POLITECNICO DI TORINO
Repository ISTITUZIONALE

Combining cytotoxic agents with continuous dose levels in seamless phase I-II clinical trials

Original
Combining cytotoxic agents with continuous dose levels in seamless phase I-II clinical trials / Jiménez, J.L., Tighiouart, M.. - In: JOURNAL OF THE ROYAL STATISTICAL SOCIETY SERIES C-APPLIED STATISTICS. - ISSN 0035-9254. - 71:5(2022), pp. 1996-2013. [10.1111/rssc.12598]

Availability:
This version is available at: 11583/2973291 since: 2022-11-22T15:28:53Z

Publisher:
Wiley

Published
DOI:10.1111/rssc.12598

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright
Wiley postprint/Author's Accepted Manuscript

This is the peer reviewed version of the above quoted article, which has been published in final form at http://dx.doi.org/10.1111/rssc.12598.This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions.

(Article begins on next page)

07 November 2026

arXiv:2109.14231v3 [stat.ME] 26 Oct 2022

COMBINING CYTOTOXIC AGENTS WITH CONTINUOUS DOSE LEVELS IN SEAMLESS PHASE I-II CLINICAL TRIALS

A PREPRINT

José L. Jiménez
Novartis Pharma A.G.
Basel, Switzerland
jose_luis.jimenez@novartis.com

Mourad Tighiouart
Cedars-Sinai Medical Center
Los Angeles, USA
Mourad.Tighiouart@cshs.org

October 27, 2022

ABSTRACT

Phase I-II cancer clinical trial designs are intended to accelerate drug development. In cases where efficacy cannot be ascertained in a short period of time, it is common to divide the study in two stages: i) a first stage in which dose is escalated based only on toxicity data and we look for the maximum tolerated dose (MTD) set and ii) a second stage in which we search for the most efficacious dose within the MTD set. Current available approaches in the area of continuous dose levels involve fixing the MTD after stage I and discarding all collected stage I efficacy data. However, this methodology is clearly inefficient when there is a unique patient population present across stages. In this article, we propose a two-stage design for the combination of two cytotoxic agents assuming a single patient population across the entire study. In stage I, conditional escalation with overdose control (EWOC) is used to allocate successive cohorts of patients. In stage II, we employ an adaptive randomization approach to allocate patients to drug combinations along the estimated MTD curve, which is constantly updated. The proposed methodology is assessed with extensive simulations in the context of a real case study.

1 Introduction

Integrated phase I-II clinical trial designs are efficient approaches to accelerate drug development. This seamless process uses a single protocol to investigate two aspects of the development of a compound, namely safety and efficacy.

With the use of chemotherapy compounds, a dose limiting toxicity (DLT) is generally ascertained after one cycle of therapy for the purpose of estimating the maximum tolerated dose (MTD). If the efficacy outcome can be ascertained in a short period of time (e.g., after one or two cycles of therapy), it is common to employ one-stage designs where the joint probability of toxicity and efficacy is sequentially updated after each cohort of patients. Extensive methodology to carry out such designs can be found in [24, 3, 8, 23, 5, 20, 33] for single agent trials and in [7, 32] for drug combinations. In contrast, when efficacy cannot be ascertained in a short period of time, we may employ two-stage designs where a set of maximum tolerated dose combinations are first selected, and then tested for efficacy in a subsequent second stage (see e.g., [19, 13, 6, 22, 34, 35, 25, 10]).

This work is motivated by the cisplatin-cabazitaxel clinical trial design proposed by [25]. In that phase I-II study design, stage I was inspired by a previous phase I trial published by [15] in patients with advanced solid tumors, where a single MTD of cisplatin/cabazitaxel 15/75 mg/m^2 was recommended. However, based on their results and other preliminary efficacy data, it was hypothesized by the clinical team that there could be a series of dose combinations which could be tolerable and efficacious in prostate cancer. In stage II of [25], the design enrolls 30 additional patients to identify dose combinations with high probability of efficacy, along the MTD curve estimated from the first stage data.

A limitation of the original cisplatin-cabazitaxel design by [25] is that uncertainty about the estimated MTD curve is not taken into account in stage II. Moreover, efficacy data from stage I were not used to model the dose-efficacy relationship in stage II. The rational for this disconnection was that stage I and stage II populations were not exactly the same, and

therefore it was hypothesised that the dose-efficacy profiles could also be different. Recently, [12] proposed a method to robustly combine efficacy data across stages in this particular setting. In practice however, it is not uncommon to design phase I-II clinical trials with the same patient population across stages. In such a situation, it would be more sensible to use all the toxicity and efficacy data collected from stage I as well as update the MTD curve during stage II. To the best of our knowledge, such a design does not exist in the phase I-II clinical trial designs literature with combination of cytotoxic agents and continuous dose levels.

In this article, we extend the work of [25] by allowing the use of all the safety and efficacy data collected throughout the entire trial, which implies that the MTD curve is sequentially updated during stage II. Since we assume that efficacy cannot be ascertained in a short period, the design proceeds in two stages. The design's operating characteristics are evaluated by extensive simulations in the context of the cisplatin-cabazitaxel clinical trial.

The rest of the manuscript is organized as follows. In section 2 we describe the clinical trial that motivates this research. In section 3 we introduce the dose-toxicity and dose-efficacy models as well as the study design. In section 4 we present the operating characteristics under several scenarios for the true dose-toxicity and dose-efficacy relationships. Last, in section 5, we discuss the major contributions of this manuscript together with potential areas where there may be room for improvement.

2 The cisplatin-cabazitaxel trial

The two-stage phase I-II cisplatin-cabazitaxel trial design proposed by [25] investigates the combination of cisplatin and cabazitaxel, given every 3 weeks, in patients with metastatic, castration resistant prostate cancer. A continuous set of doses that range from 10 to 25 mg/m^2 for cabazitaxel and from 50 to 100 mg/m^2 for cisplatin were administered intravenously. In stage I, the study enrolls 30 patients using conditional escalation with overdose control (EWOC) algorithm (see [26]) to estimate the MTD curve. In stage II, the study enrolls another 30 patients from the same population of patients but with visceral metastasis to identify dose combinations with high probability of efficacy along the MTD curve estimated at the end of stage I. These patients are allocated to dose combinations along the MTD curve using a Bayesian adaptive design after modeling the dose-efficacy curve with cubic splines.

The first stage of the cisplatin-cabazitaxel trial was motivated by the phase I trial published by [15], where the combinations of cisplatin/cabazitaxel 15/75 mg/m^2, 20/75 mg/m^2 and 25/75 mg/m^2 were explored in patients with advanced solid tumors, using a "3+3" design. The recommended MTD was 15/75 mg/m^2 on the basis of data from 9 evaluable patients in phase I and 15 patients in the expansion cohort, although only 2 patients out of 18 treated at the recommended MTD had DLT. Based on this low toxicity rate and preliminary efficacy data, the investigators hypothesized that a series of tolerable dose combinations which could be efficacious exist in prostate cancer.

The original cisplatin-cabazitaxel trial design was tailored for a situation with different patient populations across stages. As previously mentioned, it is not uncommon to have the same patient population across the entire study. Under this assumption, having two disconnected stages would not be sensible since the MTD could be constantly updated and the efficacy data collected in stage I could be used in stage II. Therefore, we propose an extension of the original cisplatin-cabazitaxel trial design to a situation where we have either the same patient population across the two stages or very similar populations with respect to their susceptibility to toxicities and clinical benefit, the primary response endpoint in the cisplatin-cabazitaxel trial.

3 Method

3.1 Formulation of the problem

Let $\{X_{\min}, Y_{\min}, X_{\max}, Y_{\max}\}$ be the lower and upper bound of the continuous dose levels x, y of the compounds X and Y, respectively. The doses are standardized to fall within the interval [0,1] using the transformations $h_1(x) = (x - X_{\min})/(X_{\max} - X_{\min})$ and $h_2(y) = (y - Y_{\min})/(Y_{\max} - Y_{\min})$, so that $x \in [0, 1]$ and $y \in [0, 1]$.

Let $Z = \{0, 1\}$ be the binary indicator of DLT where $Z = 1$ represents the presence of a DLT after a predefined number of treatment cycles, and $Z = 0$ otherwise. Let $E = \{0, 1\}$ be the indicator of treatment response where $E = 1$ represents a positive response after a predefined number of treatment cycles, and $E = 0$ otherwise. Following Tighiouart (2019) [25], let $\theta_Z = 0.33$ be the target probability of DLT and $\theta_E = 0.15$ be the probability of efficacy of the standard of care treatment. The goal of this phase I-II trial is to identify dose combinations (x, y) with probability of DLT $\pi_Z(x, y) \leq \theta_Z$ and probability of efficacy $\pi_E(x, y) \geq \theta_E$. When employing synergistic cytotoxic agents, it is common to assume that both the dose-toxicity and dose-efficacy relationship are monotonically increasing functions. This implies that the optimal dose will lie in the MTD set, defined as $\mathcal{M} = \{(x, y) : \pi_Z(x, y) = \theta_Z\}$.

3.2 Definition of the marginal dose-toxicity and dose-efficacy models

Following the work of [9, 4, 16], we assume that the binary outcomes of toxicity and efficacy are independent. Let the marginal probability of DLT and marginal probability of efficacy models be defined as

$$\pi_Z(x, y) = P(Z = 1|x, y) = F(\alpha_0 + \alpha_1 x + \alpha_2 y + \alpha_3 xy), \tag{1}$$

and

$$\pi_E(x, y) = P(E = 1|x, y) = F(\beta_0 + \beta_1 x + \beta_2 y + \beta_3 xy + \beta_4 x^2 + \beta_5 y^2), \tag{2}$$

respectively, with $\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3 > 0$, and where $F(.)$ represents the cumulative distribution function of the standard normal distribution. Note that in (1), DLT is assumed to be attributed to either drug X or Y, or both, i.e., we do not take into account toxicity attributions (see [11]) because the number of attributable DLTs is expected to be very low given the cytotoxic nature of cisplatin and cabazitaxel.

In (2), we added quadratic terms $\beta_4 x^2, \beta_5 y^2$ to allow more flexibility in the marginal efficacy model $\pi_E(x, y)$. The rationale for this is that the dose-efficacy profiles along the MTD curve can become quite complex (see Figure 2) when having synergistic compounds (i.e., $\alpha_3 > 0$) even if we only consider linear terms in both marginal models. If we also account for the uncertainty in the MTD curve estimation and the relatively small sample size available, in the development of this research we have found that a marginal efficacy model $\pi_E(x, y)$ with only linear terms was, in certain situations, not able to capture the dose-efficacy profile along the MTD curve. In our simulation scenarios, the optimal dose combination is guaranteed to be on the MTD curve because the underlying true dose-efficacy relationship (i.e., the data-generating model) is a monotonically increasing function. However, in a setting with dose-toxicity profiles without drug-drug interactions (i.e., $\alpha_3 = 0$), this additional flexibility would not be necessary and a marginal efficacy model with only linear terms would be sufficient to capture any dose-efficacy profile along the MTD curve.

Following [26], we reparameterize the marginal probability of DLT defined in equation (1) in terms of parameters that clinicians can easily interpret. Let ρ_{uv} denote the probability of DLT when the levels of agents $X = u$ and $Y = v$, with $u = 0, 1$, and $v = 0, 1$, so that $\alpha_0 = F^{-1}(\rho_{00})$, $\alpha_1 = (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))$, and $\alpha_2 = (F^{-1}(\rho_{01}) - F^{-1}(\rho_{00}))$. It then follows that the MTD set takes the form

$$\mathcal{M} = \left\{ (x, y) : y = \frac{(F^{-1}(\theta) - F^{-1}(\rho_{00})) - (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))x}{(F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})) + \alpha_3 x} \right\}. \tag{3}$$

3.3 Likelihood and posterior distributions

We assume that ρ_{10}, ρ_{01} and α_3 are independent *a priori* with $\rho_{01} \sim$ beta$(1, 1)$, $\rho_{10} \sim$ beta$(1, 1)$, and conditional on (ρ_{01}, ρ_{10}), $\rho_{00}/\min(\rho_{01}, \rho_{10}) \sim$ beta$(1, 1)$. Also, let the interaction parameter $\alpha_3 \sim$ gamma(0.1, 0.1). Let $D_n = \{(x_i, y_i, Z_i, E_i), i = 1, \ldots, n\}$ be the data collected after n patients have been enrolled. The posterior distribution of the dose-toxicity model parameters is

$$\begin{aligned}\pi(\rho_{00}, \rho_{10}, \rho_{01}, \alpha_3|D_n) \propto \prod_{i=1}^{n} G((\rho_{00}, \rho_{10}, \rho_{01}, \alpha_3; x_i, y_i))^{Z_i} \times (1 - G(\rho_{00}, \rho_{10}, \rho_{01}, \alpha_3; x_i, y_i))^{1-Z_i} \\ \times \pi(\rho_{01})\pi(\rho_{10})\pi(\rho_{00}|\rho_{01}, \rho_{10})\pi(\alpha_3),\end{aligned}$$

where

$$\begin{aligned}G(\rho_{00}, \rho_{10}, \rho_{01}, \alpha_3; x_i, y_i) = F(F^{-1}(\rho_{00}) + (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))x_i \\ + (F^{-1}(\rho_{01}) - F^{-1}(\rho_{00}))y_i + \alpha_3 x_i y_i).\end{aligned}$$

In the dose-efficacy model, we do not incorporate any prior knowledge. Thus, we assume that $\beta_0, \beta_4, \beta_5 \sim$ N$(0, 10^2)$, $\beta_1, \beta_2, \beta_3 \sim$ Gamma$(0.1, 0.1)$ and that $\beta_j, j = 0, \ldots, 5$ are independent *a priori*. The posterior distribution of the dose-efficacy model parameters is

$$\begin{aligned}\pi(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5|D_n) \propto \prod_{i=1}^{n} \pi_E(x_i, y_i)^{E_i} \times (1 - \pi_E(x_i, y_i))^{1-E_i} \\ \times \pi(\beta_0)\pi(\beta_1)\pi(\beta_2)\pi(\beta_3)\pi(\beta_4)\pi(\beta_5)\end{aligned}$$

All Bayesian computations are done using `R` (see [18]) and `JAGS` (see [17]).

3.4 Study design

Stage I will enroll a total of $N_1 = C_1 \times m_1$ patients, where C_1 represents the total number of cohorts in phase I with equal number of patients m_1. Stage II will enroll a total of $N_2 = C_2 \times m_2$ patients. Let $N = N_1 + N_2$ be the total number of patients that the entire study will enroll and $\widehat{\mathcal{M}}_N$ be the estimated MTD set based on data from all N patients. At the end of stage II, we test the following null and alternative hypotheses

$$\begin{aligned} &H_0 : \pi_E(x,y) \leq \theta_E \text{ for all } (x,y) \in \widehat{\mathcal{M}}_N \quad \text{vs.} \\ &H_1 : \pi_E(x,y) > \theta_E \text{ for some } (x,y) \in \widehat{\mathcal{M}}_N, \end{aligned} \tag{4}$$

and we reject the null hypothesis if

$$\max_{(x,y)\in\widehat{\mathcal{M}}_N} P(\pi_E(x,y) > \theta_E | D_N) > \delta_u, \tag{5}$$

where δ_u is a pre-specified design parameter. Moreover, the dose combination

$$(x,y)_{opt} = \underset{(x,y)\in\widehat{\mathcal{M}}_N}{\text{argmax}} \; P(\pi_E(x,y) > \theta_E | D_N), \tag{6}$$

is recommended as the optimal dose combination and is selected for further phase II or III studies.

Stage I is based on the escalation with overdose control (EWOC) principle (see [1, 29, 30, 26, 28, 21]) where the posterior probability of overdosing the next cohort of patients is bounded by a feasibility bound α. In a cohort with two patients, the first one would receive a new dose of compound X given that the dose y of compound Y that was previously assigned. The other patient would receive a new dose of compound Y given that dose x of compound X was previously assigned. These steps are described in Stage I of Algorithm 1. Using EWOC, these new doses are at the α-th percentile of the conditional posterior distribution of the maximum tolerated dose combinations. The feasibility bound α increases from 0.25 up to 0.5 in increments of 0.05 (see [31]).Accrual continues until the maximum sample size in stage I is reached or the trial is stopped early for safety.

Stage II follows the response-adaptive randomization principle. This type of Monte Carlo algorithm uses the current parameter estimates to sample a cohort of m_2 dose combinations from the estimated dose-efficacy standardized density of $\widehat{\pi}_E(x,y)$ along the estimated MTD curve. Note that $\widehat{\pi}_E(x,y)$ uses the Bayes estimates of the dose-efficacy model parameters. Because stage II selects doses on the estimated MTD curve $\widehat{\mathcal{M}}$, and there is a one-to-one correspondence between $(x,y) \in \widehat{\mathcal{M}}$ and $x \in X'$, where $X' \subset X$, we may write $\widehat{\pi}_E(x,y) = \widehat{\pi}_E(x)$ for $(x,y) \in \widehat{\mathcal{M}}$. The standardized density of the estimated efficacy curve is $\widehat{\pi}_E^{\text{STDZ}}(x) = \frac{\widehat{\pi}_E(x)}{\int_{x\in X'} \widehat{\pi}_E(x)dx}$. A rejection sampling algorithm is then used to sample m_2 dose combinations from this density.

These steps are are described in Stage II of Algorithm 1. Note that in this stage, the MTD curve is updated after the DLT status of each cohort of m_2 patients is resolved and a new cohort of m_2 patients are randomly selected from the new MTD curve according to the standardized density $\widehat{\pi}_E(x,y)$.

Another contribution of this article with respect to [25, 10, 12] is that we do not require the first cohort of stage II to be homogeneously distributed along the estimated MTD curve. The justification is that, in this article, we do not have two disconnected stages and thus we use all the available stage I efficacy data to already have posterior estimates of $\beta_0, \ldots, \beta_5$ at the end of stage I as displayed in line 13 of Algorithm 1. This allows to allocate the first cohort of patients in stage II to dose combination following the adaptive randomization principle.

To facilitate the understanding of Algorithm 1, we define the variables that will be used. In stage I, c_1 denotes the current cohort indicator with each cohort containing 2 patients. For example, $c_1 = 2$ (i.e., the second cohort) contains the dose combinations from patients 3 and 4. In a generalized way, one cohort contains patients $(2c_1 - 1)$ and $(2c_1)$ and so, if $c_1 = 2$, we know that the second cohort contains the patients $2c_1 - 1 = 3$ and $2c_1 = 4$. Now let (x_{2c_1-1}, y_{2c_1-1}) and (x_{2c_1}, y_{2c_1}) denote the dose combinations to which the two patients from cohort c_1 were (or will be) allocated. Using the same example, if $c_1 = 2$, we see that $(x_{2c_1-1}, y_{2c_1-1}) = (x_3, y_3)$ and $(x_{2c_1}, y_{2c_1}) = (x_4, y_4)$ correspond to the dose combinations from patients 3 and 4. Also, let $\pi(\Gamma_{X|Y=y}|D_n)$ represent the posterior distribution of the MTD of drug X given that the level of drug Y is equal to y (i.e., given that Y is fixed) and given the data D_n (see equation (3) for the definition of the MTD). Last, let $\Pi^{-1}_{\Gamma_{X|Y=y}}(\alpha|D_n)$ denote the α-th percentile of $\pi(\Gamma_{X|Y=y}|D_n)$.

Algorithm 1 Stage I and stage II algorithms

– STAGE I –
In the first cohort ($c_1 = 1$) patients 1 and 2 receive the dose combination $(x_1, y_1) = (x_2, y_2) = (0, 0)$
In the second cohort ($c_1 = 2$) patients 3 and 4 receive doses (x_3, y_3) and (x_4, y_4) respectively, where $y_3 = y_1$, $x_4 = x_2$, x_3 is the α-th percentile of $\pi(\Gamma_{X|Y=y_1}|D_2)$, y_4 is the α-th percentile of $\pi(\Gamma_{Y|X=x_2}|D_2)$.
for $c_1 = 3 : C_1$ **do**
  if c_1 is an even number **then**
    Patient $2c_1 - 1$ receives the dose combination (x_{2c_1-1}, y_{2c_1-3}) where $x_{2c_1-1} = \Pi^{-1}_{\Gamma_{X|Y=y_{2c_1-3}}}(\alpha|D_{2c_1-2})$
    Patient $2c_1$ the dose combination (x_{2c_1-2}, y_{2c_1}), where $y_{2c_1} = \Pi^{-1}_{\Gamma_{Y|X=x_{2c_1-2}}}(\alpha|D_{2c_1-2})$
  else
    Patient $2c_1 - 1$ receives the dose combination (x_{2c_1-3}, y_{2c_1-1}) where $y_{2c_1-1} = \Pi^{-1}_{\Gamma_{Y|X=x_{2c_1-3}}}(\alpha|D_{2c_1-2})$
    Patient $2c_1$ receives the dose combination (x_{2c_1}, y_{2c_1-2}) where $x_{2c_1} = \Pi^{-1}_{\Gamma_{X|Y=y_{2c_1-2}}}(\alpha|D_{2c_1-2})$
  end if
end for
– STAGE II –
Calculate the posterior median of the parameters $(\widehat{\rho}_{00}, \widehat{\rho}_{10}, \widehat{\rho}_{01}, \widehat{\alpha}_3, \widehat{\beta}_0, \widehat{\beta}_1, \widehat{\beta}_2, \widehat{\beta}_3, \widehat{\beta}_4, \widehat{\beta}_5)$ given data D_{N_1}
for $c_2 = 1 : C_2$ **do**
  Calculate estimated MTD set $\widehat{\mathcal{M}} = \left\{(x, y) : y = \left(\frac{F^{-1}(\theta_Z) - F^{-1}(\widehat{\rho}_{00}) - (F^{-1}(\widehat{\rho}_{10}) - F^{-1}(\widehat{\rho}_{00}))x}{(F^{-1}(\widehat{\rho}_{01}) - F^{-1}(\widehat{\rho}_{00})) + \widehat{\alpha}_3 x}\right)\right\}$
  Generate a sample of dose combinations of size m_2 that belong to $\widehat{\mathcal{M}}$ from the (estimated) standardized density of $\widehat{\pi}_E(x, y)$, and assign it to the subsequent cohort of m_2 patients
  Calculate the posterior median of the parameters $(\widehat{\rho}_{00}, \widehat{\rho}_{10}, \widehat{\rho}_{01}, \widehat{\alpha}_3, \widehat{\beta}_0, \widehat{\beta}_1, \widehat{\beta}_2, \widehat{\beta}_3, \widehat{\beta}_4, \widehat{\beta}_5)$ given data $D_{N_1 + c_2 \times m_2}$
end for

Stage II can also be viewed as an extension of a response-adaptive randomization procedure with a finite number of doses (see [2]) to a response-adaptive randomization procedure with an infinite number of dose combinations. Mind that one could easily adapt Algorithm 1 to have a single stage by starting directly with the adaptive randomization part of the design (i.e., stage II). The dose finding algorithm contains the following stopping rules for safety and futility:

- **Futility stopping rule**
  For ethical considerations and to avoid exposing patients to sub-therapeutic dose combinations, we would stop the trial for futility if

$$\max_{(x,y)\in\widehat{\mathcal{M}}_n} \left[P(F(\beta_0+\beta_1 x+\beta_2 y+\beta_3 xy+\beta_4 x^2+\beta_5 y^2)>\theta_E \mid D_n)\right] < \delta_0,$$

  where δ_0 is a pre-specified threshold. For the purpose of this article, we choose $\delta_0 = 0.1$. Mind that this stopping rule applies only after the initial cohort of patients in stage II.

- **Safety stopping rule**
  The design contains two stopping rules for safety, one for stage I and a different one for stage II. During stage I, we would stop the trial if

$$P(\pi_Z(x=0,y=0)>(\theta+0.1) \mid D_n) > \delta_{\theta_1}, \tag{7}$$

  where $\delta_{\theta_1} = 0.5$. In contrast, during stage II we would stop the trial if

$$P(\Theta>(\theta+0.1) \mid D_n) > \delta_{\theta_2}, \tag{8}$$

  where Θ represents the rate of DLTs for both stages of the design regardless of dose and $\delta_{\theta_2} = 0.7$ represents the confidence level (i.e., 70%) that a prospective trial results in an excessive DLT rate. A non-informative Jeffrey's prior Beta$(0.5, 0.5)$ is placed on the parameter Θ. Mind that equation (8) does not depend on a pre-specified dose combination like equation (7). The rational behind this choice is that it would be sub-optimal to condition the stage II safety stopping rule on a pre-specifed dose combination, even if selected at the end of stage I, given that, during stage II, the MTD curve is updated after each new cohort of patients.

4 Simulation study

4.1 Operating characteristics

Stages I and II are evaluated through the following operating characteristics:

1. Pointwise average bias of the estimated MTD curve:
   Following [26], let $\mathcal{M}$ be the true MTD curve of a pre-specified scenario and $\widehat{\mathcal{M}}_j$ be the estimated MTD set at the j-th simulated trial, $j = 1, \ldots, J$. The pointwise average bias when estimating the MTD is defined as follows. For every point $(x, y) \in \mathcal{M}$, let

$$d^{(j)}_{(x,y)} = \text{sign}(y'-y) \min_{(x^*,y^*):(x^*,y^*)\in\widehat{\mathcal{M}}_j} \{(x-x^*)^2+(y-y^*)^2\}^{1/2}, \tag{9}$$

   where y' is such that $(x, y') \in \widehat{\mathcal{M}}_j$. This is the minimum relative distance of the point (x, y) of $\mathcal{M}$ to $\widehat{\mathcal{M}}_j$. The pointwise average bias is thus calculated as $d_{(x,y)} = \frac{1}{J}\sum_{j=1}^{J} d^{(j)}_{(x,y)}$.

2. Percentage of correct MTD curve recommendation:
   Following [26], we define the percentage of correct MTD recommendation as the pointwise percentage of trial for which the minimum distance of the point (x, y) on $\mathcal{M}$ to $\widehat{\mathcal{M}}_j$ is no more than $(100 \times p)\%$ of the true MTD. It is calculated using the definition of the minimum relative distance given by equation (9) and it is defined as

$$R_{(x,y)} = \frac{1}{J}\sum_{j=1}^{J} I\{|d^{(j)}_{(x,y)}| \le p\Delta(x,y)\},$$

   where $0 < p < 1$, and $\Delta(x, y)$ represents the euclidean distance between the minimum dose combination $(0, 0)$ and the point (x, y) on the MTD curve. Intuitively, this metric is equivalent to drawing a circle with radius $p\Delta(x, y)$ and calculating the percentage of trials with $\widehat{\mathcal{M}}_j$ falling inside the circle, where p is a tolerance parameter.

Figure 1: True dose-toxicity profiles. The black dot in the dose combination of cisplatin and cabazitaxel 15/75 represents the MTD found by [15] using a '3+3' design.

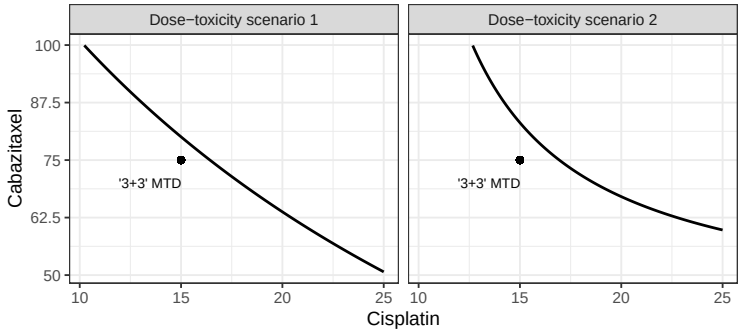

3. Average DLT rate and percentage of trials with DLT rate higher than $\theta_Z + 0.1$.
4. Distribution of recommended optimal dose combinations.
5. Proportion of recommended optimal dose combinations with true probability of efficacy above θ_E.
6. Bayesian power (type-I error probability under H_0):

$$\text{Power} \approx \frac{1}{J}\sum_{j=1}^{J} \mathbf{1}\left\{ \max_{(x,y)\in\widehat{\mathcal{M}}_N} \left(P(F(\beta_{0j} + \beta_{1j}x + \beta_{2j}y + \beta_{3j}xy + \beta_{4j}x^2 + \beta_{5j}y^2) > \theta_E | D_N) \right) > \delta_u \right\},$$

where '$\mathbf{1}(.)$' represents an indicator function.

For scenarios favoring the null hypothesis as defined in (4), the above formula represents the Bayesian type-I error probability.

7. Proportion of patients in stage II allocated to dose combinations with true probability of efficacy above θ_E.
8. Average posterior probability of early stopping for futility and safety.

4.2 Dose-toxicity and dose-efficacy scenarios

Following [25], we use the two dose-toxicity scenarios that were expected by the principal investigator of the cisplatin-cabazitaxel trial, which are displayed in Figure 1. To generate dose-toxicity scenarios 1 and 2, we have used the parameter values $\{\rho_{00} = 1 \times e^{-7}, \rho_{01} = 0.3, \rho_{10} = 0.3, \alpha_3 = 2\}$ and $\{\rho_{00} = 1 \times e^{-5}, \rho_{01} = 0.01, \rho_{10} = 0.005, \alpha_3 = 9\}$, respectively. The target probability of DLT is $\theta = 0.33$.

Within each of these dose-toxicity scenarios, we study four different dose-efficacy profiles, both under the alternative (H_1) and the null (H_0) hypotheses, that place the most efficacious dose combination in a different region of the MTD curve. These profiles are obtained using model (2) for (x, y) in the true MTD curve as defined by dose-toxicity scenarios 1 and 2. These are displayed in Figure 2. The parameter values used to generate the dose-efficacy profiles under H_1 and H_0 are displayed in Table 1. It is important to mention that all the dose-efficacy scenarios have been generated with the parameters $\beta_4, \beta_5 = 0$ (i.e., using monotonic dose-efficacy profiles) since we are combining two cytotoxic agents. These two parameters are incorporated for the sole purpose of adding some flexibility since, as mentioned in section 3.2, in the development of this article we have found that, because of the assumed synergistic effect of the compounds in terms of the dose-toxicity profile, an efficacy model with only linear terms was not able to fully capture the dose-efficacy profile in one of the dose-efficacy scenarios we present (scenario 3).

The probability of response (clinical benefit) for the standard of care treatment is $\theta_E = 0.15$, and the cohort size in stage II is $m_2 = 5$. Under H_1, we assume an effect size of 0.25 (i.e., in all dose-efficacy profiles, the highest probability of efficacy is equal to $\theta_E + 0.25 = 0.4$). Under H_0, the highest probability of efficacy is equal to θ_E. The sample sizes for Stages I and II are $N_1 = N_2 = 30$ and we simulated $J = 1000$ trials using Algorithm 1. The DLT and efficacy responses were generated from models (1) and (2), respectively.

Figure 2: True dose-efficacy profiles under the null and alternative hypotheses varying with the dose of Cisplatin along the true MTD curves. The red dotted line represents θ_E and the green dotted line represent the highest probability of efficacy.

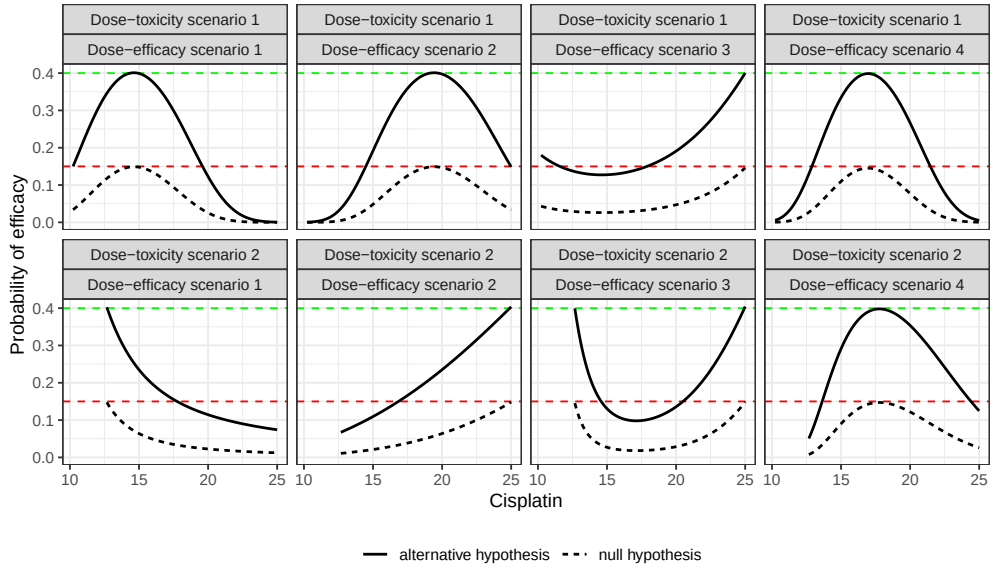

Table 1: Dose-efficacy parameters values used to generate the dose-efficacy profiles under H_0 and H_1 presented in Figure 2. The parameters $\beta_0(H_0)$ and $\beta_0(H_1)$ belong to the null and alternative hypotheses, respectively. The parameters $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5$ are the same under the null and alternative hypotheses.

Dose-toxicity scenario	1				2			
Dose-efficacy scenario	1	2	3	4	1	2	3	4
$\beta_0(H_0)$	-6.3	-6.3	-7.3	-4.8	-2.8	-2.8	-6.6	-7.28
$\beta_0(H_1)$	-5.51	-5.51	-6.5	-4	-2	-2	-5.8	-6.49
β_1	2	4.3	6.17	1.25	0.05	1.55	4.63	0.2
β_2	4.3	2	5.5	1.25	1.57	0.05	4.73	0.2
β_3	10	10	0	12	1	1	0	26
β_4	0	0	0	0	0	0	0	0
β_5	0	0	0	0	0	0	0	0

4.2.1 Results

In Figure 3, we present the pointwise average bias of the estimated MTD curve at the end of the trial with respect to the true MTD curve under the null (dashed line) and alternative (solid line) hypotheses, respectively. In line with the results presented by [26, 10], the design yields a good estimation of the MTD curve, specially around its central region whereas this estimation is less precise at the extremes of the MTD curve. However, because of the use of all toxicity data from stages I and II, the average bias we observe in the MTD curve estimation is, overall, lower than the values reported in the supplementary material of [10] with also vague prior distributions. In Figures S1 and S2-S3 in the supplementary material, we display the estimated MTD curves and the percentage of correct MTD recommendation under the null and alternative hypotheses, respectively. In most scenarios, the pointwise percent selection varies between 60% and 90% except near the edges of the true MTD curve using a threshold value of $p = 0.1$ and between 80% and 100% using a tolerance threshold $p = 0.2$. These results are consistent with the conclusions we draw from the pointwise average bias.

In Table 2 we present the safety operating characteristics and we observe that, on average, the DLT rate is between 29 - 33% and the percentage of trials with proportion of DLT above $\theta_Z + 0.1$ ranges between 0 - 0.2%. Therefore, the design is considered safe.

In Figure 4, we show the probability of rejecting the null hypothesis, which corresponds to the (Bayesian) power under the alternative hypothesis and to the (Bayesian) type-I error probability under the null hypothesis. This probability

Figure 3: Pointwise average bias of the estimated MTD curve with respect to the true MTD curve under the null (dashed line) and alternative (solid line) hypotheses, respectively. Results are displayed from the point of view of the standardized Cisplatin.

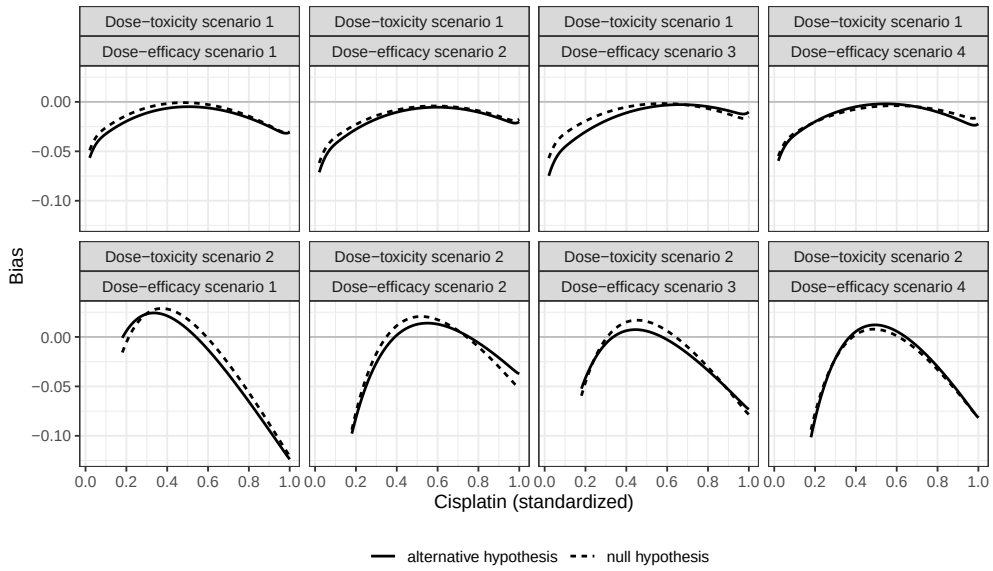

Table 2: Average DLT rate and percentage of trials with DLT above $\theta_Z + 0.1$.

Dose-toxicity scenario	1				2			
Dose-efficacy scenario	1	2	3	4	1	2	3	4
Average DLT rate (H_0)	0.333	0.337	0.333	0.336	0.289	0.304	0.299	0.330
Average DLT rate (H_1)	0.336	0.339	0.337	0.337	0.289	0.308	0.300	0.332
Percentage (%) of trial with DLT rate above $\theta_Z + 0.1$ (H_0)	0.0	0.1	0.2	0.1	0.0	0.0	0.0	0.1
Percentage (%) of trial with DLT rate above $\theta_Z + 0.1$ (H_1)	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.1

ranges between 0.7 and 0.85 under the alternative hypothesis (i.e., power), and between 0.07 and 0.29 under the null hypothesis for all but one scenario (i.e., type-I error). These values are in line with the results presented by Tighiouart (2019) [25] and Jiménez et al. (2020) [10]. We observe one interesting result under the null hypothesis in dose-toxicity scenario 2 - dose-efficacy scenario 4, where the type-I error probability is close to 0.5. This value is a consequence of a slight overestimation of the MTD curve in its central region (see Figure 3). This behavior is not present in dose-toxicity scenario 1 - dose-efficacy scenario 4 where the MTD curve is practically unbiased in its central region. By overestimating the MTD, we are no longer having a maximum probability of 0.4 and 0.15 under the alternative and null hypotheses, respectively. This causes an increase of both the power and the type-I error probability. In fact, dose-toxicity scenario 2, dose-efficacy scenario 4 is the scenario that achieves both the highest power and highest type-I error probability because of the overestimation of the MTD curve in the region of highest efficacy. By the same token, an underestimation of the MTD curve in the region of highest efficacy will lead to both lower power and lower type-I error probability.

In Figure 5, we present the density (i.e., the distribution) of the dose combinations recommended at the end of each simulated trial as the optimal dose combination. In this article, the MTD curve is never fixed which translates into recommended optimal dose combinations above and below the true MTD curve. However, we can see how the density of these recommended optimal dose combinations is, generally speaking, very close to the true MTD curve, with the region of higher mass (i.e., bright yellow) containing the true optimal dose combination. Mind that in dose-toxicity

Figure 4: Bayesian power and type-I error probability for scenarios under the alternative and null hypotheses, respectively.

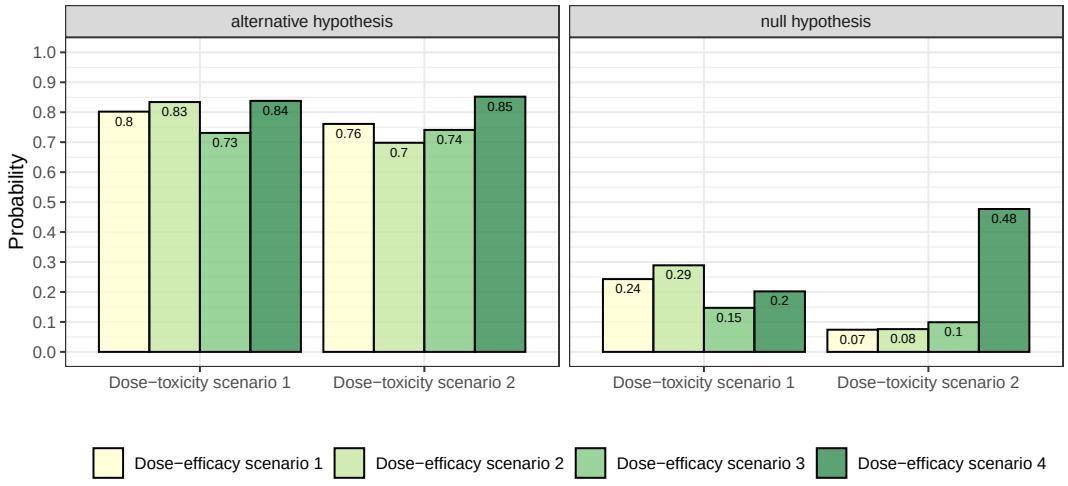

scenario 2 - dose-efficacy scenario 3, the true probability of efficacy at the left extreme of the true MTD curve is intended to be very close to the one from the optimal dose combination located at the right extreme of the MTD curve (see Figure 2). The bi-modality observed in the density of the recommended optimal dose combinations in this scenario shows that the design is able to capture the characteristic of this dose-efficacy scenario.

In Figure S4 (supplementary material), we present the distribution of dose combinations allocated to all patients during stage II, which works under the adaptive randomization principle. We display all scenarios under the alternative hypothesis only. Results clearly show how the mode of the distribution corresponds to the dose combination with highest probability of efficacy under each scenario. Moreover, the proportion of patients during stage II that were allocated to dose combinations with a true probability of efficacy above θ_E ranges between 69-77% (see Figure S5 in the supplementary material).This result also holds under the null hypothesis (results not shown).

In Figure S6, we present the average posterior probability of early stopping for the stopping rules we pre-specified in section 3.4. Under the alternative hypothesis, we see that the probability of early stopping for futility is always below 0.02. Under the null hypothesis, this probability ranges generally between 0.067 and 0.27, depending on the scenario. We see how the probability of early stopping for futility in dose-toxicity scenario 2 - dose-efficacy scenario 4 is much lower (0.067) than in other scenarios because of the overestimation of the central region of the MTD curve and the allocation of patients with higher probability of efficacy. In terms of the probability of early stopping for safety, we see that under both the null and alternative hypotheses, this probability is practically equal to 0. The values presented here are in line with the operating characteristics from [25].

4.2.2 Model Robustness

In Figure S9 of the supplementary material, we present the power and type-I error obtained under model misspecification. In this assessment, the true underlying dose-toxicity and dose-efficacy relationships' link functions are different from the probit link function used in the design's working marginal models (see section 3.2). In simple terms, toxicity and efficacy data will be generated using a common link function (e.g., logistic) and the design will proceed under the (wrong) assumption that such common link function is a probit link function. The purpose of this robustness assessment is to test whether the design is able to maintain the operating characteristics presented in previous sections under small deviations between the true underlying marginal models and the working marginal models. To limit the number of comparisons, the toxicity and efficacy marginal models employ the same link function.

For the assessment, we reproduced the first dose-toxicity scenario presented in Figure 1 and the first two dose-efficacy scenarios presented in Figure 2 using the logistic and the complementary log-log link functions. Because the link functions are different, the dose-toxicity and dose-efficacy surfaces have to be different, specially around the extremes of the the surfaces. We refer the reader to [27] to see how the dose-toxicity surfaces differ using different link functions. Under model misspecification, the true models' parameters are selected in a way that there is an overlap of the toxicity

Figure 5: Density of the dose combinations recommended at the end of each simulated trial as the optimal dose combination under the alternative hypothesis. The solid black line represents the true MTD curve of each dose-toxicity scenario and the red dot represents the true optimal dose combination.

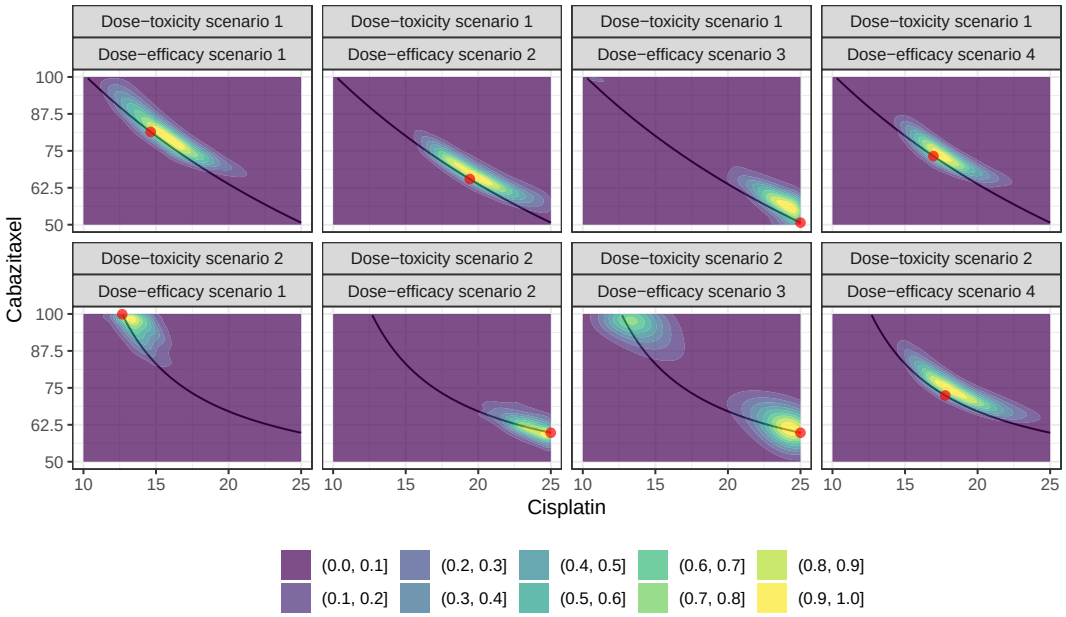

and efficacy profiles of all link functions in the region where the optimal dose combination is located. In Figures S7 and S8 of the supplementary material we display the MTD curves as well as the efficacy profiles on the MTD curves generated with the logistic and complementary log-log link functions.

The results presented in Figure S9 show that our design is able to maintain its performance under mode misspecification, with power and type-I error values with minor differences between the three tested link functions. This is not unexpected since the quadratic terms included in the dose-efficacy model allow for higher flexibility.

Last, and for illustration purposes, in Figures S10 and S11 of the supplementary material we present the contours of the marginal toxicity and efficacy models, respectively, with the logit, logistic and complementary log-log link functions. For this illustration we selected dose-toxicity scenario 1 and dose-efficacy scenario under the alternative hypothesis. In this robustness assessment, the true toxicity and efficacy model parameters are selected so that the MTD curves and the efficacy profiles within the MTD curves match across all the link functions (see Figures S7 and S8). However, outside of the MTD curves, we observe how each link function induces slightly different dose-toxicity and dose-efficacy surfaces.

5 Discussion

We present a Bayesian two-stage design for cancer phase I-II clinical trials combining two cytotoxic agents with continuous dose levels. Motivated by the cisplatin-cabazitaxel trial, we extend the work of [25] by allowing the continuous update of the maximum tolerated dose (MTD) curve throughout the entire study as well as the use of stage I efficacy data under the assumption of a single patient population across stages. In stage I we aim to estimate the MTD curve following the escalation with overdose control (EWOC) principle and enrolling up to 30 patients. In stage II we model treatment efficacy as a binary indicator of treatment response using a probit quadratic model, and we enrol 30 additional patients using response adaptive randomization. The rational for quadratic terms in a setting of two cytotoxic agents is that, under the presence of (synergistic) interaction between the compounds, the dose-efficacy profile along the MTD can become quite complex, and we found that a linear model was not able to correctly identify the region with the most efficacious dose combinations in certain scenarios. However, under the presence of non-interactive compounds the dose-efficacy profiles cannot surpass a certain level of complexity and thus a linear dose-efficacy model would be sufficient to identify the most efficacious dose combinations.

We study the properties of our design under the two (synergistic) dose-toxicity scenarios that were expected by the principal investigator of the cisplatin-cabazitaxel trial introduced by [25] with 30 patients in each stage and an effect size of 0.25 in dose-toxicity scenarios under the alternative hypothesis. We study four different dose-efficacy profiles that place the most efficacious dose combination in different regions of the MTD curve (i.e., left extreme, right extreme, central and both extremes). The estimated MTD curve tends to be slightly biased around its extremes and less biased around its central region as reported by [26, 10]. In fact, because of the continuous update of the MTD curve during stage II, we have observed a slight reduction of the overall estimation bias with respect to the results presented by [10] where the MTD curve was only updated in stage I. The proposed probit quadratic efficacy model has a good performance and is able to fully capture the different dose-efficacy profiles along the true MTD curve. This is fully reflected in Figure 5, dose-toxicity scenario 2 - dose-efficacy scenario 3, where the probability of efficacy if high at both extremes of the true MTD curve. As a matter of fact, in the development of this research, we found that a marginal efficacy model with only linear terms was not able to capture this particular profile which is heavily conditioned by the synergistic nature of the MTD curve. Thus, adding quadratic terms allows to better capture any dose-efficacy profile along any MTD curve. Overall, the design is safe and yields good operating characteristics across scenarios with power and type-I error rates that are consistent with the values reported by [25, 10]. Another contribution of this article is that we account for the uncertainty of the estimated MTD curve during stage II since we continuously update it after each cohort of patients. In fact, we have observed its behavior in stage II, which was not explored in [25, 10]. Results show that overall bias in the MTD curve is also driven by the adaptive randomization part of the design where we allocate patients to dose combination that are likely to be efficacious along the estimated MTD curve. This type of patient allocation unbalances the amount of data that we have in different regions of the estimated MTD curve, which could potentially damage the precision of the estimation in the search of higher efficacy. We believe these findings are coherent with the fact that we jointly work with i) the uncertainty around the MTD curve estimation and ii) the uncertainty around the estimation of the dose-efficacy profile. With these conclusions in mind, we strongly recommend using informative prior distributions whenever possible, and/or a slightly more conservative EWOC criterion (i.e., a lower percentile) either throughout the entire study or in stage II only. An alternative approach to hypothesis testing is to report the percent selection of correct optimal dose combinations, as is traditionally done in phase I/II dose finding studies. Since we are dealing with continuous dose levels, a reasonable definition of this metric would be the percentage of recommended optimal dose combination with (true) $\pi_Z(x, y) \leq \theta_Z + \delta_1$ and (true) $\pi_E(x, y) \geq \theta_E^* - \delta_2$, where θ_E^* represents the true maximum probability of efficacy of a given scenario (i.e., $\theta_E^* = 0.4$ and $\theta_E^* = 0.15$ in the dose-efficacy scenarios that we considered) and δ_1, δ_2 are pre-specified positive threshold values.

In this article, we have not used the informative prior distributions used by [25] and therefore our estimation of the MTD curve is consequently less precise. The operating characteristics presented in this article can be considered as the baseline operating characteristics this design can have given the use of vague prior distributions and relatively low sample size. By increasing the sample size or using more informative prior distributions in the dose-toxicity model, the design's operating characteristics are expected to improve.

Finally, model robustness was explored by using logistic and complementary log-log link functions to generate toxicity and efficacy data. Results show that the proposed design is robust to model misspecification under small deviations of the true model from the working model.

On the application of the proposed design to settings with other type of compounds (i.e., immunotherapies or targeted therapies), we note that major changes would be needed to adapt the design to the characteristics of such compounds. The most relevant adaption would be the definition of the optimal dose combination combination. With cytotoxic agents, it is generally assumed that the optimal dose combination will be somewhere along the MTD curve. With immunotherapies or targeted therapies, this is not guaranteed and utility functions become necessary to find the right balance between toxicity and efficacy (see e.g., [14]). Also, because we would no longer restrict the search to the estimated MTD, we would need to extend the stage II adaptive randomization algorithm so that we are allowed to allocate patients to any dose combination with estimated probability of toxicity below the target probability of toxicity.

Also, the design could be implemented in a setting with discrete dose combinations either with cytotoxics or immunotherapies with minor modifications with respect to the version of the design with continuous dose levels. To do so, in stage I, the continuous dose combinations would be rounded to nearest discrete dose combinations whereas, in stage II, patients would only be allocated to the initial set of discrete combinations based on the updated probability of efficacy estimates (i.e., the higher estimated probability of efficacy of a dose combination the higher the probability of being allocated to that dose combination).

These two extensions are part of our research plan and we intent to evaluate the operating characteristics of the proposed design under these settings.

Data availability

The `R` and `JAGS` scripts necessary to fully reproduce the results presented in this article are available at `https://github.com/jjimenezm1989/Phase-I-II-design-combining-two-cytotoxics`

Funding

Mourad Tighiouart is funded by NIH the National Center for Advancing Translational Sciences (NCATS) UCLA CTSI (UL1 TR001881-01) and NCI P01 CA233452-02.

Disclaimer

José L. Jiménez is employed by Novartis Pharma A.G. who provided support in the form of salary for the author, but did not have any additional role in the preparation of the manuscript. Also, the views expressed in this publication are those of the authors and should not be attributed to any of the funding institutions or organisations to which the authors are affiliated.

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] Scott M Berry, Bradley P Carlin, J Jack Lee, and Peter Muller. *Bayesian adaptive methods for clinical trials*. CRC press, 2010.

[3] 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.

[4] Chunyan Cai, Ying Yuan, and Yuan Ji. A bayesian dose-finding design for oncology clinical trials of combinational biological agents. *Journal of the Royal Statistical Society. Series C, Applied statistics*, 63(1):159, 2014.

[5] 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.

[6] 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.

[7] 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.

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

[9] Anastasia Ivanova, Ken Liu, Ellen Snyder, and Duane Snavely. An adaptive design for identifying the dose with the best efficacy/tolerability profile with application to a crossover dose-finding study. *Statistics in medicine*, 28(24):2941–2951, 2009.

[10] José L Jiménez, Sungjin Kim, and Mourad Tighiouart. A bayesian seamless phase I–II trial design with two stages for cancer clinical trials with drug combinations. *Biometrical Journal*, 62(5):1300–1314, 2020.

[11] Jose L Jimenez, Mourad Tighiouart, and Mauro Gasparini. Cancer phase I trial design using drug combinations when a fraction of dose limiting toxicities is attributable to one or more agents. *Biometrical Journal*, 61(2):319–332, 2019.

[12] José L Jiménez and Haiyan Zheng. A bayesian adaptive design for dual-agent phase I-II cancer clinical trials combining efficacy data across stages. *arXiv preprint arXiv:2106.08277*, 2021.

[13] 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.

[14] Suyu Liu, Beibei Guo, and Ying Yuan. A bayesian phase I/II trial design for immunotherapy. *Journal of the American Statistical Association*, 113(523):1016–1027, 2018.

[15] 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.

[16] Jiaying Lyu, Yuan Ji, Naiqing Zhao, and Daniel VT Catenacci. Aaa: triple adaptive bayesian designs for the identification of optimal dose combinations in dual-agent dose finding trials. *Journal of the Royal Statistical Society. Series C, Applied statistics*, 68(2):385, 2019.

[17] Martyn Plummer et al. Jags: A program for analysis of bayesian graphical models using gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing*, volume 124, pages 1–10. Vienna, Austria., 2003.

[18] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2019.

[19] 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.

[20] 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.

[21] Yun Shi and Guosheng Yin. Escalation with overdose control for phase I drug-combination trials. *Statistics in medicine*, 32(25):4400–4412, 2013.

[22] 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.

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

[24] 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.

[25] 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)*, 68(1):235–250, 2019.

[26] 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.

[27] Mourad Tighiouart, Steven Piantadosi, and André Rogatko. Dose finding with drug combinations in cancer phase I clinical trials using conditional escalation with overdose control. *Statistics in medicine*, 33(22):3815–3829, 2014.

[28] 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.

[29] 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.

[30] 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.

[31] Graham M Wheeler, Michael J Sweeting, and Adrian P Mander. Toxicity-dependent feasibility bounds for the escalation with overdose control approach in phase i cancer trials. *Statistics in medicine*, 36(16):2499–2513, 2017.

[32] 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.

[33] 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.

[34] Ying Yuan and Guosheng Yin. Bayesian phase I/II adaptively randomized oncology trials with combined drugs. *The annals of applied statistics*, 5(2A):924, 2011.

[35] Liangcai Zhang and Ying Yuan. A practical bayesian design to identify the maximum tolerated dose contour for drug combination trials. *Statistics in medicine*, 35(27):4924–4936, 2016.