Research Article

Statistics in Medicine

Received 14 May 2012, Accepted 23 December 2013 Published online 27 January 2014 in Wiley Online Library

(wileyonlinelibrary.com) DOI: 10.1002/sim.6094

Competing designs for drug combination in phase I dose-finding clinical trials

M.-K. Riviere,[a,b*†] F. Dubois[b] and S. Zohar[a]

The aim of phase I combination dose-finding studies in oncology is to estimate one or several maximum tolerated doses (MTDs) from a set of available dose levels of two or more agents. Combining several agents can indeed increase the overall anti-tumor action but at the same time also increase the toxicity. It is, however, unreasonable to assume the same dose–toxicity relationship for the combination as for the simple addition of each single agent because of a potential antagonist or synergistic effect. Therefore, using single-agent dose-finding methods for combination therapies is not appropriate.

In recent years, several authors have proposed novel dose-finding designs for combination studies, which use either algorithm-based or model-based methods. The aim of our work was to compare, via a simulation study, six dose-finding methods for combinations proposed in recent years. We chose eight scenarios that differ in terms of the number and location of the true MTD(s) in the combination space. We then compared the performance of each design in terms of correct combination selection, patient allocation, and mean number of observed toxicities during the trials. Our results showed that the model-based methods performed better than the algorithm-based ones. However, none of the compared model-based designs gave consistently better results than the others. Copyright © 2014 John Wiley & Sons, Ltd.

Keywords: dose-finding studies; phase I trial; drug combination; Bayesian inference; oncology; cytotoxic

1. Introduction

The majority of phase I cytotoxic dose-finding studies seek to establish a dose high enough to be able to observe potential efficacy while maintaining the toxicity rate within certain pre-defined acceptable limits. In oncology, phase I studies focus on determining the maximum tolerated dose (MTD) that will be used in further phase II clinical trials of which the central interest is on potential efficacy. Following Storer [1], the statistical formulation of the problem is to select a dose level from several available doses, with a toxicity probability closest to a given target [2–4]. Most statistical model-based or algorithm-based methods were developed for cytotoxic single-agent phase I dose-finding clinical trials [5]. In this context, it is assumed that the toxicity of a single agent is monotonic and increases with the dose, as does the efficacy.

In the field of oncology, it is currently rare to find new molecules that perform better than existing therapeutic strategies. When combining two or more agents, there may be a potential synergistic effect in terms of efficacy. That is why investigators wish to increase overall anti-tumor action and survival by combining several agents, either cytotoxics or targeted molecules, or both. As a result, it is difficult to suppose that each molecule will act independently in terms of toxicity. In dose-finding studies, physicians aim to gradually increase toxicity during the dose-escalation procedure. However, when combining several agents, the ordering of toxicity probabilities is not fully known. For instance, the combination of two cytotoxics can induce an ordered subset of toxicity (Figure 1(a)). Even when a partial ordering is known, it is still difficult to decide how to escalate or de-escalate a combination of doses. Indeed, on a diagonal, there is no knowledge about which combination is more toxic; it is not known prior to the trial

[a]*INSERM, U1138, Equipe 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6, Paris, France*
[b]*IRIS (Institut de Recherches Internationales Servier), Suresnes, France*
**Correspondence to: M-K. Riviere, INSERM, U1138, Equipe 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6, Paris, France.*
[†]*E-mail: marie-karelle.riviere@crc.jussieu.fr*

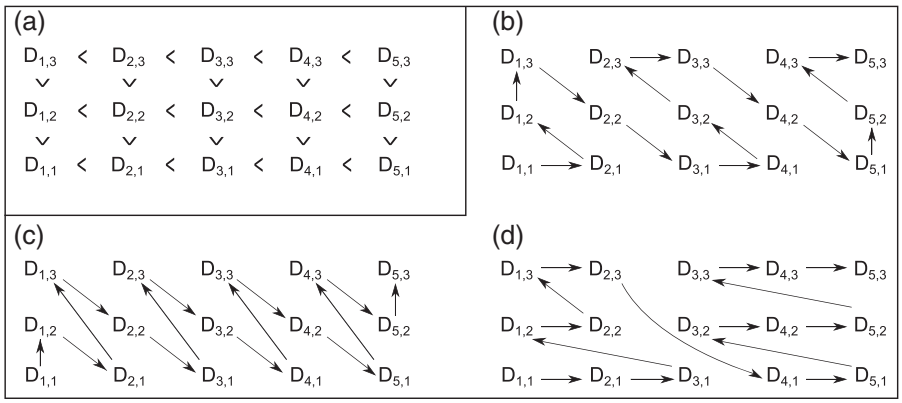

Figure 1. (a) Partial order relationships between combinations. (b) Three possible orderings consistent with the partial order retained for POCRM among all possible orderings for the simulation study.

which of $D_{1,2}$ and $D_{2,1}$ is more toxic. Therefore, it is senseless to use single-agent dose-finding methods in combination studies.

Several authors have recently addressed this issue by proposing new methods for combination studies, which are either algorithm-based or model-based. Ivanova and Wang have proposed an 'up-and-down algorithm-type' method with isotonic regression for the estimation of a set of possible MTDs [6]. Furthermore, Wang and Ivanova have developed a three-parameter model-based method for which the parameters are estimated using Bayesian inference [7]. Mandrekar *et al*. have proposed incorporating both the toxicity and efficacy of each agent into the identification of an optimal dosing region for the combination using a continuation ratio model to separate each agent's toxicity and efficacy curves [8, 9]. Fan *et al*. have proposed a '2 + 4 + 3' algorithm-based dose-allocation scheme as well as the performance of two-dimensional isotonic regression to estimate the MTD [10]. Yin and Yuan have developed a Bayesian adaptive design based on latent 2×2 tables in which the combinations' toxicity probabilities in the two-dimensional space are estimated using a Gumbel-type model [11]. Additionally, Yin and Yuan have extended their method by changing to a copula-type model to simulate the effect of two or more drugs in combination [12]. Bailey *et al*. have introduced a second agent as a covariate in a logistic model [13]. Recently, Wages *et al*. have considered a continual reassessment method (CRM) based approach considering different orderings with partial order between combinations. In this case, the MTD is estimated for the order associated with the highest model-selection criterion [14].

For this simulation study, we focused on designs dealing only with toxicity. We selected six designs, which were either algorithm-based or model-based, and whose statistical estimation methods and allocation rules differed. The aim of this work was to provide statisticians involved in dose-finding studies with tools to evaluate combinations in order to select the most suitable design according to clinical trial's combination and indication. We selected the methods of Ivanova and Wang [6], Ivanova and Kim [18], Wang and Ivanova [7], Yin and Yuan [11, 12], and Wages *et al*. [14], which we consider representative of the methods that can be found in the literature. We had initially considered the selection of only one MTD at the end of the trial but then extended the comparison to the selection of multiple MTDs (Section 3).

2. Methods

2.1. Notations

Let there be a two-agent combination used in a phase I dose-finding clinical trial for which the dose–toxicity relationship is monotonic with respect to both dose levels. Let $D_{j,k}$ denote the dose level of a combination in which j refers to agent 1 ($j = 1, \ldots, J$), and k to agent 2 ($k = 1, \ldots, K$).

Toxicity refers to a dose-limiting toxicity (DLT), that is, an adverse event of grade 3 or higher. We represent the observation of a toxicity for patient i ($i = 1, \ldots, N$) by a Bernoulli random variable y_i equal to 1 if a DLT is observed for patient i and 0 otherwise. Let us assume that the combination dose $D_{j,k}$ is administered to $n_{j,k}$ patients and that we observe a total of $t_{j,k}$ DLTs at that combination dose

level. We then denote the observed proportion of DLTs by $\varphi_{j,k}=\frac{t_{j,k}}{n_{j,k}}$ and define $\pi_{j,k}$ as the toxicity probabilities.

The target (probability of) toxicity, θ, is fixed prior to trial onset as well as the initial guesses of toxicity (also called working model or skeleton). The combination X_i received by patient i $(i=1,\ldots,N)$ can be seen as a random variable taking values $x_i\in\{\mathrm{D}_{j,k};j=1,\ldots,J;k=1,\ldots,K\}$. For simplicity purposes, the selected combination will be referred to as an MTD in order to maintain the same designation as in single-agent trials.

2.2. *Up-and-down design for combinations* [6]

Ivanova and Wang proposed an up-and-down design for two-agent combinations associated with an isotonic regression for the estimation of the MTD [6]. The aim of this method was to identify a set of MTDs for each dose of agent 2. Within this context, we proposed some modifications in order to ensure comparability with other methods. This two-dimensional approach was based on a non-parametric design, which is an extension of the Narayana group's design [15, 16]. If the last allocation was at combination $\mathrm{D}_{j,k}$, the dose-allocation rule for the next combination is defined as follows: (i) $\mathrm{D}_{j+1,k}$ if $\varphi_{j,k}<\theta$ and there was no toxicity observed in the last cohort; (ii) $\mathrm{D}_{j-1,k}$ if $\varphi_{j,k}>\theta$ and there was at least one toxicity in the last cohort; (iii) $\mathrm{D}_{j+1,k-1}$ if $\varphi_{j,k}<\theta$ and there was at least one toxicity in the last cohort; and (iv) $\mathrm{D}_{j-1,k+1}$ if $\varphi_{j,k}>\theta$ and there was no toxicity in the last cohort.

The number of patients assigned to the lowest level of agent 2 is restricted to $\frac{N}{K}$ to enable the method to explore other levels of agent 2. Ivanova and Wang [6] proposed a 'start-up phase' in order to gather enough information before estimating $\varphi_{j,k}$'s. The start-up phase is conducted according to the following algorithm: (i) if, in the last cohort, no toxicity was observed, agent 1 is increased by one dose level; and (ii) if, in the last cohort, at least one toxicity was observed, agent 1 is decreased by two dose levels and agent 2 increased by one dose level. This process is repeated until all levels of agent 2 have been explored, and alternative combinations are proposed when reaching the boundary of the combination space.

When the overall sample size is reached, the estimate of the set of maximum tolerated combinations is calculated after using bivariate isotonic regression [17].

In order to ensure comparability between all of the methods presented in this manuscript, a decision rule with the selection of one MTD was proposed. The recommended combination at the end of the trial was the one with the toxicity probability closest to the target after isotonic regression; and, if there were several, then the one with the highest level of agent 2 was recommended. Moreover, the start-up phase was modified to avoid safety concerns. Indeed, we supposed that when combining two cytotoxic agents (and due to the potential synergistic effect in terms of toxicity between them), it is unreasonable to explore all levels of agent 2 during the start-up phase. Therefore, when toxicity was observed and the original rule was not possible, the start-up phase was stopped.

2.3. *Up-and-down design using the T-statistic* [18]

Using a newer approach proposed by Ivanova and Kim, a modification of the previous algorithm-based method can be implemented by replacing the Narayana design-based allocation rule with the T-statistic [18]. With the other parts of the method remaining the same, we defined the T-statistic at combination $\mathrm{D}_{j,k}$ by the following:

$$T_{j,k}=\frac{\varphi_{j,k}-\theta}{\frac{s_{j,k}}{\sqrt{n_{j,k}}}},\quad\text{where } s^2_{j,k}=\frac{t_{j,k}-2t_{j,k}\varphi_{j,k}+n_{j,k}\varphi^2_{j,k}}{n_{j,k}-1}$$

Then, according to the recommendation on parameter values in [18], the dose-allocation rule would be as follows: (i) $\mathrm{D}_{j+1,k}$ if $T_{j,k}\leq -1$, (ii) $\mathrm{D}_{j-1,k}$ if $T_{j,k}\geq 1$, (iii) $\mathrm{D}_{j+1,k-1}$ if $-1<T_{j,k}\leq 0$, and (iv) $\mathrm{D}_{j-1,k+1}$ if $0<T_{j,k}<1$.

2.4. *Two-dimensional dose finding in discrete dose space* [7]

Furthermore, Wang and Ivanova proposed a new two-dimensional model-based design of which the aim was to identify a set of MTDs for each dose of agent 2 [7]. As presented in the previous section, some

minor changes were proposed for this method in order to ensure comparability and respect clinical practice. The dose-combination model was defined as follows with an interaction term proposed by Wang and Ivanova:

$$\pi_{j,k} = 1 - (1 - p_j)^{\alpha} (1 - q_k)^{\beta} \exp\left(-\gamma \log(1 - p_j) \log(1 - g_k)\right)$$

where $\alpha > 0, \beta > 0, \gamma > 0$, and $p_j (j = 1, \ldots, J), q_k (k = 1, \ldots, K)$ are the working models for agents 1 and 2, respectively. The interaction term γ was introduced in the model to consider the possible synergistic effects.

After each cohort of patients, the estimation of $\hat{\pi}_{j,k}$ was updated using Monte Carlo method with exponential prior distributions centered in 1 for all parameters. At each step, the combination chosen to be allocated to the next cohort is the closest to the target belonging to $\left(\mathrm{D}_{j+1,k}, \mathrm{D}_{j,k+1}, \mathrm{D}_{j-1,k+1}, \mathrm{D}_{j-1,k}, \mathrm{D}_{j,k-1}, \mathrm{D}_{j+1,k-1}, \mathrm{D}_{j,k}\right)$.

For comparison purposes, the method was restricted to select only one MTD such that $\mathrm{D}_{j,k}$ was the combination with a toxicity probability closest to the target: $(j,k) = \mathrm{argmin}_{j,k} |P(Y = 1|\mathrm{D}_{j,k}) - \theta|$ and among the $\mathrm{D}_{j,k}$'s that have already been administered to patients, as proposed by Yin and Yuan [11, 12], without decreasing the performance of the method. Again, for the same reasons as previously outlined, the start-up phase was modified as detailed in Section 2.2.

2.5. *Continual reassessment method for partial ordering* [14]

Wages *et al*. proposed a dose-finding approach based on the CRM that considers orderings between combinations [14]. The ordering between agents is assumed to be monotonic and increases with the dose.

Assuming there are M possible ways to order combinations that are consistent with the non-decreasing assumption, let $w_{\ell} (\ell = 1, \ldots, J \times K)$ be the working model (that is the initial guesses of toxicity in ascending order), and $\alpha_{i,\ell}$ the initial guess w corresponding to the position of the combination received by patient i $(i = 1, \ldots, N)$ in the ordering ℓ.

For each ordering $m = 1, \ldots, M$, the dose–toxicity model is defined as a function of the dose and a parameter $a \in \mathbb{A}$: $\forall m, P(Y_i = 1|X_i = x_i) = \pi_m(x_i, a)$, where the 'empiric' model $\pi_m(x_i, a) = \alpha_{i,m}^{a}$ was chosen by the authors. Following O'Quigley *et al*. [2], a prior probability distribution $g(a)$ for a was assigned. Let $\{p(1), \ldots, p(M)\}$ denote the prior probability of each ordering representing their plausibility, where $\sum_m p(m) = 1$ and $\forall m, p(m) \geq 0$.

After the inclusion of I^{th} patient, for each ordering m, the posterior mean $\hat{a}_m$ and the posterior probabilities of m are estimated by

$$\hat{a}_m = \frac{\int_{\mathbb{A}} a . L_m(a|\text{data}) g(a) da}{\int_{\mathbb{A}} L_m(a|\text{data}) g(a) da} \quad \text{and} \quad \tilde{p}(m|\text{data}) = \frac{p(m) \int_{\mathbb{A}} L_m(a|\text{data}) g(a) da}{\sum_{m=1}^{M} \left[p(m) \int_{\mathbb{A}} L_m(a|\text{data}) g(a) da \right]}$$

where $L_m(a|\text{data})$ is a binomial likelihood.

The order h $(h = 1, \ldots, M)$ with the greatest posterior probability, $\tilde{p}(m|\text{data})$, is chosen for the next cohort. Nevertheless, when a uniform prior distribution is chosen for the ordering probabilities, as the trial proceeds, the difference between the posterior probabilities of m takes some time to differentiate. Therefore a start-up phase could be set up, where, for the first few patients, the ordering is sampled randomly with the weights of posterior probabilities. Once the order h has been chosen, the combination assigned to the next patient (or cohort) is the one closest to the target toxicity: $\mathrm{argmin}_{\ell} |\pi_h(w_{\ell}, \hat{a_h}) - \theta| = \mathrm{argmin}_{\ell} \left| w_{\ell}^{\hat{a}_h} - \theta \right|$ with $\ell = j \times k$. The original method recommended trial initiation at the combination assumed to be the MTD. Nevertheless, in our simulations, the first administered combination was the lowest for comparison purpose.

2.6. *Dose-finding design based on copula regression* [12]

Yin and Yuan proposed a Bayesian method using copula regression for combinations [12]. The authors have made the assumption that each single agent had already been evaluated in separate phase I trials.

As a consequence, physicians have a reasonable prior knowledge of the MTD of each drug alone. Let $p_1, \ldots p_J$ and $q_1, \ldots q_K$ be the prior toxicity probabilities of each dose level of agents 1 and 2 alone, respectively, and $\pi_{j,k}$ the probability of toxicity in combination $(\mathrm{D}_{j,k})$.

A Clayton-copula regression type that enables expressing the joint toxicity probability of combination $\mathrm{D}_{j,k}$ with marginal true probabilities of toxicity $\left(p_j^{\alpha} q_k^{\beta}\right)$ was used:

$$\pi_{j,k} = 1 - \left(\left(1 - p_j^{\alpha}\right)^{-\gamma} + \left(1 - q_k^{\beta}\right)^{-\gamma} - 1 \right)^{-\frac{1}{\gamma}}$$

where $\gamma, \alpha, \beta > 0$ are unknown parameters. The parameter γ characterizes the drug interaction effect, and α and β characterize the uncertainty on the initial guesses.

Let c_e and c_d, with $c_e + c_d > 1$, denote fixed probability cutoffs for dose escalation and de-escalation respectively that need to be calibrated through simulations studies. Prior distributions of model parameters are assumed to be independent with a prior distribution centered on 1 for α and β and a non-informative prior distribution for γ. Adaptive rejection Metropolis sampling within Gibbs sampling [19] was used to sample (α, β, γ) from the posterior joint distribution in order to calculate posterior estimates of $\pi_{j,k}$ and $P(\pi_{j,k} < \theta)$. In practice, the dose-allocation method is as follows: (i) if, at the current dose combination, $\mathrm{D}_{j,k}$, $P\left(\pi_{j,k} < \theta\right) > c_e$, then the combination is escalated to an adjacent combination $(\mathrm{D}_{j+1,k}, \mathrm{D}_{j,k+1}, \mathrm{D}_{j+1,k-1}, \mathrm{D}_{j-1,k+1})$ with the probability of toxicity higher than the current value and closest to θ; (ii) if, at the current dose combination, $\mathrm{D}_{j,k}$, $P\left(\pi_{j,k} > \theta\right) > c_d$ then the combination is de-escalated to an adjacent combination $(\mathrm{D}_{j-1,k}, \mathrm{D}_{j,k-1}, \mathrm{D}_{j+1,k-1}, \mathrm{D}_{j-1,k+1})$ with the probability of toxicity lower than the current value and closest to θ; and (iii) otherwise, the next cohort of patients continues to be treated at the current combination. Once the maximum sample size is reached, the combination associated with probability of toxicity that is closest to θ is selected as the MTD combination (from the dose tested on at least one cohort).

A start-up phase was proposed in order to gather enough information for estimating the $\pi_{j,k}$ where each agent's dose level is increased until at least one toxicity is observed while the other agent remains at its lowest level.

2.7. *Dose-finding design based on latent contingency table* [11]

The method proposed by Yin and Yuan in [11] is the same as that in [12] with another model for toxicity probability. A Gumbel model was chosen to model the probability of toxicity at combination $\mathrm{D}_{j,k}$, given by

$$\pi_{j,k} = 1 - \left(1 - p_j^{\alpha}\right)\left(1 - q_k^{\beta}\right)\left[1 + p_j^{\alpha} q_k^{\beta} \frac{e^{\gamma} - 1}{e^{\gamma} + 1}\right]$$

3. Simulations

We simulated 2000 independent replications of phase I trials evaluating two-agent combination trials in which five dose levels for agent 1 and three dose levels for agent 2 were chosen. Eight scenarios were studied (Table I) with several number and locations of the MTDs in the combination space. The chosen scenarios seemed to cover a wide variety of underlying realities. The toxicity target was fixed at 0.3, and the overall sample size was 60. To ensure comparability, the cohort size was chosen equal to 3 for all methods, and no stopping rules were used. Because of practical concerns, each trial started at the lowest combination $\mathrm{D}_{1,1}$.

For simplicity, the dose-finding methods are denoted in Section 3 as follows: (i) AISO for Ivanova and Wang [6], (ii) TSTAT for Ivanova and Kim [18], (iii) I2D for Wang and Ivanova [7], (iv) POCRM for Wages *et al*. [14], (v) BCOPULA for Yin and Yuan [12], and (vi) BGUMBEL for Yin and Yuan [11].

In order to be able to compare all dose-finding designs, modifications and assumptions were made (see Supporting information). All designs were optimized using the model average best-setting choices to improve the percentage of correct selections (PCS) when recommending one combination at the end of the trial. Indeed, we studied the influence of working models for each model-based design. Moreover, for each method with a start-up phase, we studied its influence. For example, for the POCRM, we studied the influence of the number of orderings retained for POCRM and the impact of those chosen on PCS. For I2D, we introduced the interaction term between the two agents suggested by Wang and Ivanova [7] and so on.

Table I. Toxicity scenarios for the two-agent combinations.

	Agent 1									
Agent 2	1	2	3	4	5	1	2	3	4	5
	Scenario 1					Scenario 2				
3	0.15	**0.30**	0.45	0.50	0.60	0.45	0.55	0.60	0.70	0.80
2	0.10	0.15	**0.30**	0.45	0.55	**0.30**	0.45	0.50	0.60	0.75
1	0.05	0.10	0.15	**0.30**	0.45	0.15	**0.30**	0.45	0.50	0.60
	Scenario 3					Scenario 4				
3	0.10	0.15	**0.30**	0.45	0.55	0.50	0.60	0.70	0.80	0.90
2	0.07	0.10	0.15	**0.30**	0.45	0.45	0.55	0.65	0.75	0.85
1	0.02	0.07	0.10	0.15	**0.30**	**0.30**	0.45	0.60	0.70	0.80
	Scenario 5					Scenario 6				
3	0.07	0.09	0.12	0.15	**0.30**	0.15	**0.30**	0.45	0.50	0.60
2	0.03	0.05	0.10	0.13	0.15	0.09	0.12	0.15	**0.30**	0.45
1	0.01	0.02	0.08	0.10	0.11	0.05	0.08	0.10	0.13	0.15
	Scenario 7					Scenario 8				
3	**0.30**	0.50	0.60	0.65	0.75	0.08	0.15	0.45	0.60	0.80
2	0.15	**0.30**	0.45	0.52	0.60	0.05	0.12	**0.30**	0.55	0.70
1	0.07	0.10	0.12	0.15	**0.30**	0.02	0.10	0.15	0.50	0.60
	Scenario 9					Scenario 10				
3	0.15	**0.30**	0.45	0.55	0.65	0.70	0.75	0.80	0.85	0.90
2	0.02	0.05	0.08	0.12	0.15	0.45	0.50	0.60	0.65	0.70
1	0.005	0.01	0.02	0.04	0.07	0.05	0.10	0.15	**0.30**	0.45

The MTD(s) combination are given in bold.

At the end of this optimization phase, in the simulation study, the marginal initial guesses of toxicities for agent 1, $p_{j=1}$ were chosen as (0.12, 0.2, 0.3, 0.4, 0.5), and for agent 2, q_k, as (0.2, 0.3, 0.4) for I2D, BCOPULA, and BGUMBEL using the 'getprior' function of the 'dfcrm' R package according to Lee and Cheung [20]. For BCOPULA and BGUMBEL, the dose-allocation thresholds were equal to $c_{e=1}=0.8, c_{d=1}=0.55$ and $c_{e=1}=0.7, c_{d=1}=0.55$, respectively, as proposed by Yin and Yuan [11, 12]. For POCRM, following Wages *et al.* [14], the number of possible orderings was restricted to 3 after a sensitivity analysis, and the working model was set up using the 'getprior' function with the length of indifference interval $\delta=0.03$ and the initial guessed MTD $\ell=13$ near the last combinations: (0.0001, 0.0006, 0.002, 0.005, 0.01, 0.02, 0.04, 0.06, 0.1, 0.14, 0.19, 0.24, 0.3, 0.36, 0.42). (Other working models were investigated; see Supporting information.)

At each simulated trial, we computed (i) the PCS at the end of the trial; (ii) the percentage of patients allocated at the true MTD(s) during the trial; (iii) the mean number of observed DLTs throughout the trial; and (iv) the mean number of patients allocated to each combination throughout the trial.

Designs were programmed using R version 2.13 [21] for AISO, TSTAT, I2D, and POCRM, and in C++ for BCOPULA and BGUMBEL.

3.1. *Dealing with multiple MTDs*

In this manuscript, we have proposed the recommendation of only one MTD at the end of the trial. In our case, we believe that the existence of one MTD for each row of agent 2 is not always true, but more than one MTD in the entire combination space is possible. Following this, we proposed some decision rules in order to identify at least one MTD at the end of the trial. We then evaluated its performance using the same scenarios as in the previous section. We first identified an MTD by level of agent 2; at the end of the trial, for $k=1, \ldots, K$, the MTD, $D_{j_k k}$ is the combination closest to the target, as follows: $j_k = \text{argmin}_j |P(Y=1|D_{j,k}) - \theta|$. Then we applied the following decision rules in order to identify MTDs that are too toxic or not toxic 'enough' by level of agent 2.

3.1.1. Decision rule for algorithm-based methods. The following decision rule was applied at the end of the trial: (i) if the combination selected to be the MTD $\in \{D_{1,k}, k=1, \ldots, K\}$ and $\pi^{\text{MTD}}_{1,k} - \theta > \tau_1$, then no combination was recommended on a row at the end of the trial; or (ii) if the combination selected

to be recommended $\in \{D_{J,k}, k = 1, \ldots, K\}$ and $\theta_\psi - \pi^{\text{MTD}}_{J,k\psi} > \tau_2$, then, once more, no combination was recommended.

3.1.2. Decision rule for model-based methods. The following decision rule was applied at the end of the trial: (i) if the combination selected to be recommended $\in \{D_{1,k\psi} k = 1, \ldots, K\}$ and $P_\psi\left(\pi^{\text{MTD}}_{1,k\psi} > \theta\right)_\psi > \tau_3$, then no combination was recommended on a row at the end of the trial; or (ii) if the combination selected to be recommended $\in \{D_{J,k}, k = 1, \ldots, K\}$ and $P_\psi\left(\pi^{\text{MTD}}_{J,k\psi} < \theta\right)_\psi > \tau_4$, then, once more, no combination was recommended. This rule was not applied to the POCRM as this method transforms a multidimensional combination space into an addition of several possible uni-dimensional orders.

In this simulation study, the thresholds were chosen as follows: $\tau_{1\psi} = \tau_{2\psi} = 0.15, \tau_{3\psi} = 0.90$, and $\tau_{4\psi} = 0.95$.

4. Results

4.1. Selection of one MTD at the end of the trial

Table II shows that the algorithm-based methods did not perform as well as the model-based ones. When comparing the performance of model-based methods, no design seemed to really stand out (Table II).

Scenarios 1 and 3 included three possible MTDs that were on one diagonal of the combination space; combinations $D_{2,3}$, $D_{3,2}$, and $D_{4,1\psi}$ in scenario 1 and combinations $D_{3,3}, \psi D_{4,2}, \psi$ and $D_{5,1\psi}$ in scenario 3. For these scenarios, all model-based designs gave a high PCS (over 66%), whereas POCRM seemed to

Table II. Comparison of all dose-finding designs in terms of percentage of correct selection, percentage of patients allocated at the true MTD(s) during the trial, and mean number of observed DLTs throughout the trial when the aim is to select only one MTD.

	Scenario									
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8	sc9	sc10
	Percentage of correct selections									
AISO	46.9	57.9	57.4	36.4	67.7	39.0	50.7	27.0	24.8	31.8
TSTAT	55.0	52.5	62.3	36.8	67.1	45.3	46.5	30.9	20.8	27.0
I2D	68.0	73.7	66.9	89.7	83.7	37.2	41.9	50.4	5.1	13.0
POCRM	72.7	64.4	72.5	73.8	81.8	49.1	47.7	55.1	3.4	8.2
BCOPULA	66.2	71.8	71.7	84.1	78.1	30.7	49.6	43.5	5.0	16.3
BGUMBEL	67.1	72.5	68.4	87.5	77.9	33.6	48.0	49.5	6.0	8.6
CRM anti-diag1	73.7	74.8	71.9	84.9	80.0	71.4	73.2	84.3	0.0	0.0
CRM anti-diag2	73.7	74.8	71.9	84.9	80.9	75.1	75.4	83.9	0.0	0.0
	Percentage of patient allocated at a true MTD(s) during the trial									
AISO	32.9	59.1	28.7	21.3	18.4	23.0	29.8	16.4	9.9	8.4
TSTAT	40.3	52.6	36.9	22.0	25.2	23.7	31.6	14.8	9.1	5.7
I2D	44.1	55.6	38.9	79.8	34.6	23.0	32.0	24.0	3.9	12.1
POCRM	46.8	39.6	51.6	57.4	66.1	28.8	34.1	28.5	3.1	8.8
BCOPULA	40.0	50.1	40.3	84.1	27.8	16.6	38.3	23.6	3.0	14.0
BGUMBEL	40.8	52.8	39.5	81.6	30.5	20.0	34.1	26.0	3.5	10.5
CRM anti-diag1	49.2	55.3	45.2	74.5	45.5	43.5	52.6	58.4	0.0	0.0
CRM anti-diag2	49.1	55.3	45.2	74.5	46.4	46.4	54.3	57.5	0.0	0.0
	Mean number of observed DLTs all over the trial									
AISO	13.8	19.5	12.0	26.4	8.4	12.7	15.3	14.0	12.2	22.3
TSTAT	16.1	20.6	13.9	26.4	9.1	15.2	18.2	16.1	14.4	23.3
I2D	15.3	17.6	14.1	19.9	10.1	14.3	16.1	15.3	14.4	17.0
POCRM	20.1	22.8	18.2	23.3	14.4	17.7	19.9	20.5	17.4	22.5
BCOPULA	14.2	16.1	12.7	19.5	9.2	12.8	14.6	14.3	12.4	15.5
BGUMBEL	14.6	16.6	13.2	19.7	9.5	13.5	15.6	14.7	13.1	16.5
CRM anti-diag1	16.5	18.4	15.3	20.3	11.5	14.8	17.6	17.1	0.0	0.0
CRM anti-diag2	16.4	18.4	15.3	20.3	11.5	14.9	17.9	16.9	-	-

perform better in terms of PCS. In scenario 2, in which two possible MTDs were located in the lower part of the combination space, the highest PCS values were observed for I2D, BCOPULA, and BGUMBEL (over 70%). When the correct combination was in the lower ($D_{1,1}$) or higher ($D_{5,3}$) extremity of the combination space, as with scenarios 4 and 5, I2D performed better. In scenarios 6 and 7, when the true MTDs were randomly located in the combination space, the performance in terms of combination selection was low, less than 40% for most designs in scenario 6 and less than 50% in scenario 7. When there was only one true MTD and it was located in the middle of the combination space, the PCS was less than 55%, whichever the design. Finally, in scenarios 9 and 10, where the true MTD was unique and at the border of the combination space, the algorithm-based methods performed better than the model-based methods. For scenario 9, the PCS was above 20% for algorithm-based methods (AISO and TSTAT) but was always below 6% for model-based methods. This could be due to the way in which the combination space was explored: AISO and TSTAT provided better adjacent combination exploration owing to their dose-allocation method. Most PCS values remained, however, relatively low.

Table II shows that POCRM generated more DLTs than the other methods. It also tended to overtreat more patients than the other methods, and at higher combinations. In fact, the gain in PCS for POCRM

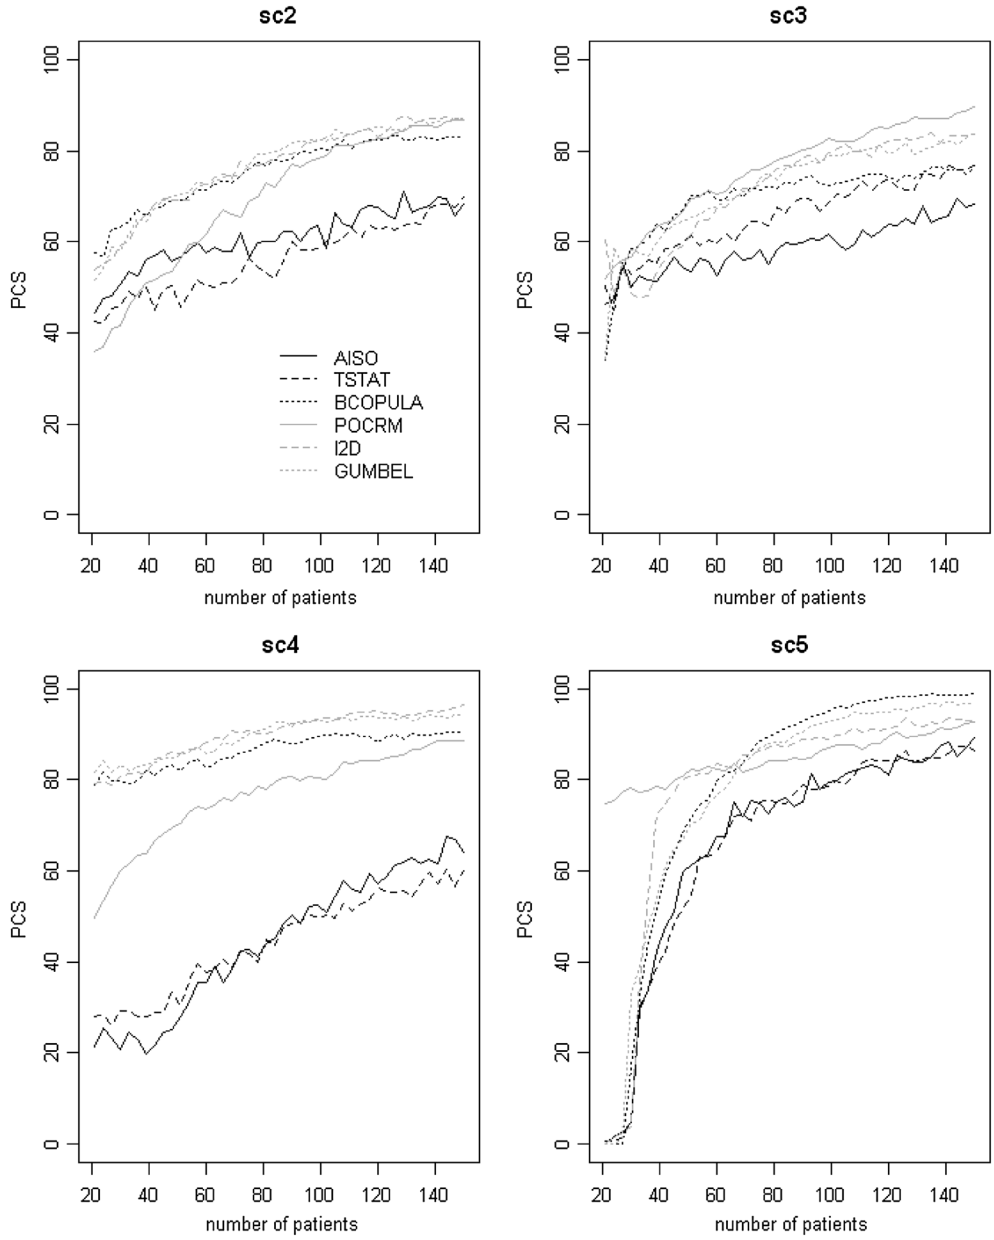

Figure 2. Convergence curves for scenarios 2–5.

when using a certain working model rather than others (see Supporting information) increased the mean number of DLTs. In the simulation study, we considered that this was acceptable, as the mean DLTs observed in all the scenarios was 19.7. This result is close to the expected number of DLTs (18), corresponding to a 0.3 toxicity combination in 60 patients. If investigators judged this possibility unacceptable, they would need to adjust the design parameterization prior to trial onset, or to consider a dose-allocation design with overdose control. The BCOPULA method gave similar PCS than BGUMBEL and had good properties in terms of mean DLTs number.

We then studied the convergence to the true MTDs while increasing the number of patients for all designs (Figure 2). We chose to show only four scenarios out of the eight presented in our simulation study. In scenario 2, all model-based designs were similar, but the algorithm-based design seemed to converge slowly. This finding was observed in all cases, where the convergence of AISO and TSTAT was slower than that of model-based methods. In scenario 3, the difference in PCS between the algorithm-based methods and the model-based methods tended to diminish with the increasing number of patients. In this scenario, POCRM, BGUMBEL, and I2D showed the best convergence, whereas in scenario 4, I2D and BGUMBEL approached nearly 95%. In scenario 5, BCOPULA and BGUMBEL had by far the best convergence and reached 90% very quickly. Nevertheless, in general, all methods (excepted BCOPULA and BGUMBEL in scenario 5) showed difficulties in attaining 100%, even with 300 patients. Overall, the convergence was rather slow.

4.2. Comparison with one-dimensional CRM

An important point is the contribution of multidimensional methods versus one-dimensional methods. As suggested during the review of this paper, we performed a one-dimensional CRM on a subset of combinations selected in an anti-diagonal of the dose-combination space where the toxicity probabilities order was known between combinations. We chose the following two different anti-diagonal paths chosen, as follows:

CRM anti-diag1: $D_{1,1\psi} \rightarrow D_{1,2\psi} \rightarrow D_{2,2\psi} \rightarrow D_{3,2\psi} \rightarrow D_{4,2\psi} \rightarrow D_{5,2\psi} \rightarrow D_{5,3\psi}$

CRM anti-diag2: $D_{1,1\psi} \rightarrow D_{2,1\psi} \rightarrow D_{2,2\psi} \rightarrow D_{3,2\psi} \rightarrow D_{4,2\psi} \rightarrow D_{4,3\psi} \rightarrow D_{5,3\psi}$

Using the 'dfcrm' package, we performed 2000 simulations on the scenarios corresponding to these anti-diagonals with restrictions to avoid skipping doses. The target toxicity, patient number, and cohort size were the same as for multidimensional methods. The working model was generated using the 'getprior' function with an indifference interval $\delta\psi = 0.05$, a initial guessed MTD at dose level 4 for a trial with seven doses. For scenarios 1–5, where at least one of the true MTD was included in anti-diagonals, PCSs were similar between multidimensional model-based methods and CRM. For scenarios 6–8, where the true MTDs were not located on the same diagonal, CRM on a reduced ordered subset of combinations containing at least one MTD had clearly higher performances than multidimensional designs. But in practice, the true MTD(s) is (are) not necessarily contained in the chosen anti-diagonal of the ordered

Table III. Comparison of AISO, TSTAT, I2D, BCOPULA, and BGUMBEL designs in terms of percentage of correct combination selection for each level of agent 2 when selecting multiple MTDs.

	Scenario 1			Scenario 2		Scenario 3			Scenario 4
	$D_{4,1\psi}$	$D_{3,2\psi}$	$D_{2,3\psi}$	$D_{2,1\psi}$	$D_{1,2\psi}$	$D_{5,1\psi}$	$D_{4,2\psi}$	$D_{3,3\psi}$	$D_{1,1\psi}$
AISO	49.0	55.5	49.3	49.6	71.2	44.4	54.1	49.5	51.8
TSTAT	46.1	54.9	61.5	46.4	64.8	44.6	56.8	64.4	57.5
I2D	55.6	71.3	63.5	76.9	84.6	75.4	64.6	64.5	90.4
BCOPULA	41.0	47.9	35.2	70.9	82.5	47.0	55.4	33.8	73.2
BGUMBEL	46.1	61.3	38.5	76.5	83.2	40.9	67.0	41.2	74.5

	Scenario 5	Scenario 6		Scenario 7			Scenario 8	Scenario 9	Scenario 10
	$D_{5,3\psi}$	$D_{4,2\psi}$	$D_{2,3\psi}$	$D_{5,1\psi}$	$D_{2,2\psi}$	$D_{1,3\psi}$	$D_{3,2\psi}$	$D_{2,3\psi}$	$D_{4,1\psi}$
AISO	67.3	47.6	42.2	34.4	47.8	41.8	76.3	37.7	46.9
TSTAT	69.5	49.9	48.4	33.2	40.8	62.5	75.4	40.3	39.8
I2D	77.9	60.9	32.4	18.1	42.6	53.9	89.3	11.8	23.4
BCOPULA	80.7	23.8	20.5	23.4	33.1	63.4	86.2	5.9	31.4
BGUMBEL	83.2	35.6	25.0	7.8	41.4	64.0	91.5	8.2	19.6

combinations retained, which was the case of both scenarios 9 and 10. For these scenarios, the CRM (anti-diad1 and anti-diag2) could obviously never select a true MTD as it was not contained in the chosen path. In this case, algorithm-based multidimensional methods performed better than model-based ones, even if PCS remained quite low.

4.3. Selection of multiple MTDs at the end of the trial

When selecting one MTD per level of agent 2 (Table III), PCSs of all methods were good on each row for scenarios 1–5 and 8 (higher than 40% in all cases and up to 91.5%). For scenario 6, the algorithm-based methods (AISO and TSTAT) and I2D performed well, whereas BCOPULA and BGUMBEL had rather low PCSs (between 20% and 35%). For scenario 7, $D_{2,2\psi}$ and $D_{1,3\psi}$ were well identified by all designs, but the PCSs for $D_{5,1\psi}$ were lower for model-based methods.

5. Discussion

The aim of this manuscript was to compare several dose-finding designs for cytotoxic combination studies. Based on this simulation study, model-based methods seemed to perform better than algorithm-based methods in terms of the percentage of correct combination selections (PCSs) when targeting a single MTD at the end of the trial. In general, the model-based methods gave a high PCSs in this case, and there was no major difference between the model-based methods compared. When one MTD per row was targeted, algorithm-based methods performed better than model-based methods but with low PCS.

For comparison purposes, several choices were made, which merit discussion. According to the combination dimensional space, we arbitrarily fixed the sample size at 60, as in Yin and Yuan [11, 12]. In this study, we chose five dose levels of agent 1 and three dose levels of agent 2, which resulted in 15 possible combinations to evaluate. Nevertheless, when using a different dimensional space ($J\psi\times K$), further investigations need to be carried out to find the optimal sample size for each method. In practice, it seems unreasonable to have such a large number of available combinations to evaluate, and only a subset of the dimensional space could be relevant. For this reason, we decided to compare the methods on a more realistic basis. Therefore, we chose 10 scenarios on a $5\psi\times 2\psi$ dimensional space and performed 2000 simulations of trials with 40 patients (data not shown). As in our manuscript, all model-based designs performed well.

Some authors have made the assumption that using one MTD for each level of agent 2 is possible when exploring a large number of combinations [6, 7]. We thus proposed decision rules designed to detect when at least one MTD existed in the combination space. These decision rules were implemented at the end of the trial and were found to maintain the performance of the designs. In this case, how should the most appropriate combination for further investigation be chosen? Phase II trials can study several combinations, and if they require the selection of a unique combination, other criteria such as efficacy or pharmacokinetics should be taken into account in the decision process. Indeed, when two cytotoxic agents are combined, the resulting pharmacokinetic (PK) profiles of two MTDs are not necessarily similar. In this case, the investigators could base their final decision on the maximization of exposure, or on the maximization of an efficacy surrogate.

Some issues are raised by the design modifications proposed in this paper. Some methods are designed to select only one MTD. For instance, BCOPULA and BGUMBEL have a conservative allocation algorithm that explores a restricted subset of the combination space and focuses on one combination when it is estimated to be the correct one. In these methods, patients are often allocated to one or few combinations, and the other combinations are allocated to very few or to no patients. As a result, the estimation in a row of agent 2 can be poor. Moreover, for decision rules, we decided to keep the same tau values ($\tau_{1\psi}=\tau_{2\psi}=0.15, \tau_{3\psi}=0.90$, and $\tau_{4\psi}=0.95$) for all of the designs. But some designs could have performed better if we had calibrated these values specifically.

The partial ordering method (POCRM) [14] is based on determining the most appropriate combination ordering in terms of toxicity, from a set of possible orderings. Nevertheless, the number of possible orderings increases with the combination space. In our simulation study, we restricted the choice to three reasonable orderings, as in Wages *et al*. [14] (see Supporting information). It should be noted that the method does not contain a 'non-skipping' rule and that in theory the combination allocated to the next cohort can 'skip' more than one combination (that is, selecting a combination, which is not in the immediate adjacent space of the current combination). Especially, if investigators necessarily

wish to begin the trial at the lowest combination, as we did in our simulation study, and no toxicity is observed, then the next combination will gravitate towards the initial guessed MTD. This will cause a huge skipping depending on the ordering selected and on the working model. It is clear that in practice this should not be allowed. As most real clinical trials begin at the lowest dose, we think that this method should be modified to allow initiating the trial at the lowest combination while adding 'non-skipping' rules or a start-up phase. For instance, as in 'classical' one-dimensional CRM, after selecting the ordering with the highest posterior probability, the method could restrict the dose allocation to combinations, which are next to the current one in the most probable ordering. Another possibility could be, if the current combination is $D_{j,k}$, to restrict the next combination when escalating to combinations $D_{j+1,k}$ and $D_{j,k+1}$ or to implement one of the starting phase proposed by Yin and Yuan [12] or by Wang and Ivanova [7].

During the review of this manuscript, Wages and Conaway [22] have published a paper proposing some guidelines for the POCRM. In their paper, they have suggested to place the initial guessed MTD at the middle of the working model to ensure that there is enough spacing both below and above this dose. In our simulation study, according to our sensitivity analyses, we have placed the initial guessed MTD near the third quartile of the dose range. As published by O'Quigley and Zohar [23], there is no sharp answer about what is the definition of a reasonable against a non-reasonable working model, although it may well be the notion of robustness itself. In the Supporting information, we have tried to point out how a non-reasonable or mis-specified choice can dramatically lessen the performance of the method (this choice was not robust for all scenarios; see Supporting information) [24]. This is why in this manuscript our choice was driven by this finding; thus we have selected a reasonable working model as it has shown to give good performances on average for all scenarios. Another important modification that we have added to the POCRM is the recommendation to start at the lowest dose. We have based our decision on common practice in phase I for a single agent or a combination of agents. This modification has shown to have equal performance than if the POCRM started at the initial guessed MTD (data not shown).

Another important issue relates to the performance of multidimensional methods versus one-dimensional methods. As suggested during the review of this paper, we performed a basic CRM on a subset of combinations selected in an anti-diagonal of the dose-combination space. When the MTD was included in the anti-diagonal, the one-dimensional basic CRM worked as well or better than any multidimensional method. This finding points out that if the MTD exists in the selected anti-diagonal, a one-dimensional method is preferable to a more complex one. In practice, the entire combination space is not often studied in combination trials; it can increase the number of combinations to be evaluated, and '$3+3$' dose-allocation rules, which are still used by investigators, are not valid for such trials. Using one-dimensional approaches involves a choice by the investigators in the determination of the combinations to study, and this can be a difficult question. The most important issue will then be whether the chosen subset of combinations contains an MTD. If it does, a one-dimensional method would perform better than multidimensional ones.

In our comparative simulation study, none of the model-based designs gave consistently better results than the others. Each method requires several choices prior to trial beginning, such as the choice of the working model, of the start-up phase, and of the prior distributions. According to our simulation results (see also Supporting information), it seemed that some choices can tumble the performance of a design. The issue of using a single MTD or multiple MTDs when evaluating a large combination space is challenging. Statisticians should propose combination methods that could identify the presence of one or more MTDs in the combination space, in their assumption, and in the dose-allocation process. These methods should also identify at which levels of agent 2 MTDs are located. Statisticians and investigators should be aware of the pros and cons of these designs in planning future trials. Our work was to enlighten multidimensional methods by comparing them using the same scenarios and the same (or very close) features.

Acknowledgements

We would like to thank Drs Ying Yuan and Raphael Porcher as well as reviewer 1 and the associate editor for their advice and constructive comments that have served to substantially improve the manuscript. This work was partially funded by grants from the ANRT (Association Nationale de la Recherche et de la Technologie), 'Laboratoires Servier' (CIFRE number 2011/0900), and Denis Diderot University-Paris 7 ('Aide à la mobilité internationale 2012').

References

1. Storer B. Design and analysis of phase I clinical trials. *Biometrics* 1989; **45**:925–937.
2. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase 1 clinical trials in cancer. *Biometrics* 1990; **46**:33–48.
3. Faries D. Practical modifications of the continual reassessment method for phase I cancer clinical trials. *Journal of Biopharmaceutical Statistics* 1994; **4**:147–164.
4. Korn E, Midthune D, Chen T, Rubinstein L, Christian M, Simon R. A comparison of two phase I trial designs. *Statistics in Medicine* 1994; **13**:1799–1806.
5. Chevret S. *Statistical Methods for Dose-Finding Experiments*, Statistics in Practice. John Wiley and Sons Ltd.: Chichester, 2006.
6. Ivanova A, Wang K. A non-parametric approach to the design and analysis of two-dimensional dose-finding trials. *Statistics in Medicine* 2004; **23**:1861–1870.
7. Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics* 2005; **61**:217–222.
8. Mandrekar S, Cui Y, Sargent D. An adaptive phase I design for identifying a biologically optimal dose for dual agent drug combinations. *Statistics in Medicine* May 2007; **26**:2317–2330.
9. Mandrekar S, Qin R, Sargent D. Model-based phase I designs incorporating toxicity and efficacy for single and dual agent drug combinations: methods and challenges. *Statistics in Medicine* May 2010; **29**:1077–1083.
10. Fan S, Venook A, Lu Y. Design issues in dose-finding phase I trials for combinations of two agents. *Journal of Biopharmaceutical Statistics* 2009; **19**:509–523.
11. Yin G., Yuan Y. A latent contingency table approach to dose finding for combinations of two agents. *Biometrics* 2009; **65**:866–875.
12. Yin G, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society* 2009; **58**:211–224.
13. Bailey S, Neuenschwander B, Laird G, Branson M. A Bayesian case study in oncology phase I combination dose-finding using logistic regression with covariates. *Journal of Biopharmaceutical Statistics* 2009; **19**:469–484.
14. Wages N, Conaway M, O'Quigley J. Continual reassessment method for partial ordering. *Biometrics* 2011; **67**:1555–1563.
15. Narayana T. Sequential procedures in the probit analysis. *Dissertation*, The University of North Carolina, Chapel Hill, 1953.
16. Ivanova A, Montazer-Haghigh A, Mohanty S, Durham S. Improved up-and-down designs for phase I trials. *Statistics in Medicine* 2003; **22**:69–82.
17. Dykstra R, Robertson T. An algorithm for isotonic regression for two or more independent variables. *Annals of Statistics* 1982; **10**:708–716.
18. Ivanova A, Kim S. Dose finding for continuous and ordinal outcomes with a monotone objective function: a unified approach. *Biometrics* Mar 2009; **65**(1):307–315.
19. Gilks W, Best N, Tan K. Adaptive rejection metropolis sampling within Gibbs sampling. *Journal of the Royal Statistical Society* 1995; **44**:455–472.
20. Lee S, Cheung Y. Model calibration in the continual reassessment method. *Clinical Trials* 2009; **6**:227–238.
21. R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing: Vienna, Austria, 2011. http://www.R-project.org/, ISBN 3-900051-07-0.
22. Wages N, Conaway M. Specifications of a continual reassessment method design for phase I trials of combined drugs. *Pharmaceutical Statistics* 2013; **12**(4):217–224.
23. O'Quigley J, Zohar S. Retrospective robustness of the continual reassessment method. *Journal of Biopharmaceutical Statistics* Sep 2010; **20**(5):1013–1025.
24. Neuenschwander B, Branson M, Gsponer T. Critical aspects of the Bayesian approach to phase I cancer trials. *Statistics in Medicine* Jun 2008; **27**(13):2420–2439.

Supporting information

Additional supporting information may be found in the online version of this article at the publisher's web site.