

2. Phase I/II Trial: Stage 1

2.1 Model

Following [16], consider the generic form of a dose-toxicity model

$$P(T = 1|x, y) = F(\eta_0 + \eta_1 x + \eta_2 y + \eta_3 xy), \tag{1}$$

where $T = 1$ represents an observed DLT at the dose combination (x, y), $T = 0$ otherwise, $x \in [X_{\min}, X_{\max}]$ is the dose level of agent A, $y \in [Y_{\min}, Y_{\max}]$ is the dose level of agent B and $F(.)$ is a known cumulative distribution function. We assume that the dose combinations are continuous and standardized to be in the interval [0,1], the interaction parameter $\eta_3 > 0$, and $\eta_1, \eta_2 > 0$ in order to guarantee that the probability of DLT increases with the dose of any agent when the other one is held constant.

The MTD is defined as any dose combination (x^*, y^*) such that

$$P(T = 1|x^*, y^*) = \theta. \tag{2}$$

As described in [16], we reparameterize equation (1) as follows. Let ρ_{10}, the probability of DLT when the levels of drugs A and B are 1 and 0, respectively, ρ_{01}, the probability of DLT when the levels of drugs A and B are 0 and 1, respectively, ρ_{00}, the probability of DLT when the levels of drugs A and B are both 0. Hence, it is possible to show that MTD takes the form

$$C = \left\{(x^*, y^*) : y^* = \left[(F^{-1}(\theta) - F^{-1}(\rho_{00})) - (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))x^*\right] \div \left[(F^{-1}(\rho_{01}) - F^{-1}(\rho_{00}))\eta_3 x^*\right]\right\}. \tag{3}$$

We assume that ρ_{10}, ρ_{01} and η_3 are independent *a priori* with $\rho_{01} \sim \text{beta}(a_1, b_1)$, $\rho_{10} \sim \text{beta}(a_2, b_2)$, and conditional on (ρ_{01}, ρ_{10}), $\rho_{00}/\min(\rho_{01}, \rho_{10}) \sim \text{beta}(a_3, b_3)$. The prior distribution on the interaction parameter η_3 is a gamma distribution with mean a/b and variance a/b^2. Let $D_n = \{(x_i, y_i, T_i)\}$ be the data gathered after enrolling n patients. The posterior distribution of the model parameters is

$$\begin{aligned} \pi(\rho_{00}, \rho_{10}, \rho_{01}, \eta_3) \propto & \prod_{i=1}^{n} G((\rho_{00}, \rho_{10}, \rho_{01}, \eta_3; x_i, y_i))^{T_i} \\ & \times (1 - G(\rho_{00}, \rho_{10}, \rho_{01}, \eta_3; x_i, y_i))^{1-T_i} \\ & \times \pi(\rho_{01})\pi(\rho_{10})\pi(\rho_{00}|\rho_{01}, \rho_{10})\pi(\eta_3), \end{aligned} \tag{4}$$

where

$$G(\rho_{00},\rho_{10},\rho_{01},\eta_3;x_i,y_i) = F(F^{-1}(\rho_{00}) + (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))x_i + (F^{-1}(\rho_{01}) - F^{-1}(\rho_{00}))y_i + \eta_3 x_i y_i). \tag{5}$$

Note that the operating characteristics of this stage are evaluated using informative prior distributions (see [15]).

2.2 Trial Design

Dose escalation / de-escalation proceeds using the same algorithm described in [16]. It is based on escalation with overdose control (EWOC), where after each cohort of enrolled patients, the posterior probability of overdosing the next cohort of patients is bounded by a feasibility bound α, see e.g. [1, 18, 19, 17]. In a cohort with two patients, the first one would receive a new dose of agent A given that the dose y of agent B was previously assigned. The other patient would receive a new dose of agent B given that dose x of agent A was previously assigned. Using EWOC, these new doses are at the α-th percentile of the conditional posterior distribution of the MTDs. The algorithm continues until the maximum sample size is reached or until the trial is stopped for safety. A detailed description of this algorithm can be found in [16]. At the end of the trial, the MTD curve is estimated as

$$C_{\text{est}} = \left\{(x^*,y^*) : y^* = \left[(F^{-1}(\theta) - F^{-1}(\widehat{\rho}_{00})) - (F^{-1}(\widehat{\rho}_{10}) - F^{-1}(\widehat{\rho}_{00}))x^*\right] \div \left[(F^{-1}(\widehat{\rho}_{01}) - F^{-1}(\widehat{\rho}_{00}))\widehat{\eta}_3 x^*\right]\right\}, \tag{6}$$

where $\widehat{\rho}_{00},\widehat{\rho}_{10},\widehat{\rho}_{01},\widehat{\eta}_3$ are the posterior medians given the data D_n. This method has been extensively studied by [16] and hence we only present the operating characteristics in the context of the CisCab trial described in section 4.

3. Phase I/II Trial: Stage 2

3.1 Model

Let x be a dose of drug A such that $(x,y) \in C_{\text{est}}$. Also, assume that x is standardized to be in the interval [0,1]. We model the time to progression as a Weibull distribution with probability density function

$$f(t;x) = \frac{k}{\lambda(x;\boldsymbol{\psi})}\left(\frac{t}{\lambda(x;\boldsymbol{\psi})}\right)^{k-1}\exp\left(-\frac{t}{\lambda(x;\boldsymbol{\psi})}\right)^{k}. \tag{7}$$

The median TTP is

$$\text{Med}(x) = \lambda(x;\boldsymbol{\psi})(\log 2)^{\frac{1}{k}}. \tag{8}$$

A flexible way of modeling the median TTP along the MTD curve is through the use of the cubic spline function

$$\lambda(x;\boldsymbol{\psi}) = \exp\left(\beta_0 + \beta_1 x + \beta_2 x^2 + \sum_{j=3}^{k}\beta_j (x-\kappa_j)_+^3\right), \tag{9}$$

where $\boldsymbol{\psi} = (\boldsymbol{\beta},\boldsymbol{\kappa})$, with $\boldsymbol{\beta} = (\beta_0,\ldots,\beta_k)$ and $\boldsymbol{\kappa} = (\kappa_3,\ldots,\kappa_k)$, being $\kappa_3 = 0$. Let $D_m = \{(x_i,t_i,\delta_i), i = 1\ldots,m\}$ be the data after enrolling m patients in the trial where t represents the TTP or last follow-up, and δ the censoring status, and let $\pi(\boldsymbol{\psi},k)$ be the joint prior density on the parameter vector $\boldsymbol{\psi}$ and k. The posterior distribution is

$$\pi(\boldsymbol{\psi},k|D_m) \propto \pi(\boldsymbol{\psi},k)\prod_{i=1}^{m}\left[\frac{k}{\lambda(x_i;\boldsymbol{\psi})}\left(\frac{t_i}{\lambda(x_i;\boldsymbol{\psi})}\right)^{k-1}\right]^{\delta_i} \times \exp\left(-\frac{t_i}{\lambda(x_i;\boldsymbol{\psi})}\right)^{k}. \tag{10}$$

Let Med_x be the median TTP at dose combination x and let Med_0 be the median TTP of the standard of care treatment. We propose an adaptive design in order to test the hypothesis

$$\begin{aligned} &\text{H}_0 : \text{Med}_x \leq \text{Med}_0 \text{ for all } x \quad \text{vs.} \\ &\text{H}_1 : \text{Med}_x > \text{Med}_0 \text{ for some } x. \end{aligned} \tag{11}$$

3.2 Trial Design

1. We first treat n_1 patients at dose combinations $x_1, \ldots, x_{n_1}$, which are equally spaced along the estimated MTD curve C_{est}.
2. Obtain Bayes estimates, $\hat{\boldsymbol{\psi}}$ of $\boldsymbol{\psi}$ and $\hat{k}$ of k given the data D_{n_1} using equation (10).
3. Generate n_2 dose combinations from the standardized density $\hat{\text{Med}}(x) = \lambda(x; \hat{\boldsymbol{\psi}})(\log 2)^{\frac{1}{\hat{k}}}$ and assign them to the next n_2 patients.
4. Repeat steps 2 and 3 until a total of n patients have been enrolled to the trial subject to pre-specified stopping rules.

Decision Rule: At the end of the trial, we reject the null hypothesis if $\text{Max}_x\{P(\text{Med}(x; \boldsymbol{\psi}_i) > \text{Med}_0 | D_{n,i})\} > \delta_u$, where δ_u is a design parameter.

Stopping Rule (Futility): For ethical reasons and to avoid treating patient at sub-therapeutic dose levels, we will stop the trial for futility if there is strong evidence that none of the dose combinations are promising, i.e., $\text{Max}_x\{P(\text{Med}(x; \boldsymbol{\psi}_i) > \text{Med}_0 | D_{n,i})\} < \delta_0$, where δ_0 is a design parameter.

Stopping Rule (Efficacy): For ethical reasons, if the investigator considers there is enough evidence in favor of one or more dose combinations being tested, and no further patients need to be enrolled, the trial can be terminated if $\text{Max}_x\{P(\text{Med}(x; \boldsymbol{\psi}_i) > \text{Med}_0 | D_{n,i})\} > \delta_1$, where $\delta_1 \geq \delta_u$ is a study parameter and the dose combination $x^{\text{opt}} = \arg\max_v\{P(\text{Med}(v; \boldsymbol{\psi}_i) > \text{Med}_0 | D_{n,i})\}$ is selected for further randomized phase II or phase III clinical trials.

3.2.1 Design Operating Characteristics

We assess the operating characteristics of the proposed design by assuming that $\lambda(x; \boldsymbol{\psi})$ is a cubic spline with two knots placed between 0 and 1. This class of modeling is very flexible and is able to adapt to scenarios where the median TTP is either constant or skewed toward one of the edges. Vague priors are placed on the model parameters by assuming $\boldsymbol{\beta} \sim N(\boldsymbol{\mu}, \sigma^2 \boldsymbol{I}_6)$, where $\boldsymbol{\mu} = \{0, 0, 0, 0, 0, 0\}$ and $\sigma^2 = 100$, and $(\kappa_4, \kappa_5) \sim \text{Unif}\{(u, v) : 0 \leq u < v \leq 1\}$. Note that the parameters of the prior distribution of $\boldsymbol{\beta}$ are always the same regardless of the value of Med_0.

For each scenario favoring the alternative hypothesis, we estimate the Bayesian power, which is defined as

$$\text{Power} \approx \frac{1}{M} \sum_{i=1}^{M} I[\text{Max}_x\{P(\text{Med}(x; \boldsymbol{\psi}_i) > \text{Med}_0 | D_{n,i})\} > \delta_u], \tag{12}$$

where

$$P(\text{Med}(x; \boldsymbol{\psi}_i) > \text{Med}_0 | D_{n,i}) \approx \frac{1}{L} \sum_{j=1}^{L} I\left[\text{Med}(x; \boldsymbol{\psi}_{i,j}) > \text{Med}_0\right] \tag{13}$$

and $\boldsymbol{\psi}_{i,j}$ is the j-*th* MCMC sample for the i-*th* trial.

For scenarios favoring the null hypothesis, (12) is the estimated Bayesian type-I error probability. The optimal dose from the i-*th* trial is defined as

Table 1: True parameter values for $\rho_{00}, \rho_{01}, \rho_{10}, \eta$ and θ as well as safety results for the two simulated scenarios of the first stage where EWOC is employed.

	Scenario 1	Scenario 2
ρ_{00}	1e-5	1e-8
ρ_{01}	0.10	0.00005
ρ_{10}	0.10	0.00008
η	20	20
θ	0.33333	0.33333
Average number of DLTs	0.34	0.27
Number of trials with DLT rate $> \theta + 0.1$	7.30	0.00

$$x_i^{\text{opt}} = \arg\max_v \{P(\text{Med}(v; \boldsymbol{\psi}_i) > \text{Med}_0 | D_{n,i})\}. \quad (14)$$

We also report the estimated TTP median by replacing $\boldsymbol{\psi}$ in (8) by the average posterior median across all simulated trials. Last, we also report the mean posterior probability of $\text{Med}_x > \text{Med}_0$ for any dose combination x.

4. Application to the CisCab Phase I/II Trial

We illustrate the methods proposed in sections 2 and 3 with a phase I/II trial referred as the "CisCab trial" where TTP is a secondary endpoint of the trial. We are motivated by a phase I trial published by [8], that combines cisplatin and cabazitaxel in patients with advanced solid tumors, where the MTD was established at 15/75 mg/m^2. In a first part of this motivating trial, doses were escalated according to a standard "3+3" design and no DLTs were observed at dose combination which was found to be the MTD. During the second part of the trial, 15 additional patients were treated at the MTD and 2 DLTs were observed. In total, 18 patients were treated at the MTD.

Considering these results, there may be other active dose combinations that are tolerable and active in prostate cancer with visceral metastasis. The CisCab trial considers doses that range from 10 to 25 mg/m^2 for cabazitaxel, and from 50 to 100 mg/m^2 for cisplatin, that will be administered intravenously. In a first stage, the CisCab trial will enroll 30 patients in order to obtain the MTD curve. This stage of the design proceeds as explained in section 3, with a target probability of DLT $\theta = 0.33$, and a logistic link fuction $F(.)$ in equation (1). The starting dose combination for the first cohort of two patients is 15/75 mg/m^2, and DLTs are to be resolved within 1 cycle of treatment (3 weeks). Prior distributions are calibrated such that the prior mean probability of DLT at dose combination 15/75 mg/m^2 equals θ (see [15]). The operating characteristics of this first stage are obtained by simulating 1000 trials replicates following [16].

In Figures 1 and 2 we show the true and estimated MTD curves obtained with equation (6) respectively in two different scenarios. In the scenario presented in Figure 1, the true MTD curve passes through the dose combination 15/75 mg/m^2, whereas in the scenario presented in Figure 2, the true MTD curve is significantly above the dose combination 15/75 mg/m^2. In both scenarios the estimated MTD curves are very close to the true MTD curves. These results are supported the pointwise average bias shown in Figures 6 and 7 in Appendix. In these scenarios, the pointwise average bias fluctuates between -0.01 and 0.01, and between -0.05 and 0.1 respectively. In terms of safety, the percent of trials with DLT rate above $\theta + 0.1$ is below 10% in both scenarios with average number of DLTs of 34% and 27%. We also present results regarding percent correct recommendation. These results are shown in Figures 8 and 9 in Appendix and overall we observe that in the two proposed scenarios, the percent of correct recommentation is between 70% and 100% in the scenario where the MTD passes through the dose combination 15/75 mg/m^2, and between 50% and 100% in the scenario where the MTD is above the dose combination 15/75 mg/m^2. Note that these results depend on a design parameter p that takes the values 0.05 and 0.1 and that states how strict we are when considering a correct recommendation, being $p = 0.1$ less strict than $p = 0.05$. The true parameter values as well as the safety results are shown in Table 1.

In the second stage, 30 additional patients are enrolled to identify the dose combinations along the MTD curve from the first stage, that are likely to have high posterior median TTP. The TTP of the standard care of treatment, which is necessary to perform the hypothesis testing procedure, is chosen to be 4 months since this is the radiographic median TTP in a placebo arm in a previous phase III trial. We present simulations based on 4 scenarios supporting the alternative hypothesis and 4 scenarios favoring the null hypothesis. For each scenario favoring the alternative hypothesis, effect sizes of 1.5 and 2 months and accrual rates of 1 and 2 patients per month will be used. In order to correctly assess the operating characteristics of the design, the 4 scenarios will have the same TTP of the standard of care treatment, the same effect size and the same accrual rate. This way, the only difference between scenarios will be the shape of the TTP median curve, allowing to see the behavior of the design when the optimal dose level is located at different dose levels.

The simulations were carried out using the model and prior distributions presented in sections 3.1 and 3.2.1 respectively, with $n_1 = 10$, $n_2 = 5$, $\delta_u = 0.8$ and $\delta_u = 0.9$.

In Figures 3 and 4 we present the 4 simulated scenarios favoring the alternative hypothesis with effect sizes of 1.5 and 2 months respectively, and an accrual rate of 1 patient per month. In Figure 5 we present the same 4 simulated scenarios favoring the null hypothesis with an accrual rate of 1 patient per month. Results for the same 4 simulated scenarios with an accrual rate of 2 patients per month can be found in Figures 10, 11 and 12 in the Appendix. In these figures, we present the true median TTP curve, the null median TTP, the average recommended dose, the estimated median TTP curve and the posterior probability that the median TTP at a dose level x is greater than the null median TTP as measurements of efficiency. Overall, we observe that the design captures the shape of the median TTP curve. However, the estimated median TTP curve is not a very informative measurement of efficiency since we are using adaptive randomization and hence much more patients are allocated in certain dose levels. The efficiency measurement we believe is more interesting is the posterior probability that the median TTP at a dose level x is greater than the TTP of the standard of care treatment. In all figures, we observe that the recommended optimal dose is very close to the true optimal dose regardless of the shape of the TTP median, the effect size or the accrual rate.

In Table 2, we present the Bayesian power, the probability of the type-I error as well as the probability of type-I + type-II errors for different effect sizes and different accrual rates. With an accrual rate of 1 patient per month, the probability of type-I error remains between 0.104 and 0.227 when $\delta_u = 0.8$ and between 0.235 and 0.308 when $\delta_u = 0.9$. However, with an accrual rate of 2 patients per month, the probability of type-I error is much smaller overall and it remains between 0.035 and 0.107 when $\delta_u = 0.8$ and between 0.008 and 0.048 when $\delta_u = 0.9$.

In terms of power, with effect size of 1.5 months and an accrual rate of 1 patient per month, we observe that the power remains between 0.706 and 0.924 when $\delta_u = 0.8$ and between 0.52 and 0.844 when $\delta_u = 0.9$. If the effect size increases up to 2 months and maintaining the same accrual rate, the power remains between 0.931 and 0.972 when $\delta_u = 0.8$ and between 0.846 and 0.932 when $\delta_u = 0.9$. In contrast, if we fix the accrual rate to 2 patients per month, we observe that overall the power decreases considerately. With an effect size of 1.5 months, the power remains between 0.522 and 0.824 when $\delta_u = 0.8$ and between 0.338 and 0.674 when $\delta_u = 0.9$. If the effect size increases up to 2 months and maintaining the same accrual rate, the power remains between 0.766 and 0.92 when $\delta_u = 0.8$ and between 0.615 and 0.829 when $\delta_u = 0.9$.

Because it is difficult to find the right balance between power and type-I error, and since it is not unusual to find probabilities of type-I error between 0.15 and 0.2 in phase II trials of these characteristics where we try a large set of doses with a small sample size, we evaluate the sum of the probabilities of type-I error and type-II error. In general, a design where the sum of these two probabilities is above 0.3 is not advisable. In our proposal, with effect size of 1.5 months and an accrual rate of 1 patient per month, the sum of the probabilities of type-I error and type-II error remains between 0.235 and 0.416 when $\delta_u = 0.8$ and between 0.264 and 0.523 when $\delta_u = 0.9$. If the effect size increases up to 2 months and maintaining the same accrual rate, the sum of the probabilities of type-I error and type-II error remains between 0.15 and 0.256 when $\delta_u = 0.8$ and between 0.123 and 0.198 when $\delta_u = 0.9$. If we fix the accrual rate to 2 patients per month, with an effect size of 1.5 months, the sum of the probabilities of type-I error and type-II error remains between 0.283 and 0.513 when $\delta_u = 0.8$ and between 0.374 and 0.67 when $\delta_u = 0.9$. If the effect size increases up to 2 months and maintaining the same accrual rate, the the sum of the probabilities of

Table 2: Bayesian power, type I error probability and type-I + type-II error probability in four scenarios with effect sizes of 1.5 and 2 months, and accrual rates of 1 and 2 patients per month.

		Power (effect size of 1.5 months)		Power (effect size of 2 months)		Probability of type-I error		Probability of type-I + type-II errors (effect size of 1.5 months)		Probability of type-I + type-II errors (effect size of 2 months)	
		δ_u		δ_u		δ_u		δ_u		δ_u	
Scenario	Accrual rate	0.8	0.9	0.8	0.9	0.8	0.9	0.8	0.9	0.8	0.9
1	1	0.924	0.844	0.971	0.927	0.227	0.121	0.303	0.277	0.256	0.194
2		0.706	0.520	0.972	0.920	0.122	0.043	0.416	0.523	0.150	0.123
3		0.904	0.808	0.973	0.932	0.139	0.072	0.235	0.264	0.166	0.140
4		0.796	0.646	0.931	0.846	0.104	0.044	0.308	0.398	0.173	0.198
1	2	0.824	0.674	0.920	0.829	0.107	0.048	0.283	0.374	0.187	0.219
2		0.522	0.338	0.865	0.755	0.035	0.008	0.513	0.670	0.170	0.253
3		0.759	0.598	0.896	0.790	0.068	0.024	0.309	0.426	0.172	0.234
4		0.623	0.445	0.766	0.615	0.053	0.018	0.430	0.573	0.287	0.403

type-I error and type-II error remains between 0.17 and 0.287 when $\delta_u = 0.8$ and between 0.219 and 0.403 when $\delta_u = 0.9$.

If we focus one the sum of the probabilities of type-I error and type-II error, we observe that with an effect size of 1.5 months we observe a lot of values above our 0.3 threshold regardless of the accrual rate, which is normal since the original design's primary endpoint was not the TTP median and it is not sufficiently powered for this effect size. In contrast, with an effect size of of 2 months, we observe that if $\delta_u = 0.8$, all the values are below our 0.3 threshold regardless of the accrual rate and if $\delta_u = 0.9$, only one scenario with an accrual rate of 2 patients per month has a value above the threshold.

In Table 3, we present the probability of early stopping and average sample size at the moment of stopping in scenarios favoring the null hypothesis. Overall, we observe that an accrual rate of 2 patients per month produces a slight increase in the probability of early stopping and a decrease between 1 and 2 patients in the average sample size at the moment of stopping with respect to using an accrual rate of 1 patient per month.

Even though it is not listed in the operating characteristics, in Figure 13 in Appendix we show the dose allocation distribution in the four scenarios with the different effect sizes and accrual rates. In scenario 1, we correctly allocate more than 71% of the patients in doses that are above the TTP of the standard of care treatment. In scenarios 2, 3 and 4 we correctly allocate more than 65%, 77% and 60% of the patients respectively in dose above the TTP of the standard of care treatment. Note that from these distributions we excluded the first n_1 doses which are automatically allocated in doses equally spaced along the MTD.

We also implement scenarios 1 and 2 when the TTP of the standard of care treatment is higher than 4 months. More precisely we tuned scenarios 1 and 2 to have effect sizes of 2 months but the TTP of the standard of care treatment is now 8 months. We used accrual rates of 2 and 3 patients per month and we observed values of power, type-I error and sum of type-I and type-II errors that are consistent with the values presented in Table 2. These results are showed in Table 4 in Appendix.

Hence, we conclude that our design has overall good operating characteristics with accrual rates that are considered realistic in practice.

5. Conclusions

In this paper we propose a Bayesian adaptive two-stage design for cancer phase I/II trials using drug combinations with continuous dose levels and TTP endpoint. We are motivated by a phase I trial published by [8], that combines cisplatin and cabazitaxel in patients with prostate cancer with visceral metastasis, where the MTD was established at 15/75 mg/m^2. In the first part of this motivating trial, doses were

Table 3: Probability of early stopping under the null hypothesis in four scenarios with accrual rates of 1 and 2 patients per month.

		Probability of early stopping			Average sample size		
		δ_0			δ_0		
Scenario	Accrual rate	0.10	0.15	0.20	0.10	0.15	0.20
1	1	0.164	0.266	0.355	19.94	18.14	16.92
2		0.121	0.252	0.390	17.44	15.41	13.76
3		0.234	0.358	0.506	20.17	17.67	16.13
4		0.264	0.417	0.554	19.79	17.75	16.17
1	2	0.307	0.457	0.576	18.55	16.95	15.99
2		0.270	0.452	0.611	15.64	13.63	12.20
3		0.410	0.611	0.740	19.02	16.94	15.01
4		0.421	0.590	0.731	18.41	16.44	14.67

escalated according to a standard "3+3" design and no DLTs were observed at dose combination which was found to be the MTD. During the second part of the trial, 15 additional patients were treated at the MTD and 2 DLTs were observed, and in total 18 patients were treated at the MTD. However, considering these results, there may be other active dose combinations that are tolerable and active in prostate cancer with visceral metastasis.

In the first stage of the design a logistic model is used to model the probability of DLT. The dose escalation algorithm proceeds by using EWOC as described in [16]. At the end of this stage, an estimate of the MTD curve is obtained. In the second stage we model the median TTP along the MTD curve using a Weibull model and incorporating a cubic spline through the scale parameter of the model. In this stage of the design a hypothesis test is performed where the null hypothesis states that the median TTP corresponding to all dose levels is below or equal to a TTP of the standard of care treatment. On the other hand, the alternative hypothesis states that the median TTP corresponding to some dose levels is above the TTP of the standard of care treatment. The dose escalation in the second stage proceeds by first allocating n_1 patients in dose levels equally spaced along the MTD curve. Subsequent patients are allocated in cohorts of n_2 patients in doses with higher posterior probability of having a median TTP greater than the TTP of the standard of care treatment using adaptive randomization.

Regarding the first stage, we studied the operating characteristics in 2 scenarios. In the one scenario the true MTD curve passes through the dose combination 15/75 mg/m^2, whereas in the other scenario the true MTD curve is significantly above the dose combination 15/75 mg/m^2. We found that this stage of the trial is safe and has good operating characteristics in terms of pointwise bias and percent selection. Note that the operating characteristics of this stage were evaluated using informative prior distributions as commented in section 4.

With respect to the second stage, we studied the operating characteristics of the design in 4 scenarios in which the null median TTP is the same and so is the effect size and accrual rate and hence the only difference between them is the median TTP curve shape that places the dose level with the highest TTP at a different location in each scenario. These 4 scenarios were implemented with effect sizes of 1.5 and 2 months, and accrual rates of 1 patient and 2 patients per month, which are considered reasonable in practice. In general, we observed good operating characteristics in terms of optimal dose recommendation and sum of the probabilities of type-I and type-II errors as main measurements of efficiency. Scenarios 1 and 2 were also implemented when the TTP of the standard of care treatment is higher than 4 months. More precisely we tuned scenarios 1 and 2 to have effect sizes of 2 months but the TTP of the standard of care treatment is now 8 months. We used accrual rates of 2 and 3 patients per month and we observed values of power, type-I error and sum of type-I and type-II errors that are consistent with the values presented in Table 2.

Note that the operating characteristics in the second stage were evaluated under vague prior distributions of the model parameters and no efficacy profiles of single agent trials we used *a priori*.

Acknowledgments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 633567 (J.J.), the National Institute of Health Grant Number 1R01CA188480-01A1 (M.T.), the National Center for Research Resources, Grant UL1RR033176, and is now at the National Center for Advancing Translational Sciences, Grant UL1TR000124 (M.T.), and 2 P01 CA098912 (M.T.).

Figure 1: True and estimated MTD curve for scenario 1 in the first stage of the design.

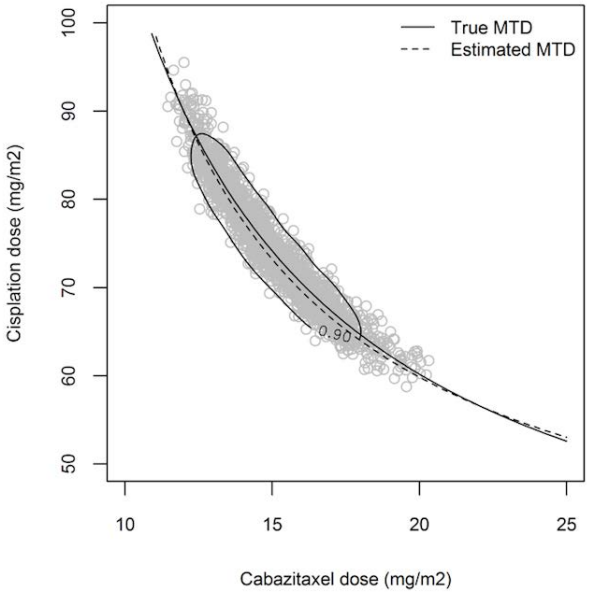

Figure 2: True and estimated MTD curve for scenario 2 in the first stage of the design.

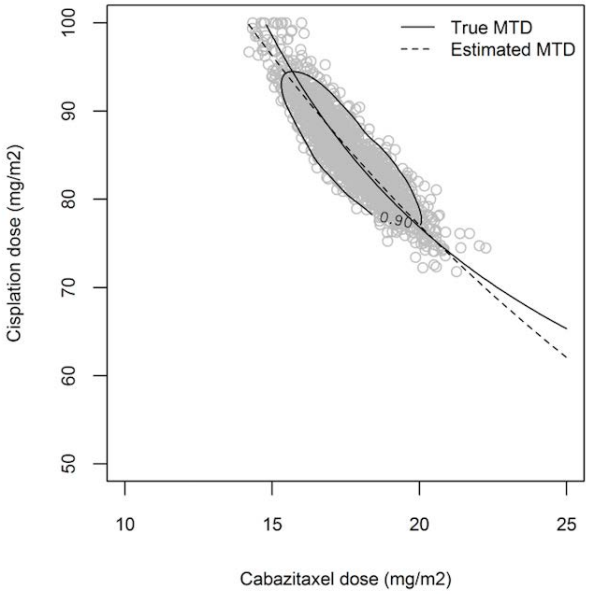

Figure 3: True and estimated TTP medians under four scenarios favoring the alternative hypothesis with effect size of 1.5 months and accrual rate of 1 patient per month.

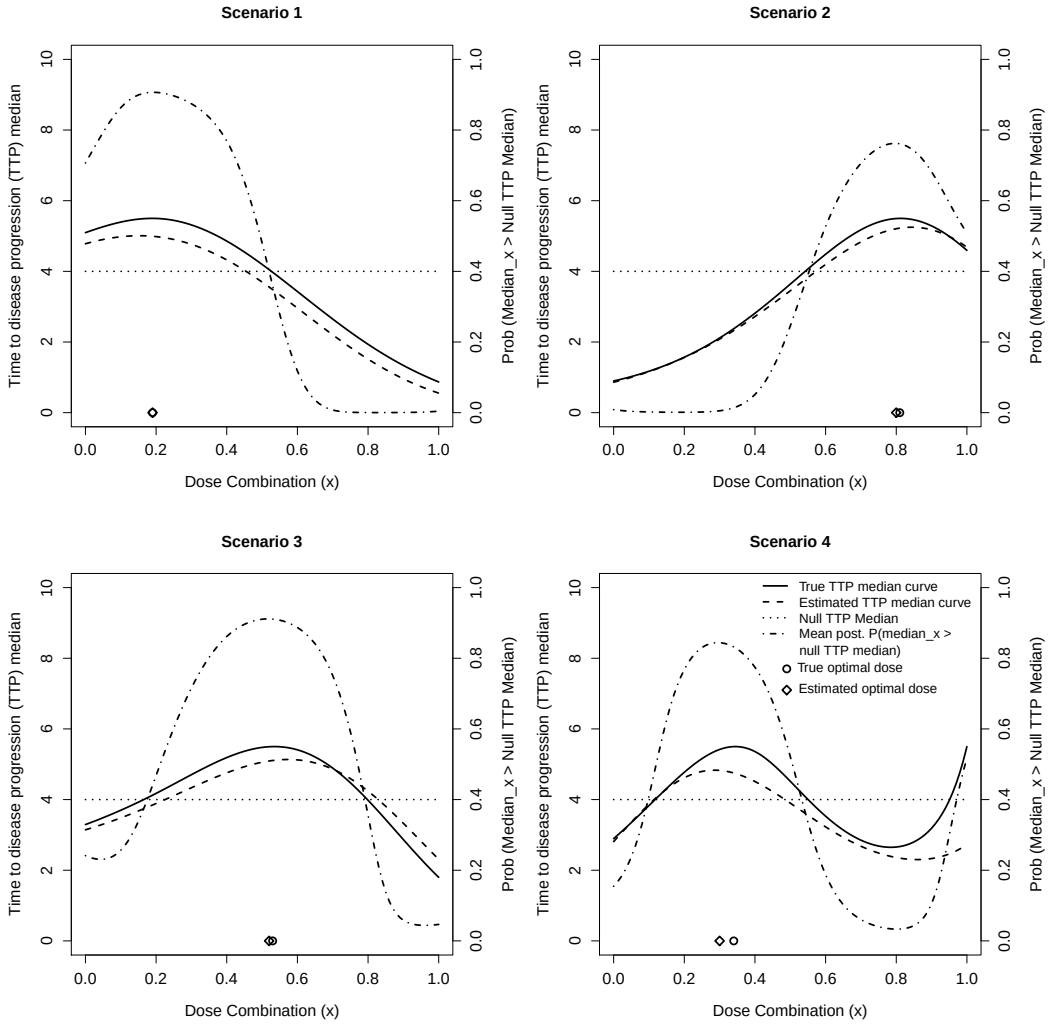

Figure 4: True and estimated TTP medians under four scenarios favoring the alternative hypothesis with effect size of 2 months and accrual rate of 1 patient per month.

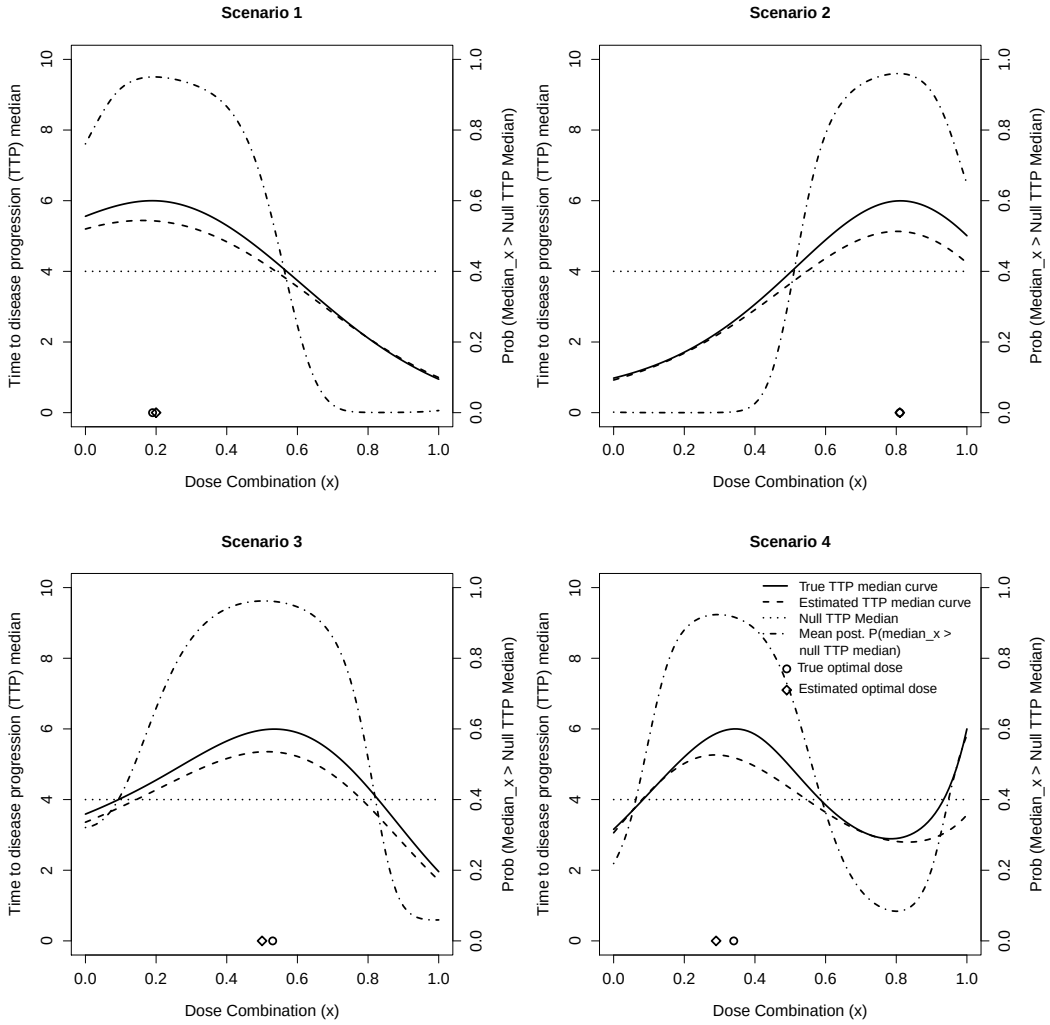

Figure 5: True and estimated TTP medians under four scenarios favoring the null hypothesis with an accrual rate of 1 patient per month.

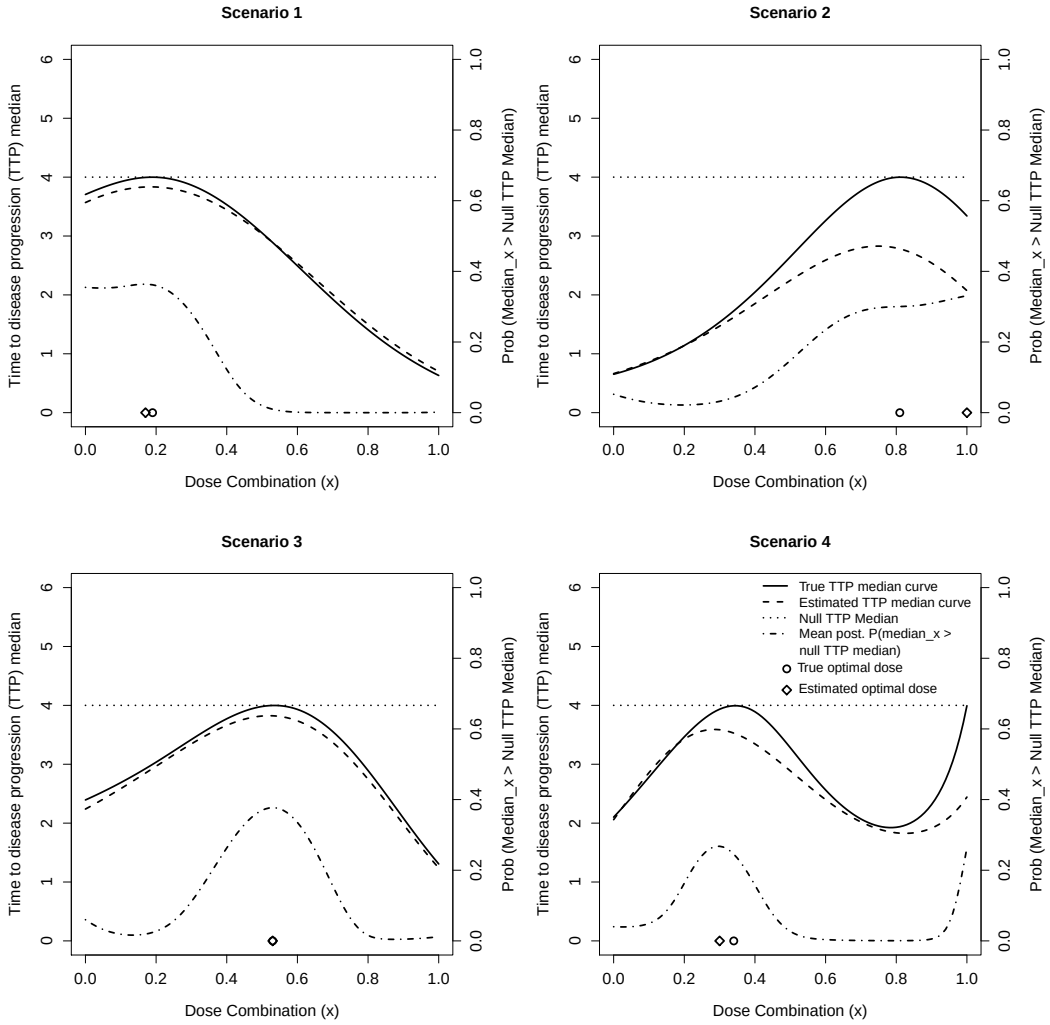

References

[1] James Babb, André Rogatko, and Shelemyahu Zacks. Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statistics in medicine*, 17(10):1103–1120, 1998.

[2] Thomas M Braun. The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Controlled clinical trials*, 23(3):240–256, 2002.

[3] Zhengjia Chen, Ying Yuan, Zheng Li, Michael Kutner, Taofeek Owonikoko, Walter J Curran, Fadlo Khuri, and Jeanne Kowalski. Dose escalation with over-dose and under-dose controls in phase I/II clinical trials. *Contemporary clinical trials*, 43:133–141, 2015.

[4] Zhengjia Chen, Yichuan Zhao, Ye Cui, and Jeanne Kowalski. Methodology and application of adaptive and sequential approaches in contemporary clinical trials. *Journal of Probability and Statistics*, 2012, 2012.

[5] Xuelin Huang, Swati Biswas, Yasuhiro Oki, Jean-Pierre Issa, and Donald A Berry. A parallel phase I/II clinical trial design for combination therapies. *Biometrics*, 63(2):429–436, 2007.

[6] Anastasia Ivanova. A new dose-finding design for bivariate outcomes. *Biometrics*, 59(4):1001–1007, 2003.

[7] Christophe Le Tourneau, J Jack Lee, and Lillian L Siu. Dose escalation methods in phase I cancer clinical trials. *JNCI: Journal of the National Cancer Institute*, 101(10):708–720, 2009.

[8] A Craig Lockhart, Shankar Sundaram, John Sarantopoulos, Monica M Mita, Andrea Wang-Gillam, Jennifer L Moseley, Stephanie L Barber, Alex R Lane, Claudine Wack, Laurent Kassalow, et al. Phase I dose-escalation study of cabazitaxel administered in combination with cisplatin in patients with advanced solid tumors. *Investigational new drugs*, 32(6):1236–1245, 2014.

[9] Paul A Murtaugh and Lloyd D Fisher. Bivariate binary models of efficacy and toxicity in dose-ranging trials. *Communications in Statistics-Theory and Methods*, 19(6):2003–2020, 1990.

[10] André Rogatko, Pulak Ghosh, Brani Vidakovic, and Mourad Tighiouart. Patient-specific dose adjustment in the cancer clinical trial setting. *Pharmaceutical Medicine*, 22(6):345–350, 2008.

[11] Hiroyuki Sato, Akihiro Hirakawa, and Chikuma Hamada. An adaptive dose-finding method using a change-point model for molecularly targeted agents in phase I trials. *Statistics in medicine*, 35(23):4093–4109, 2016.

[12] Fumiya Shimamura, Chikuma Hamada, Shigeyuki Matsui, and Akihiro Hirakawa. Two-stage approach based on zone and dose findings for two-agent combination phase I/II trials. *Journal of biopharmaceutical statistics*, pages 1–13, 2018.

[13] Peter F Thall and John D Cook. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*, 60(3):684–693, 2004.

[14] Peter F Thall and Kathy E Russell. A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics*, pages 251–264, 1998.

[15] Mourad Tighiouart. Two-stage design for phase I/II cancer clinical trials using continuous dose combinations of cytotoxic agents. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 2018.

[16] Mourad Tighiouart, Quanlin Li, and André Rogatko. A bayesian adaptive design for estimating the maximum tolerated dose curve using drug combinations in cancer phase I clinical trials. *Statistics in medicine*, 36(2):280–290, 2017.

[17] Mourad Tighiouart and Andre Rogatko. Number of patients per cohort and sample size considerations using dose escalation with overdose control. *Journal of Probability and Statistics*, 2012, 2012.

[18] Mourad Tighiouart, André Rogatko, and James S Babb. Flexible bayesian methods for cancer phase I clinical trials. dose escalation with overdose control. *Statistics in medicine*, 24(14):2183–2196, 2005.

[19] Mourad Tighiouart, André Rogatko, et al. Dose finding with escalation with overdose control (EWOC) in cancer clinical trials. *Statistical Science*, 25(2):217–226, 2010.

[20] Guosheng Yin, Yisheng Li, and Yuan Ji. Bayesian dose-finding in phase I/II clinical trials using toxicity and efficacy odds ratios. *Biometrics*, 62(3):777–787, 2006.

[21] Ying Yuan and Guosheng Yin. Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 58(5):719–736, 2009.

Appendix

In this section we display Figures that contain information regarding operating characteristics of the design and support the conclusions obtained along the manuscript.

Figure 6: Pointwise average bias for scenario 1 in the first stage of the design.

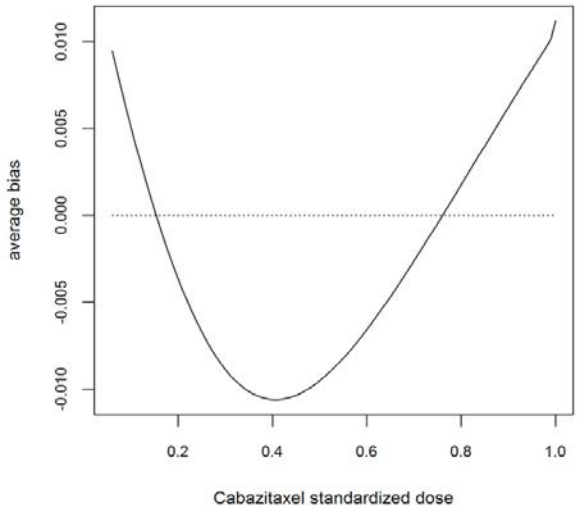

Figure 7: Pointwise average bias for scenario 2 in the first stage of the design.

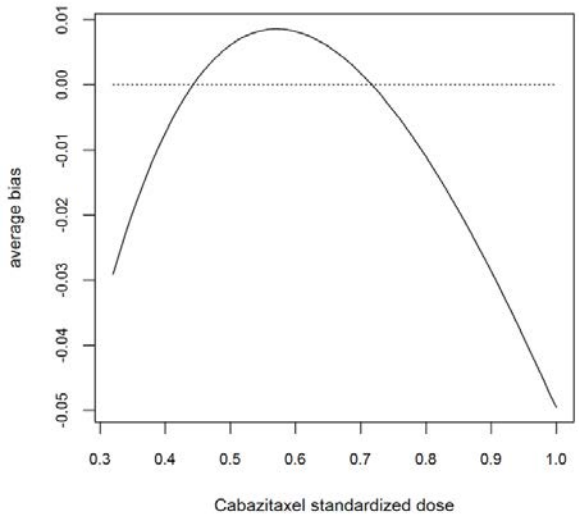

Figure 8: Percent of correct recommendation for scenario 1 in the first stage of the design.

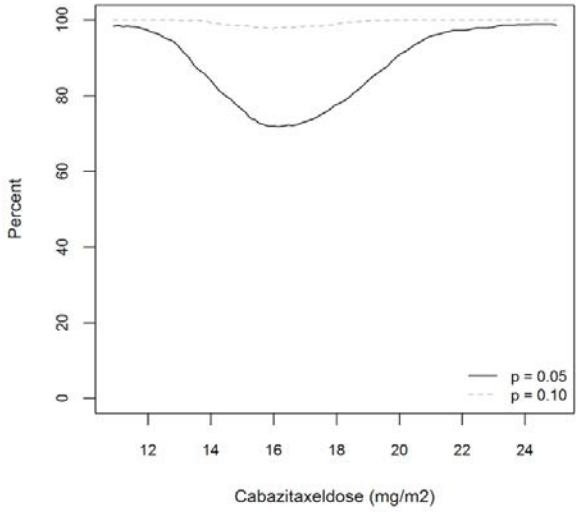

Figure 9: Percent of correct recommendation for scenario 2 in the first stage of the design.

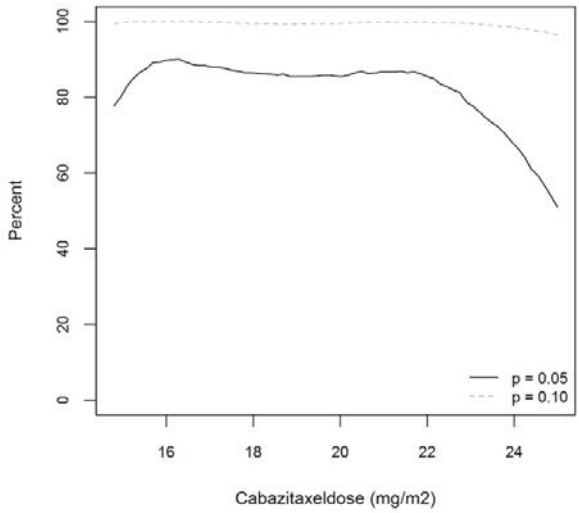

Figure 10: True and estimated TTP medians under four scenarios favoring the alternative hypothesis with effect size of 1.5 months and accrual rate of 2 patients per month.

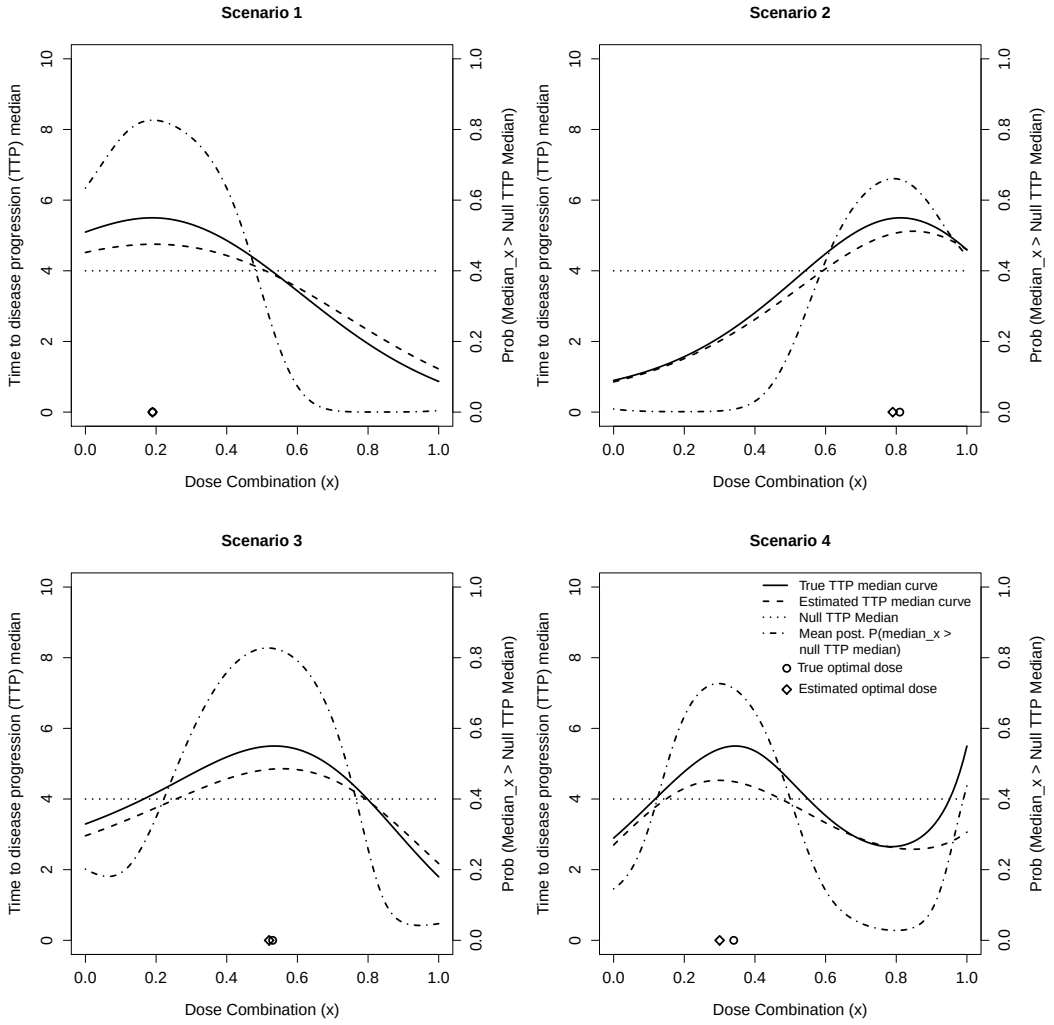

Figure 11: True and estimated TTP medians under four scenarios favoring the alternative hypothesis with effect size of 2 months and accrual rate of 2 patients per month.

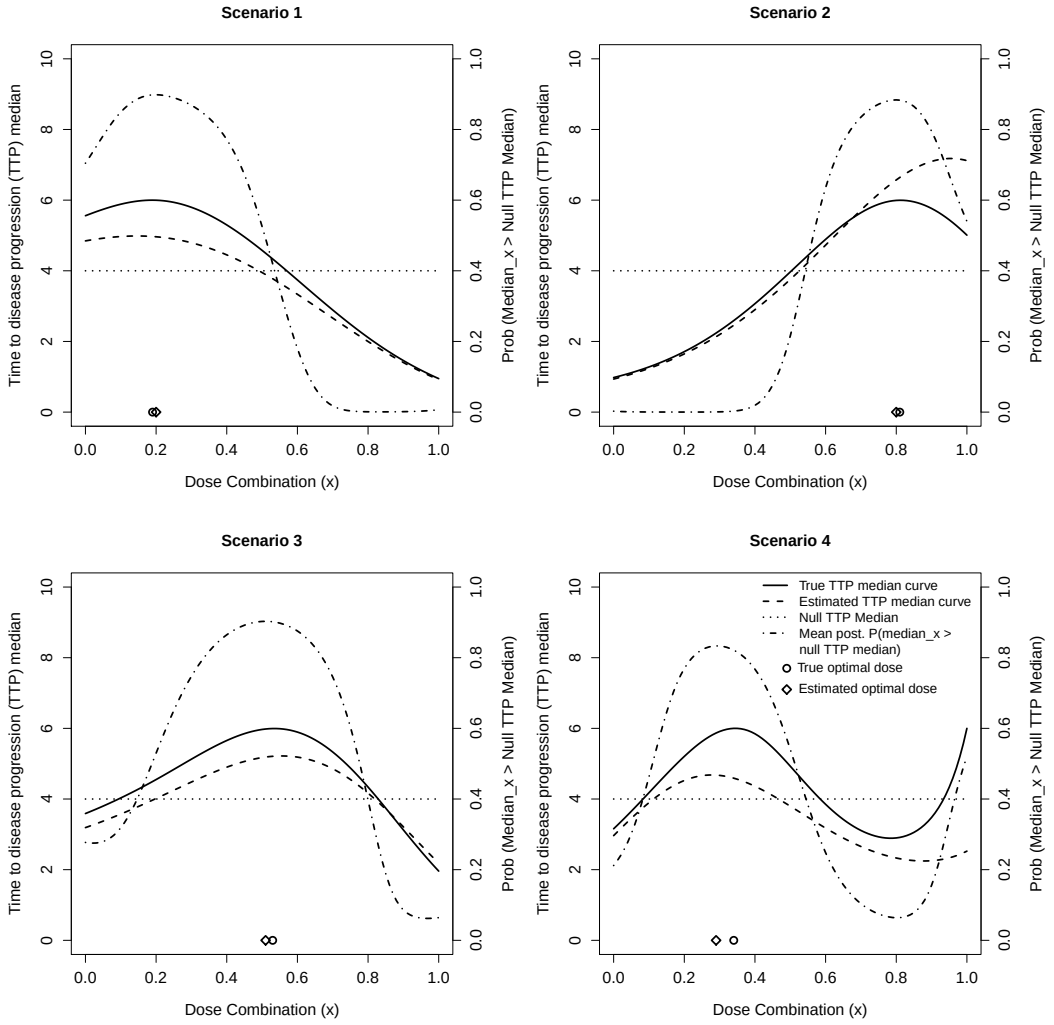

Figure 12: True and estimated TTP medians under four scenarios favoring the null hypothesis with an accrual rate of 2 patients per month.

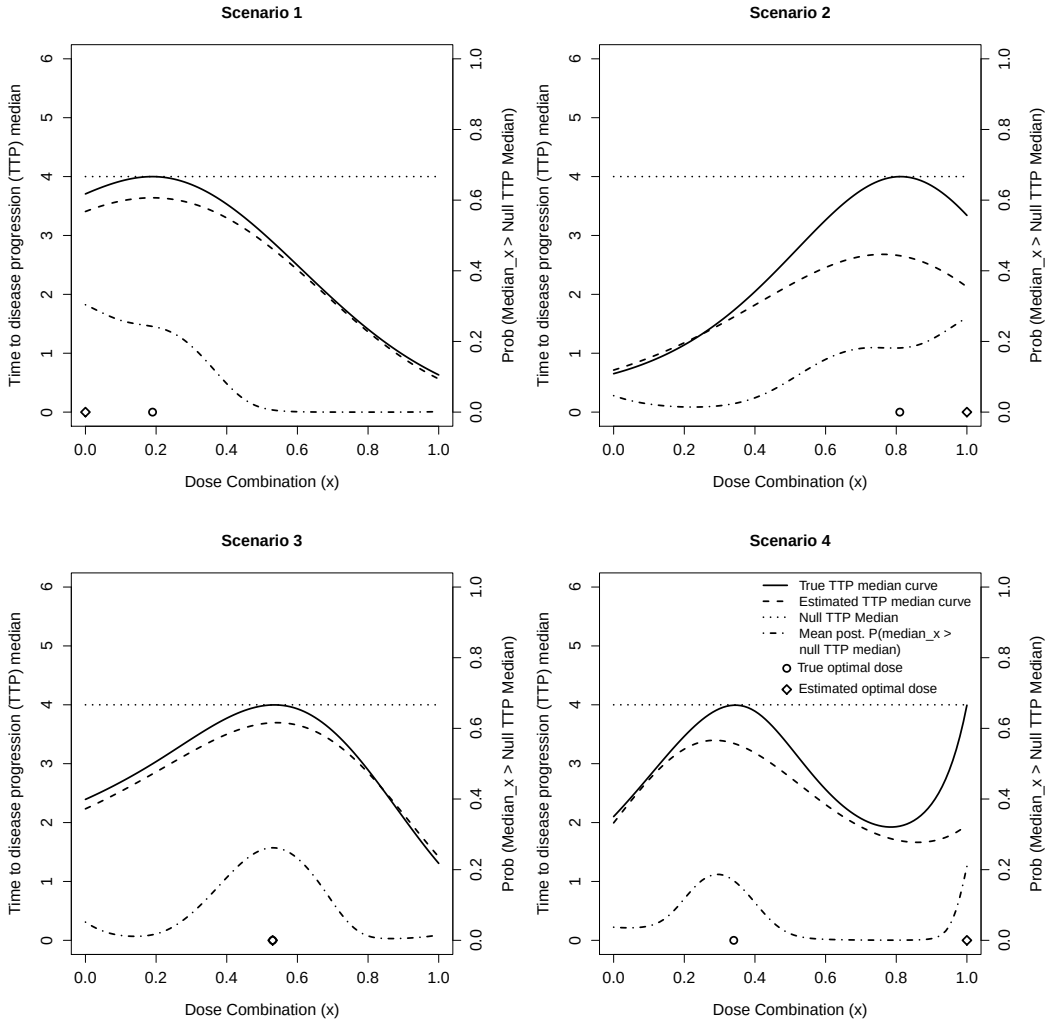

Figure 13: Dose allocation density across scenarios with effect sizes (ES) of 1.5 and 2 months and accrual rates (AR) of 1 and 2 patients per month.

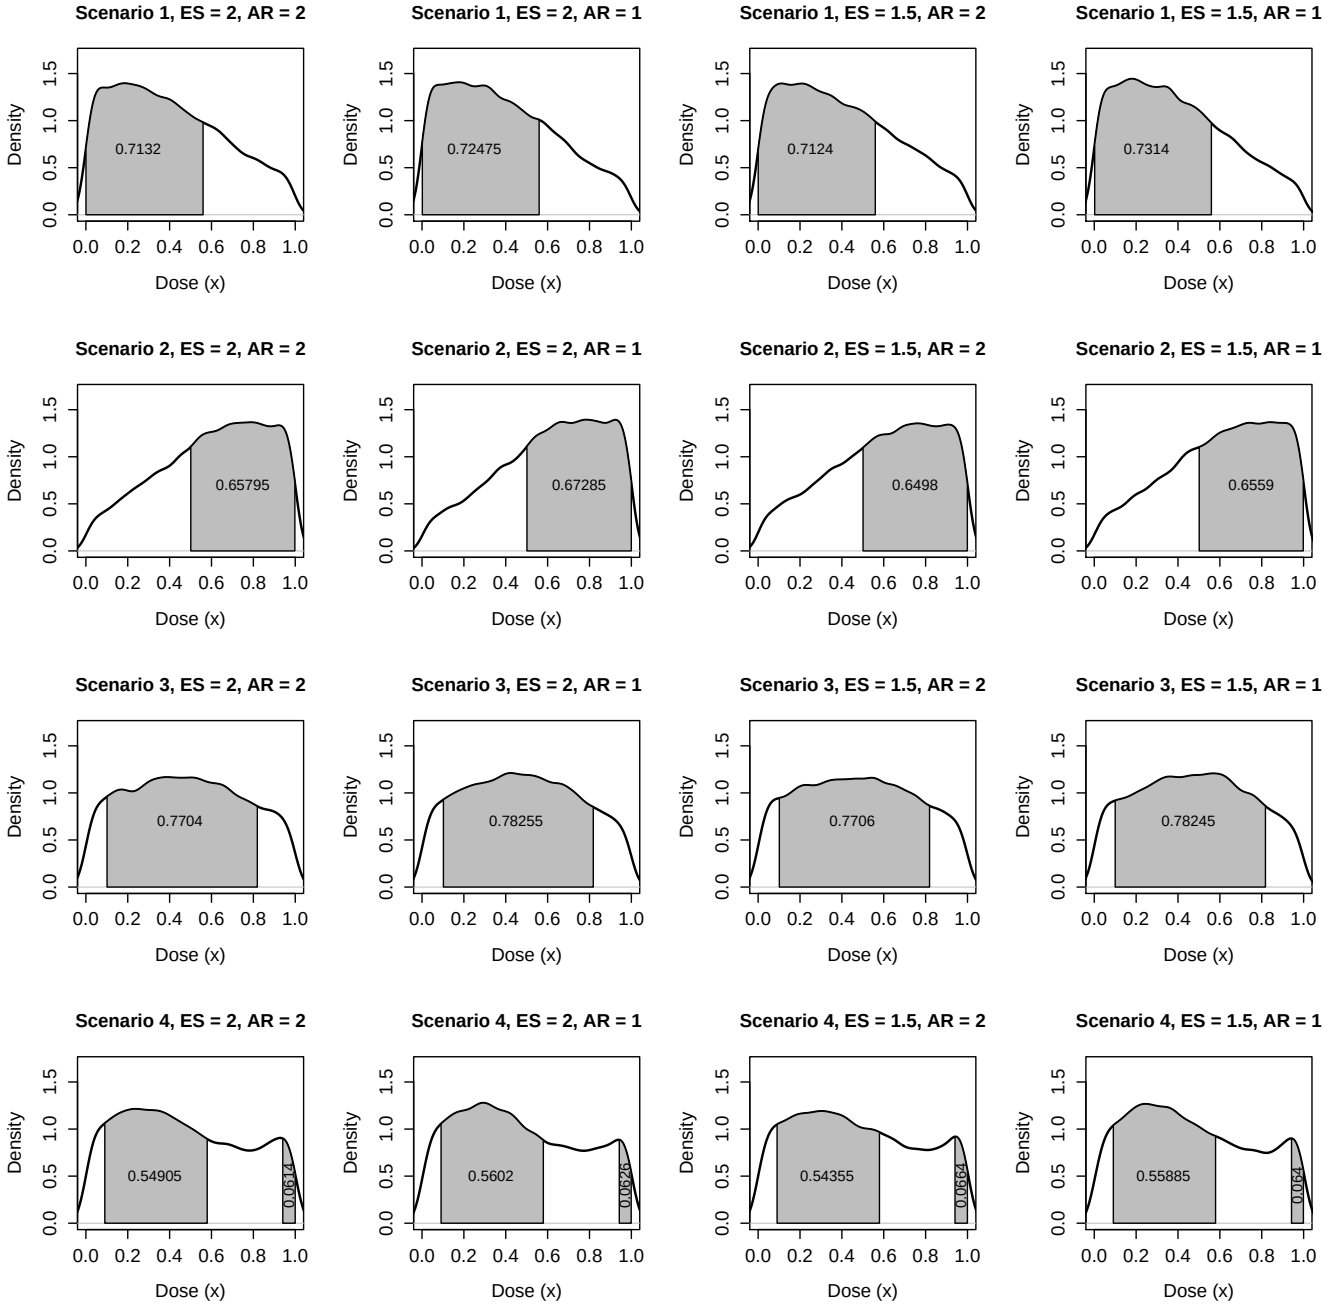

Table 4: Bayesian power, type I error probability and type-I + type-II error probability in two scenarios with effect size of 2 months, and accrual rates of 2 and 3 patients per month.

		Power (effect size of 2 months)		Probability of type-I error		Probability of type-I + type-II errors (effect size of 2 months)	
		δ_u		δ_u		δ_u	
Scenario	Accrual rate	0.8	0.9	0.8	0.9	0.8	0.9
1	2	0.836	0.677	0.162	0.080	0.326	0.403
2		0.742	0.577	0.121	0.052	0.379	0.475
1	3	0.824	0.639	0.167	0.081	0.343	0.442
2		0.747	0.572	0.109	0.046	0.362	0.474