

Bayesian Dose Finding for Combined Drugs with Discrete and Continuous Doses

Lin Huo*, Ying Yuan[†] and Guosheng Yin[‡]

Abstract. The trend of treating patients with combined drugs has grown in cancer clinical trials. Often, evaluating the synergism of multiple drugs is the primary motivation for such drug-combination studies. To enhance patient response, a new agent is often investigated together with an existing standard of care (SOC) agent. Often, a certain amount of dosage of the SOC is administered in order to maintain at least some therapeutic effects in patients. For clinical trials involving a continuous-dose SOC and a discrete-dose agent, we propose a two-stage Bayesian adaptive dose-finding design. The first stage takes a continual reassessment method to locate the appropriate dose for the discrete-dose agent while fixing the continuous-dose SOC at the minimal therapeutic dose. In the second stage, we make a fine dose adjustment by calibrating the continuous dose to achieve the target toxicity rate as closely as possible. Dose escalation or de-escalation is based on the posterior estimates of the joint toxicity probabilities of combined doses. As the toxicity data accumulate during the trial, we adaptively assign each cohort of patients to the most appropriate dose combination. We conduct extensive simulation studies to examine the operating characteristics of the proposed two-stage design and demonstrate the design's good performance with practical scenarios.

Keywords: Bayesian adaptive design, Combined drugs, Continual reassessment method, Maximum tolerated dose, Phase I trial, Toxicity probability, Two-stage design

1 Introduction

The typical goal of a phase I clinical trial in oncology is to identify the maximum tolerated dose (MTD), which is defined as the dose with a toxicity probability closest to the physician-specified target toxicity rate. Since little is known about the MTD at such an early stage of a drug study, most phase I clinical trials start with a low dose level that is presumed safe. From that point, the toxicity probability at each dose level can be continuously estimated based on the accumulating data, and patients are adaptively assigned to the appropriate dose levels.

Numerous methods have been proposed for designing phase I clinical trials, most of which have an underlying assumption that the dose-toxicity curve is monotonically increasing. For example, the algorithm-based 3+3 design is a conventional method

*Novartis Pharmaceuticals Corporation, East Hanover, NJ

[†]Department of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, TX, yyuan@mdanderson.org

[‡]Department of Statistics and Actuarial Science, The University of Hong Kong, Pokfulam, Hong Kong, gyin@hku.hk

that often locates the MTD as the dose with a toxicity probability less than 33%; see [Storer \(1989\)](#). The 3+3 design is commonly used in practice due to its simplicity; however, the estimates of the toxicity probabilities may not be stable. As an alternative, the continual reassessment method (CRM) developed by [O'Quigley, Pepe and Fisher \(1995\)](#) assumes a parametric link function between the true toxicity probabilities and the physician-specified toxicity probabilities. This model-based dose-finding method efficiently estimates the single unknown parameter in the link function, updates the posterior estimates of the toxicity probabilities, and eventually locates the MTD from the doses under consideration. Criticism of the original CRM for a tendency to allocate many patients to overly toxic doses has led to the development of several modifications to improve its performance. [Goodman, Zahurak and Piantadosi \(1995\)](#) developed a modified version of the CRM that allocates patients with a cohort size larger than one to each dose level and constrains the dose escalation by one level at a time. [Møller \(1995\)](#) combined the CRM with a preliminary up-and-down design to reach the neighborhood of the target dose while also limiting the dose escalation to a single dose level. [Piantadosi, Fisher and Grossman \(1998\)](#) implemented the CRM with a simple dose-toxicity model to guide data interpolation. [Heyd and Carlin \(1999\)](#) formulated a stopping rule to terminate the trial if the width of the 95% posterior probability interval for the MTD is within certain limits. [Ishizuka and Ohashi \(2001\)](#) extended the CRM by monitoring a posterior density function of the occurrence of the dose-limiting toxicity (DLT) at each dose level. [Leung and Wang \(2002\)](#) used the theory of decision processes to find optimal allocations that maximize the expected number of patients assigned to the MTD. [Yuan, Chappell and Bailey \(2007\)](#) developed a quasi-likelihood approach to accommodating multiple toxicity grades. To improve the robustness of the CRM, [Yin and Yuan \(2009a\)](#) proposed to use multiple prespecified toxicity probabilities in the CRM model and take a Bayesian model averaging (BMA) approach to estimating the true toxicity probabilities. Comprehensive reviews and discussions of statistical methods for phase I trial designs can be found in [Chevret \(2006\)](#) and [Yin \(2012\)](#).

All of the aforementioned methods are developed for single-agent dose-finding trials, and thus cannot address the issues in drug-combination studies. Combining different agents may induce synergistic treatment effects, target different disease pathways, and achieve high dose intensities with non-overlapping toxicities. However, complex drug-drug interactions may also lead to unknown toxicity patterns and thus make it difficult to fully rank the toxicities of all dose combinations. This is particularly true for the dose pairs along the off-diagonal directions in the two-dimensional dose-finding matrix. In the area of drug-combination studies, [Korn and Simon \(1993\)](#) introduced a tolerable dose diagram to provide guidance in targeting specific MTD combinations. [Kramar, Lebecq and Candalh \(1999\)](#) proposed searching over a selected subset of drug combinations that still maintains the monotonic toxicity order. For a trial combining paclitaxel and carboplatin to treat ovarian cancer, [Kuzuya et al. \(2001\)](#) proposed fixing one agent at each dose level and varying the dose level of the other. [Thall et al. \(2003\)](#) proposed a six-parameter model for the joint toxicity probability of the combination of gemcitabine and cyclophosphamide. The two drugs are modeled as two continuous-dose agents because their dosages can be given at any arbitrary amount through intravenous administration. [Wang and Ivanova \(2005\)](#) studied a log-linear model by using the standardized doses of

two drugs as the covariates. Yuan and Yin (2008) developed a sequential design scheme for drug-combination trials, which uses the partial order of the two-dimensional dose-finding space. Through incorporating the prior information on the toxicity probabilities when each drug is administered alone, Yin and Yuan (2009b) proposed a copula-type regression to model the joint toxicity probabilities of combined drugs. Braun and Wang (2010) developed a Bayesian hierarchical model for the probabilities of the DLTs for combined doses of two therapeutic agents. Wages, Conaway and O'Quigley (2011) investigated the CRM with partial ordering in two-dimensional dose-finding trials. Most of the drug-combination dose-finding methods focus on the cases for which the dosages of both drugs in the combination are discrete and are often prepared in tablets, except for the work of Thall et al. (2003) which investigates the combination of two intravenously administered drugs with continuous doses.

Our research is motivated by a clinical trial in metastatic renal cell carcinoma (RCC), which investigated the combination of the standard of care, temsirolimus, with a new agent for RCC patients. The new agent was administered orally at five dose levels and temsirolimus was administered by intravenous infusion. By targeting different disease pathways, the combined drugs are expected to induce a synergistic treatment effect. The primary objective of this trial is to find the MTD combination of the new therapy and temsirolimus with the target toxicity probability of 0.3. There may be multiple MTD combinations because of the dose-toxicity contour in the two-dimensional dose combination space. In this design, we focus on finding one single MTD combination, for which the toxicity probability of the identified discrete dose in combination with the minimal therapeutic continuous dose must be below but closest to the toxicity target. This allows for a further fine adjustment through the continuous dose in order to reach the target toxicity probability as closely as possible. Toward this goal, we propose a robust two-stage CRM design for clinical trials combining a discrete-dose agent and a continuous-dose agent. The first stage adopts the conventional CRM to locate the most appropriate dose for the discrete-dose agent, while fixing the continuous-dose agent at the minimum baseline dose specified by physicians. The second stage searches for the best dose combination of the continuous and discrete agents by continuously updating the posterior estimates of the toxicity probabilities for the combined doses. If the discrete dose is overestimated or underestimated in stage I, it can still be adjusted in stage II. Based on the accumulating data, we adaptively assign each cohort of patients to the most appropriate dose combination throughout the trial.

The rest of the paper is organized as follows. In Section 2, we propose the two-stage CRM design to jointly locate the maximum tolerated dose combination, and we also present the likelihood function and the posterior distribution for the unknown parameters. In Section 3, we describe the dose-finding algorithm in detail. We conduct extensive simulation studies to examine the operating characteristics of the proposed design in Section 4. Finally, we conclude with a brief discussion in Section 5.

2 Two-stage Dose Finding

2.1 Continual Reassessment Method

Suppose that there are J predetermined doses for the discrete-dose agent under study, and let (p_1, \dots, p_J) represent a set of prespecified toxicity probabilities (known as the skeleton in the CRM) at those doses. We assume that toxicity monotonically increases with respect to dose levels, i.e., $p_1 < \dots < p_J$, and denote ϕ as the target toxicity rate. A commonly used one-parameter CRM model is given as

$$\Pr(\text{toxicity at dose level } d) = \pi_d(\alpha) = p_d^{\exp(\alpha)} \quad (1)$$

for $d = 1, \dots, J$. In this model, a power function with an unknown parameter α links the true toxicity probabilities with the prespecified toxicity probabilities.

Suppose that n patients have entered the trial, and let d_i and y_i denote the received dose level and the toxicity outcome for the i th subject, respectively. The toxicity outcome is taken as a Bernoulli variable, i.e., $y_i = 1$ with probability $\pi_{d_i}(\alpha)$, and 0 with probability $1 - \pi_{d_i}(\alpha)$. The likelihood function based on the observed toxicity outcomes $\mathbf{y} = \{y_i, i = 1, \dots, n\}$ is then given by

$$L(\alpha|\mathbf{y}) \propto \prod_{i=1}^n \{p_{d_i}^{\exp(\alpha)}\}^{y_i} \{1 - p_{d_i}^{\exp(\alpha)}\}^{1-y_i}.$$

In the Bayesian paradigm, let $f(\alpha)$ denote the prior distribution for the parameter α , and we can estimate the toxicity probabilities by their posterior means,

$$\hat{\pi}_d = \int p_d^{\exp(\alpha)} \frac{L(\alpha|\mathbf{y})f(\alpha)}{\int L(\alpha|\mathbf{y})f(\alpha)d\alpha} d\alpha, \quad d = 1, \dots, J.$$

Based on these toxicity probability estimates, we assign the next cohort of patients to dose level $d^{(1)}$, which has the toxicity probability closest to the target rate. Let d^{curr} denote the current dose level, then

$$d^{(1)} = \operatorname{argmin}_{d \leq d^{\text{curr}}+1} |\hat{\pi}_d - \phi|, \quad (2)$$

which restricts dose escalation by one dose level only. The trial continues till exhausting the total sample size, and then the dose with a posterior toxicity probability closest to ϕ is selected as the MTD.

2.2 Two-stage Continual Reassessment Method

For a set of prespecified discrete doses and an unspecified continuous dose of two drugs in combination, we propose a two-stage CRM design to locate the MTD combination. For convenience, we denote the discrete-dose agent as A with prespecified doses, $A_1 < \dots < A_J$, and the continuous-dose agent as B , whose dose is unspecified but the initial minimal dose is given. In stage I, fixing B at the prespecified minimal dose z_0 , we

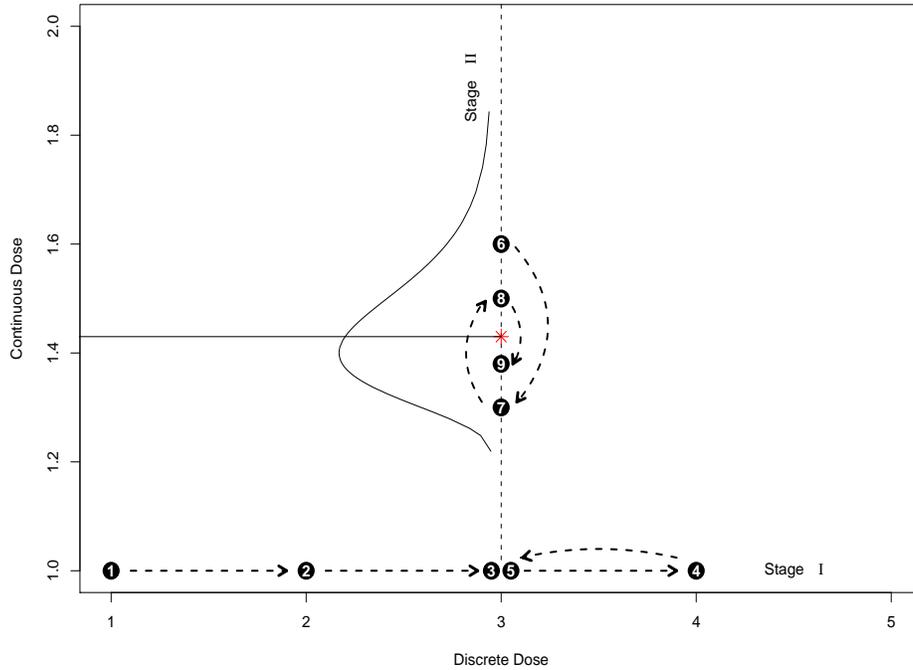


Figure 1: Illustration of the two-stage CRM. Each cohort is represented by a circle and numbered by the cohort’s order in the trial at the corresponding dose combination. The star symbol represents the target dose combination with a toxicity probability of 0.3. Stage I takes five cohorts to select the target discrete dose 3; and stage II uses the remaining four cohorts to estimate the continuous dose in combination with the discrete dose 3.

adopt the traditional CRM model to find the appropriate dose level of A to be used in combination with B such that the joint toxicity probability is closest to but not higher than the prespecified target toxicity rate ϕ . In stage II, we employ a two-parameter CRM model to calibrate the dose of B to achieve the target toxicity rate ϕ as closely as possible. The intuition behind our approach is that we first take a crude search over discrete doses of A to reach the neighborhood but lower than the target, and then make a fine adjustment by calibrating the continuous dose of B . This two-stage dose-finding procedure is illustrated in Figure 1: The first five cohorts of patients are used in stage I to locate the target discrete dose and the remaining four cohorts are treated in stage II to determine the continuous dose. Even in stage II, the discrete dose can still be adjusted if it is under- or overestimated in stage I.

Let z ($z \geq z_0$) denote the continuous dose in the drug combination. We propose a

two-parameter model in the two-stage CRM design,

$$\Pr(\text{toxicity at dose level } d \text{ of } A \text{ and dose } z \text{ of } B) = \pi_d(z) = p_d^{\exp(\alpha - \beta(z - z_0))} \quad (3)$$

for $d = 1, \dots, J$, where α and β are unknown parameters. We require $\beta > 0$, so that the monotonic toxicity assumption still holds for the continuous-dose agent when fixing the discrete dose. Via a simple transformation, model (3) can be rewritten as a generalized linear model of the form

$$\log[-\log\{\pi_d(z)\}] = \log\{-\log(p_d)\} + \alpha - \beta(z - z_0).$$

Therefore, natural choices for the prior distributions of α and β are $\alpha \sim N(0, \sigma^2)$ and $\beta \sim Ga(a, b)$, where $N(0, \sigma^2)$ denotes a normal distribution with mean 0 and variance σ^2 and $Ga(a, b)$ denotes a gamma distribution with shape and rate parameters a and b . As the effects of α and β on the toxicity probability $\pi_d(z)$ are on the complementary log-log scale, setting the values of these hyperparameters such as $\sigma^2 = 2$ and $a = b = 0.1$ would lead to adequately vague priors, such that the data will dominate the posterior distribution.

Our design consists of two sequential stages with a prespecified sample size N_1 for stage I and N_2 for stage II. In stage I, we fix the continuous-dose agent at the baseline dose z_0 and only consider dose escalation and de-escalation for the discrete dose. By taking $z = z_0$, the two-parameter CRM model (3) reduces to the usual CRM model (1), and thus the conventional CRM as described in Section 2.1 can be directly used for dose finding in stage I. After N_1 patients are used up, we identify the discrete dose together with z_0 that should have an estimated toxicity probability closest to but not greater than the target ϕ . It is worth emphasizing that the selected discrete dose together with the minimum continuous dose z_0 should have a toxicity probability smaller than ϕ , so that we will still have some room to search for the desirable continuous dose in stage II.

Assume that, at a certain time point of stage II, an additional n_2 patients have been enrolled into the trial, and let z_i denote the continuous dose assigned to the i th patient, $i = 1, \dots, N_1 + n_2$. Certainly, $z_i = z_0$ for all the first N_1 patients. Under the two-parameter CRM model, the likelihood function is given by

$$L(\alpha, \beta | \mathbf{y}) \propto \prod_{i=1}^{N_1+n_2} \{p_{d_i}^{\exp(\alpha - \beta(z_i - z_0))}\}^{y_i} \{1 - p_{d_i}^{\exp(\alpha - \beta(z_i - z_0))}\}^{1-y_i}.$$

The unknown parameters α and β can be estimated by their posterior means:

$$\hat{\alpha} = \int \int \alpha \frac{L(\alpha, \beta | \mathbf{y}) f(\alpha) f(\beta)}{\int \int L(\alpha, \beta | \mathbf{y}) f(\alpha) f(\beta) d\alpha d\beta} d\alpha d\beta,$$

$$\hat{\beta} = \int \int \beta \frac{L(\alpha, \beta | \mathbf{y}) f(\alpha) f(\beta)}{\int \int L(\alpha, \beta | \mathbf{y}) f(\alpha) f(\beta) d\alpha d\beta} d\alpha d\beta,$$

where $f(\alpha)$ and $f(\beta)$ are the prior distributions for α and β , respectively. Based on these posterior estimates, we choose the discrete dose $d^{(II)}$ as the highest dose with a

toxicity probability less than ϕ ,

$$d^{(II)} = \max\{d : p_d^{\exp(\hat{\alpha})} \leq \phi \text{ and } d \leq d^{\text{curr}} + 1\}, \quad (4)$$

in which we impose the condition $d \leq d^{\text{curr}} + 1$ to enhance patient safety by restricting to one-level dose escalation at most. We then determine the continuous dose z by solving equation (3), that is,

$$\phi = p_{d^{(II)}}^{\exp(\hat{\alpha} - \hat{\beta}(z - z_0))},$$

which yields

$$z = z_0 + \frac{\hat{\alpha} - \log\{\log(\phi)/\log(p_{d^{(II)}})\}}{\hat{\beta}}. \quad (5)$$

Thus, the identified dose combination $(d^{(II)}, z)$ has an estimated toxicity probability ϕ and will be used to treat the next cohort of patients. One of the appealing features of the two-stage procedure is that we still update the discrete dose even in the second stage by pooling all the information together, in order to achieve the goal of jointly searching for the discrete and continuous doses.

2.3 Robust Two-stage CRM with Model Selection

Like the standard CRM, the two-stage CRM design also requires prespecification of a set of arbitrary toxicity probabilities (skeleton). In an early stage trial, our knowledge on the toxicity profile of a new drug, and especially a combination of new drugs, is very limited. Hence, the prespecified skeleton could be quite subjective. If the skeleton is misspecified, the performance of the CRM can be substantially compromised. To address this issue, Yin and Yuan (2009a) proposed to use multiple skeletons in the CRM and take a Bayesian model averaging (BMA) approach to obtaining robust estimates of the toxicity probabilities. Following this route, we propose prespecifying multiple skeletons and then use the model selection procedure to overcome the possible misspecification of the skeleton. Compared with the BMA approach, model selection is computationally straightforward and easier to implement. In particular, we can elicit two or more skeletons for the discrete doses, and each skeleton represents a prior guess of the toxicity profile of the discrete-dose agent. We view the CRM model under each set of the prespecified toxicity probabilities as a separate model, and use the Bayesian information criterion (BIC) to select the best-fitting model to make inference on the unknown parameters and determine the dose assignment.

More specifically, let $\{(p_{11}, \dots, p_{1J}), \dots, (p_{K1}, \dots, p_{KJ})\}$ denote K sets of prior guesses of the toxicity probabilities (i.e., skeletons) for the discrete-dose agent. Each of the skeletons leads to a dose-toxicity model of the form

$$\pi_{kd}(z) = p_{kd}^{\exp(\alpha - \beta(z - z_0))}, \quad d = 1, \dots, J,$$

resulting in a total of K candidate models $\{M_1, \dots, M_K\}$, where M_k denotes the model based on the k th skeleton (p_{k1}, \dots, p_{kJ}) . Note that the K candidate models share the same structure and the only difference is that they use different values for the skeleton.

During the trial, after each new cohort of patients enters the trial, in light of the most recent observations, we fit each of the K candidate models independently and calculate the corresponding BIC,

$$\text{BIC}_k = -2\log L(\alpha, \beta | \mathbf{y}, M_k) + r \log(n),$$

where r is the number of model parameters and n is the total number of patients enrolled in the trial thus far. We select the best-fitting model that yields the lowest value of the BIC, and under the selected model, we can estimate the toxicity probability of each dose and guide the next dose assignment. Consequently, we are able to locate the target dose combination in a more reliable way and to treat the patients at more appropriate doses.

3 Dose-finding Algorithm

Let ϕ be the physician-determined toxicity target, and let z_0 be the prespecified minimal baseline continuous dose. Our two-stage dose-finding algorithm is described as follows.

Stage I:

1. The first cohort of patients receives the lowest dose combination (A_1, z_0) .
2. Based on the accumulated data, we first select the best-fitting model and then obtain the estimates of the dose toxicity probabilities under the selected model. The next cohort of patients will be treated at the combination of $(d^{(I)}, z_0)$, where $d^{(I)}$ is given by (2).
3. Once the maximum sample size N_1 is reached, we complete stage I and select the dose that has the toxicity probability closest to but less than ϕ as the starting dose of the discrete-dose agent for stage II.

Stage II:

1. Based on the data collected in both stages I and II, we conduct model selection and obtain the posterior estimates of the parameters α and β under the best-fitting two-parameter CRM model. We treat the next cohort of patients at the dose combination $(d^{(II)}, z)$, where $d^{(II)}$ is given by (4) and the continuous dose z is given by (5).
2. Once the total sample size of the trial, $N_1 + N_2$, is reached, we select the updated dose combination $(d^{(II)}, z)$ as the MTD based on all the observations.

For safety, we impose a stopping rule in the dose-finding algorithm: If the toxicity probability at the lowest dose combination is still overly toxic as noted by

$$\Pr(\text{toxicity rate at } A_1 \text{ and } z_0 \text{ is greater than } \phi \mid \text{data}) > 0.9,$$

we terminate the trial.

4 Numerical Studies

4.1 Simulation Studies

We investigated the performance of the proposed robust two-stage CRM design through simulation studies under eight different scenarios. We took a cohort size of 2, and allocated 8 cohorts in stage I and 10 cohorts in stage II, that is, $N_1 = 16$ and $N_2 = 20$. We assumed that toxicity monotonically increased with doses and the target toxicity probability was $\phi = 0.3$. For the continuous-dose agent, we standardized the baseline dose $z_0 = 1$; and for the discrete-dose agent, we considered five dose levels. In Table 1, we list the true toxicity probabilities for the discrete doses and the target continuous dose in the first row under each scenario. The target discrete dose is highlighted with an underline, which is the one having a toxicity probability closest to but not higher than the target toxicity probability.

To implement the proposed robust two-stage CRM, we elicited three sets of prior guesses of the toxicity probabilities:

$$(p_1, p_2, p_3, p_4, p_5) = \begin{cases} (0.05, 0.14, 0.18, 0.22, 0.30), & \text{skeleton 1,} \\ (0.08, 0.12, 0.30, 0.40, 0.50), & \text{skeleton 2,} \\ (0.20, 0.30, 0.40, 0.50, 0.60), & \text{skeleton 3.} \end{cases}$$

The three skeletons represent quite different toxicity profiles. The first skeleton is concentrated on the low toxicity levels with a toxicity probability of 0.3 at the highest dose. The second skeleton has the target in the middle and represents the case in which toxicity increases slowly at the low doses but quickly at the high doses. The toxicity probabilities in the third skeleton are evenly distributed over the range of 0.2 to 0.6. We compare the performance of the two-stage CRM using each skeleton alone and the robust version of using model selection with all three skeletons. We refer to the individual two-stage CRMs using each of these three skeletons as CRM_{2S} 1, CRM_{2S} 2, and CRM_{2S} 3, and let CRM_{2S}-MS denote the proposed robust two-stage CRM with model selection. We simulated 1,000 trials under each scenario.

Table 1 shows the simulation results, including the selection probabilities of discrete doses, and the 25th, 50th and 75th percentiles of the identified continuous doses. We also report the average number of patients treated at each discrete dose, the average number of patients who experienced toxicity, the average number of patients treated in the entire trial, and the percentage of inconclusive trials (i.e., the trials that were early stopped due to excessive toxicity, denoted by the last column “None”). In scenario 1, the first dose is the target discrete dose; the three individual two-stage CRMs using different skeletons selected the target discrete dose with quite different percentages, indicating that the two-stage CRM is sensitive to the specification of the skeleton. In particular, CRM_{2S} 2 (corresponding to the design using skeleton 2) had the lowest selection percentage of 48.7% for the target discrete dose, and the smallest estimate of the continuous dose, 1.12, compared with the target continuous dose of 1.46. By contrast, the proposed CRM_{2S}-MS selected the target discrete dose with a percentage of 63.7% and estimated the continuous dose as 1.39. Clearly in this scenario, the CRM_{2S}-MS performed much better than CRM_{2S} 2 and CRM_{2S} 3. The number of patients treated at each dose

Table 1: Simulation study of the two-stage CRM with a toxicity target of $\phi = 30\%$. We present the selection percentage for each discrete dose (the target dose in boldface), the average number of patients treated at each dose, and the median of the estimated continuous dose along with the 25th and 75th percentiles. The last three columns correspond to the average number of observed toxicities, the average number of patients in the trial, and the percentage of inconclusive trials.

Two-stage Design	Agent A (Discrete)					Agent B (Continuous)	Ave. # tox	Ave. # pts	None
	1	2	3	4	5				
Scenario 1	0.22	0.35	0.45	0.60	0.70	1.46			
CRM _{2S} 1	76.2	8.5	3.4	0.3	0	1.41 _(1.19,1.77)	13.3	32.5	11.6
# patients	21.5	6.5	2.8	1.4	0.3				
CRM _{2S} 2	48.7	37.4	2.0	0.2	0	1.12 _(1.06,1.20)	14.8	32.5	11.7
# patients	16.2	12.9	2.7	0.6	0.1				
CRM _{2S} 3	61.8	24.7	2.6	0.2	0	1.26 _(1.10,1.65)	13.9	32.8	10.7
# patients	19.0	9.6	3.3	0.7	0.2				
CRM _{2S} -MS	63.7	23.1	2.1	0.2	0	1.39 _(1.19,1.73)	14.2	32.7	10.9
# patients	19.0	9.8	3.1	0.7	0.1				
Scenario 2	0.10	0.25	0.45	0.55	0.65	1.28			
CRM _{2S} 1	56.3	27.1	13.4	2.2	0.1	1.08 _(1.04,1.21)	11.3	35.7	0.9
# patients	14.6	11.0	6.4	2.9	0.8				
CRM _{2S} 2	11.1	82.8	4.1	0.2	0	1.29 _(1.12,1.61)	11.8	35.4	1.8
# patients	6.8	21.5	5.6	1.2	0.3				
CRM _{2S} 3	24.4	62.9	10.5	0.8	0	1.23 _(1.09,1.54)	11.1	35.5	1.4
# patients	9.7	17.1	6.7	1.5	0.5				
CRM _{2S} -MS	27.2	65.1	5.7	0.8	0	1.30 _(1.14,1.65)	11.8	35.6	1.2
# patients	9.2	17.9	6.5	1.6	0.4				
Scenario 3	0.14	0.21	0.35	0.49	0.63	1.52			
CRM _{2S} 1	44.9	22.2	22.2	8.4	0.1	1.10 _(1.05,1.22)	11.2	35.3	2.2
# patients	12.6	9.4	7.4	4.7	1.2				
CRM _{2S} 2	9.9	73.7	12.2	1.1	0.1	1.46 _(1.23,1.87)	11.5	35.0	3.0
# patients	6.7	18.4	7.5	2.0	0.5				
CRM _{2S} 3	17.8	48.3	29.1	1.9	0	1.30 _(1.12,1.63)	10.9	35.1	2.8
# patients	8.5	14.0	9.5	2.5	0.7				
CRM _{2S} -MS	19.3	63.3	12.5	2.5	0	1.47 _(1.27,1.86)	11.5	35.2	2.4
# patients	8.1	16.3	7.6	2.6	0.6				
Scenario 4	0.08	0.15	0.22	0.35	0.50	1.46			
CRM _{2S} 1	10.0	10.2	29.3	45.1	4.5	1.07 _(1.03,1.14)	10.2	35.7	0.9
# patients	5.5	6.5	8.0	11.2	4.5				
CRM _{2S} 2	1.0	42.5	33.8	19.9	2.0	1.26 _(1.13,1.53)	9.9	35.7	0.8
# patients	3.6	12.2	11.2	6.5	2.2				
CRM _{2S} 3	2.1	20.6	49.1	25.8	1.8	1.33 _(1.12,1.62)	9.7	35.8	0.6
# patients	4.1	8.3	12.7	8.1	2.6				
CRM _{2S} -MS	4.0	35.2	38.1	19.8	2.0	1.25 _(1.13,1.51)	10.1	35.7	0.9
# patients	4.2	10.0	10.7	7.8	3.0				

Table 1, continued.

Two-stage Design	Agent A (Discrete)					Agent B (Continuous)	Ave. # tox	Ave. # pts	None
	1	2	3	4	5				
Scenario 5	0.05	0.10	0.20	0.40	0.54	1.58			
CRM _{2S} 1	3.8	9.6	38.8	44.8	2.6	1.07 _(1.03,1.15)	10.4	35.9	0.4
# patients	3.6	6.1	9.8	12.1	4.3				
CRM _{2S} 2	0	34.1	44.6	19.6	1.3	1.38 _(1.16,1.74)	10.0	35.9	0.4
# patients	2.6	9.7	13.9	7.5	2.2				
CRM _{2S} 3	0.1	12.5	64.4	20.8	1.8	1.35 _(1.15,1.71)	9.9	35.9	0.4
# patients	2.8	6.7	15.6	8.1	2.7				
CRM _{2S} -MS	0.3	31.7	50.2	16.5	1.0	1.35 _(1.14,1.71)	10.2	35.9	0.3
# patients	2.7	9.0	13.4	8.2	2.6				
Scenario 6	0.02	0.08	0.12	0.20	0.40	1.58			
CRM _{2S} 1	0.2	1.1	4.8	57.5	36.4	1.19 _(1.08,1.47)	9.8	36	0
# patients	2.4	3.1	3.9	12.9	13.7				
CRM _{2S} 2	0	7.1	16.1	55.3	21.5	1.34 _(1.14,1.66)	9.1	36	0
# patients	2.3	4.4	7.6	12.6	9.1				
CRM _{2S} 3	0	1.4	16.5	62.2	19.9	1.37 _(1.15,1.80)	9.3	36	0
# patients	2.3	3.3	6.8	14.3	9.3				
CRM _{2S} -MS	0.1	4.4	13.0	55.9	26.6	1.32 _(1.13,1.66)	9.4	36	0
# patients	2.3	3.6	6.5	12.9	10.7				
Scenario 7	0.02	0.05	0.08	0.12	0.25	1.28			
CRM _{2S} 1	0	0	0.2	14.2	85.6	1.37 _(1.15,1.72)	9.0	36	0
# patients	2.2	2.4	2.4	5.5	23.5				
CRM _{2S} 2	0	1.3	3.5	27.1	68.1	1.37 _(1.17,1.70)	8.4	36	0
# patients	2.2	2.8	4.2	8.1	18.7				
CRM _{2S} 3	0	0	2.2	31.3	66.5	1.36 _(1.16,1.72)	8.6	36	0
# patients	2.2	2.4	3.5	8.8	19.1				
CRM _{2S} -MS	0	0.6	2.0	23.0	74.4	1.38 _(1.17,1.72)	8.7	36	0
# patients	2.2	2.5	3.4	7.2	20.7				
Scenario 8	0.50	0.60	0.65	0.70	0.80				
CRM _{2S} 1	24.8	0	0	0	0		9.1	13.5	75.2
# patients	12.1	1.1	0.2	0.1	0				
CRM _{2S} 2	22.4	0.1	0	0	0		8.8	13.0	77.5
# patients	11.4	1.4	0.2	0	0				
CRM _{2S} 3	23.3	0	0	0	0		8.9	13.2	76.7
# patients	11.7	1.2	0.3	0	0				
CRM _{2S} -MS	24.3	0.2	0	0	0		9.1	13.5	75.5
# patients	11.9	1.3	0.3	0	0				

was similar across all of the four designs, except that CRM_{2S} 2 treated almost twice the number of patients at dose level 2 as that of CRM_{2S} 1. Scenario 2 has the target discrete dose at the second dose level and the target continuous dose of 1.28. Both the selection percentage of the target discrete dose and the estimate of the target continuous dose using the CRM_{2S}-MS ranked the second best among all of the four designs. Particularly, the proposed CRM_{2S}-MS design recommended the target discrete dose approximately 65% of the time and estimated the continuous dose as 1.30, which is very close to the target value of 1.28. The worst skeleton corresponded to CRM_{2S} 1, which yielded a very low selection percentage of the target discrete dose, only 27.1%, but incorrectly selected the first dose with a percentage of 56.3%. The continuous dose estimate of CRM_{2S} 1 is only 1.08, which is far below the target continuous dose. Therefore, if skeleton 1 had been used for the trial conduct, the first dose would have been very likely selected. This is due to overestimation of the toxicity probability, as the first dose has a very low toxicity probability of only 0.1. Thus, inappropriately selecting the dose would cause researchers to overlook an otherwise promising drug. Scenario 3 also has the target discrete dose at the second dose level and has the target continuous dose of 1.52. Again, CRM_{2S} 1 had the worst performance overall; CRM_{2S}-MS performed much better with a selection percentage of 63.3% at the target discrete dose level and an estimated continuous dose of 1.47. The third dose is the target discrete dose in both scenarios 4 and 5. Unfortunately in scenario 4, CRM_{2S} 1 selected the fourth discrete dose with 45.1%, and CRM_{2S} 2 selected the second discrete dose with 42.5%. None of these two skeletons would lead to appropriate selection of the MTD combination. Compared with the performances of CRM_{2S} 1 and CRM_{2S} 2, the CRM_{2S}-MS selected dose level 3 as the target dose with the highest percentage. In scenario 5, the CRM_{2S}-MS had the second best performance overall with respect to the target discrete dose selection percentage and continuous dose estimation. From scenarios 6 and 7, similar conclusions can be drawn: The CRM_{2S}-MS design is indeed more robust than other versions of two-stage CRMs. Scenario 8 provides an example of even the first dose being overly toxic; all the designs terminated the trial early due to the implementation of the safety rule.

To further evaluate the performance of the CRM_{2S}-MS, Figure 2 displays the distribution of the toxicity probability for the selected dose combinations across 1,000 simulated trials. We can see that under scenarios 1 to 7, the toxicity probabilities of the selected MTD combinations are all centered around the target toxicity probability of 0.3. In scenario 8, for which all the doses were overly toxic, most of the trials were terminated early. Without loss of generality, we show the contour of the toxicity probabilities of the dose combinations under scenario 4 in Figure 3, so that we can visualize the toxicity surface under the fitted model. In other scenarios, we observe similar patterns in which toxicity increases with respect to both discrete and continuous doses.

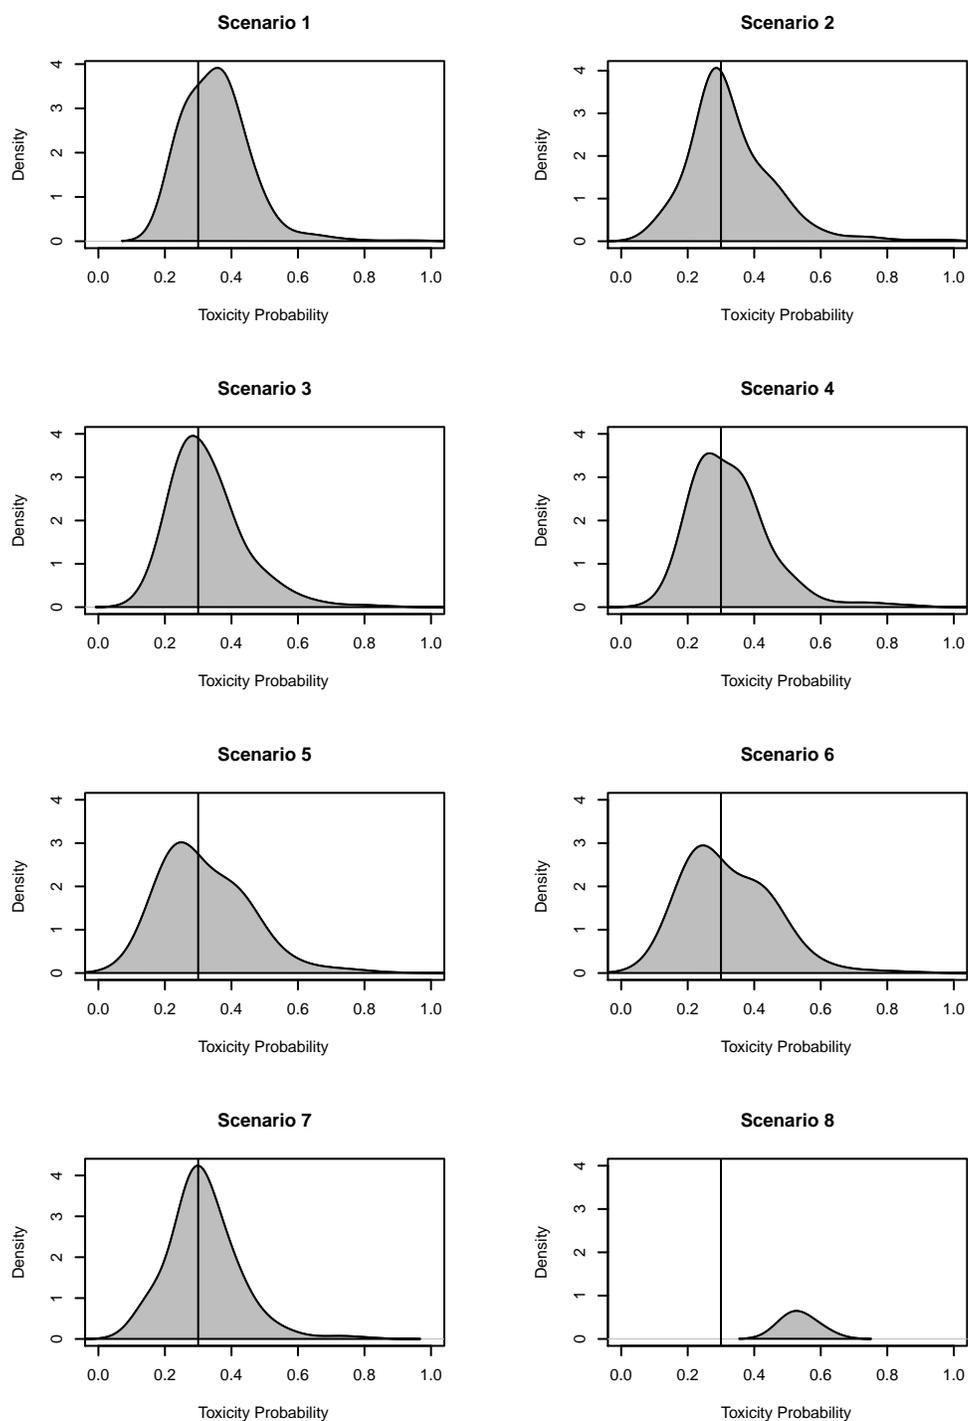


Figure 2: Distributions of the true toxicity probabilities of the selected dose combination over 1,000 simulated trials.

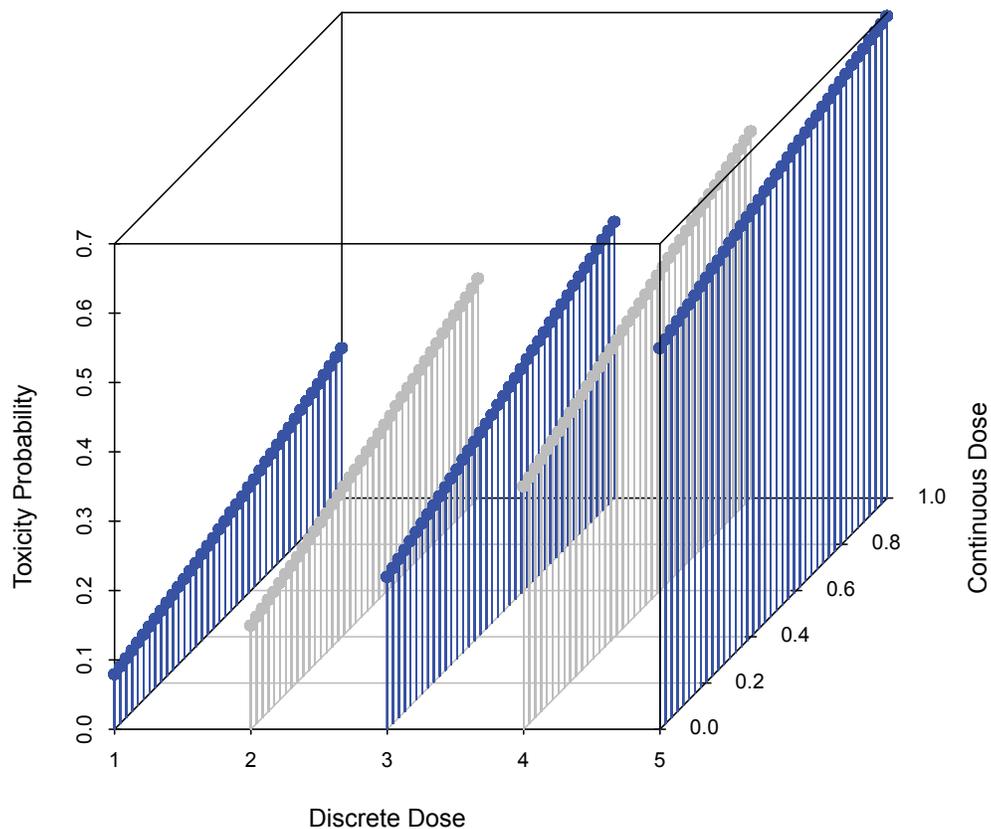


Figure 3: Toxicity probability contours of the fitted model under scenario 4.

4.2 Sensitivity Analysis

We conducted a sensitivity analysis to investigate the robustness of our two-stage CRM design when the dose-toxicity model is misspecified. Specifically, we simulated data from the model with a quadratic term,

$$\pi_d(z) = p_d^{\exp(\alpha - \beta(z - z_0) - \gamma(z - z_0)^2)}, \quad (6)$$

while we still used model (3) for dose finding. For ease of comparison, we controlled the toxicity probabilities of the discrete doses to be the same as those in Table 1, and varied the target continuous doses for different values of γ under each scenario. Table 2 shows that under all the scenarios the selection probabilities of the discrete doses were very similar to those in Table 1, and the estimated continuous doses were generally close to the target values, suggesting that the proposed two-stage CRM design is quite robust to model misspecifications.

Table 2: Sensitivity analysis of the two-stage CRM with a toxicity target of $\phi = 30\%$. We present the selection percentage for each discrete dose (the target dose in boldface), the average number of patients treated at each dose, and the median of the estimated continuous dose along with the 25th and 75th percentiles.

Two-stage Design	Agent A (Discrete)					Agent B (Continuous)		Ave. # tox	Ave. # pts	None
	1	2	3	4	5	True	Estimated			
Scenario 1	0.22	0.35	0.45	0.60	0.70					
$\gamma = 0.5$	62.0	25.2	1.5	0.1	0	1.34	1.26 _(1.13,1.41)	14.1	32.6	11.2
# patients	18.5	10.1	3.2	0.7	0					
$\gamma = 1$	62.8	24.3	1.7	0.2	0	1.29	1.21 _(1.12,1.35)	14.3	32.6	11.0
# patients	18.7	10.0	3.1	0.7	0.1					
$\gamma = 2$	63.7	22.8	1.8	0.2	0	1.24	1.20 _(1.11,1.31)	14.3	32.5	11.5
# patients	19.0	9.6	3.1	0.7	0.2					
Scenario 2	0.10	0.25	0.45	0.55	0.65					
$\gamma = 0.5$	30.8	62.1	5.6	0.6	0	1.23	1.23 _(1.11,1.40)	11.8	35.7	0.9
# patients	9.8	18.0	6.1	1.5	0.3					
$\gamma = 1$	31.8	61.2	5.3	0.2	0	1.20	1.20 _(1.11,1.35)	11.8	35.5	1.5
# patients	9.9	17.9	5.9	1.5	0.3					
$\gamma = 2$	32.4	61.9	4.4	0.3	0.1	1.17	1.16 _(1.09,1.26)	12.0	35.7	0.9
# patients	10.1	17.9	6.0	1.4	0.3					
Scenario 3	0.14	0.21	0.35	0.49	0.63					
$\gamma = 0.5$	18.5	62.1	14.3	2.4	0.1	1.38	1.31 _(1.18,1.49)	11.5	35.2	2.6
# patients	7.8	16.2	7.7	2.8	0.7					
$\gamma = 1$	19.5	62.7	13.3	2.1	0	1.32	1.26 _(1.15,1.39)	11.6	35.2	2.4
# patients	8.0	16.3	7.5	2.8	0.7					
$\gamma = 2$	20.1	61.6	13.7	2.2	0	1.26	1.22 _(1.12,1.32)	11.8	35.2	2.4
# patients	8.2	16.1	7.5	2.8	0.7					
Scenario 4	0.08	0.15	0.22	0.35	0.50					
$\gamma = 0.5$	4.3	32.1	38.9	21.4	2.7	1.34	1.20 _(1.09,1.40)	10.4	35.8	0.6
# patients	4.1	9.7	11.3	7.7	3.1					
$\gamma = 1$	4.0	31.8	39.1	21.8	2.8	1.29	1.16 _(1.08,1.31)	10.5	35.8	0.5
# patients	4.1	9.6	11.3	7.7	3.1					
$\gamma = 2$	4.4	35.1	36.9	20.3	2.7	1.24	1.14 _(1.07,1.26)	10.7	35.8	0.6
# patients	4.1	10.1	11.1	7.5	3.0					

Table 2, continued.

Two-stage Design	Agent A (Discrete)					Agent B (Continuous)		Ave. # tox	Ave. # pts	None
	1	2	3	4	5	True	Estimated			
Scenario 5	0.05	0.10	0.20	0.40	0.54					
$\gamma = 0.5$	1.1	28.2	48.0	20.9	1.6	1.41	1.23 _(1.11,1.42)	10.4	35.9	0.2
# patients	2.9	8.5	14.1	8.0	2.5					
$\gamma = 1$	0.7	31.5	49.0	17.7	1.0	1.34	1.20 _(1.10,1.36)	10.5	35.9	0.1
# patients	2.7	8.9	14.1	7.9	2.4					
$\gamma = 2$	1.0	32.5	49.7	15.7	0.9	1.28	1.16 _(1.09,1.28)	10.6	35.9	0.2
# patients	2.7	9.1	14.3	7.5	2.3					
Scenario 6	0.02	0.08	0.12	0.20	0.40					
$\gamma = 0.5$	0.1	4.2	15.2	56.3	24.2	1.41	1.22 _(1.10,1.40)	9.5	36	0
# patients	2.3	3.5	6.8	12.9	10.4					
$\gamma = 1$	0.1	4.7	17.3	54.3	23.6	1.34	1.19 _(1.09,1.34)	9.6	36	0
# patients	2.3	3.7	7.0	12.8	10.2					
$\gamma = 2$	0.1	5.2	16.5	56.9	21.3	1.28	1.14 _(1.07,1.26)	9.8	36	0
# patients	2.4	3.7	6.9	13.1	10.0					
Scenario 7	0.02	0.05	0.08	0.12	0.25					
$\gamma = 0.5$	0	0.9	2.4	23.7	73.0	1.23	1.24 _(1.12,1.38)	8.7	36	0
# patients	2.2	2.5	3.6	7.3	20.5					
$\gamma = 1$	0	1.2	2.5	23.3	73.0	1.20	1.21 _(1.10,1.35)	8.7	36	0
# patients	2.2	2.6	3.5	7.3	20.4					
$\gamma = 2$	0	0.7	3.1	23.3	72.9	1.17	1.17 _(1.08,1.27)	8.9	36	0
# patients	2.2	2.5	3.6	7.5	20.2					
Scenario 8	0.50	0.60	0.65	0.70	0.80					
$\gamma = 0.5$	24.3	0.2	0	0	0			9.1	13.5	75.5
# patients	11.9	1.3	0.3	0	0					
$\gamma = 1$	24.3	0.2	0	0	0			9.1	13.5	75.5
# patients	11.9	1.3	0.3	0	0					
$\gamma = 2$	24.3	0.2	0	0	0			9.1	13.5	75.5
# patients	11.9	1.3	0.3	0	0					

5 Conclusion

We have proposed a Bayesian adaptive two-stage dose-finding design for drug-combination trials with a continuous-dose agent and a discrete-dose agent. This two-stage procedure locates the appropriate dose for the discrete-dose agent in the first stage, then estimates the best dose for the continuous-dose agent through a two-parameter CRM model in the second stage. We have also incorporated model selection throughout the two-stage procedure to reduce the chance of a poor trial performance due to misspecification of the toxicity probabilities in the CRM. Our simulation studies have indicated that the two-parameter CRM model has desirable properties and yields good design operating characteristics.

For drug-combination trials, multiple MTDs may exist due to the toxicity equivalence contour. The two-stage CRM provides a natural and intuitive way to find one of the

MTDs, by first coarsely searching over the discrete doses and then fine searching through the continuous dose. In fact, given a known target toxicity probability, the continuous dose can be predicted at each step of dose finding. As more data are collected in the course of the trial, such prediction would become more precise. In addition, the discrete dose is not completely fixed in stage II when we search for the continuous dose. There is still some room to adjust the discrete dose in case it was over- or under-estimated in stage I. This feature inherits the spirit of jointly modeling the doses of both agents in a drug-combination trial.

References

- Braun, T. M. and Wang, S. (2010). “A hierarchical Bayesian design for phase I trials of novel combinations of cancer therapeutic agents.” *Biometrics*, 66: 805–812. [1037](#)
- Chevret, S. (2006). *Statistical Methods for Dose-Finding Experiments*. London: John Wiley & Sons Ltd. [1036](#)
- Goodman, S. N., Zahurak, M. L., and Piantadosi, S. (1995). “Some practical improvements in the continual reassessment method for phase I studies.” *Statistics in Medicine*, 14: 1149–1161. [1036](#)
- Heyd, J. M., and Carlin, B. P. (1999). “Adaptive design improvements in the continual reassessment method for phase I studies.” *Statistics in Medicine*, 18: 1307–1321. [1036](#)
- Ishizuka, N., and Ohashi, Y. (2001). “The continual reassessment method and its applications: a Bayesian methodology for phase I cancer clinical trials.” *Statistics in Medicine*, 20: 2661–2681. [1036](#)
- Korn, E. L. and Simon, R. (1993). “Using the tolerable-dose diagram in the design of phase I combination chemotherapy trials.” *Journal of Clinical Oncology*, 11: 794–801. [1036](#)
- Kramar, A., Lebecq, A. and Candalh, E. (1999). “Continual reassessment methods in phase I trials of the combination of two drugs in oncology.” *Statistics in Medicine*, 18: 1849–1864. [1036](#)
- Kuzuya, K., Ishikawa, H., Nakanishi, T., et al. (2001). “Optimal doses of paclitaxel and carboplatin combination chemotherapy for ovarian cancer: A phase I modified continual reassessment method study.” *International Journal of Clinical Oncology*, 6: 271–278. [1036](#)
- Leung, D. H-Y, and Wang, Y-G. (2002). “An extension of the continual reassessment method using decision theory.” *Statistics in Medicine*, 21: 51–63. [1036](#)
- Møller, S. (1995). “An extension of the continual reassessment methods using a preliminary up-and-down design in a dose finding study in cancer patients, in order to investigate a greater range of doses.” *Statistics in Medicine*, 14: 911–922. [1036](#)

- O’Quigley, J., Pepe, M., and Fisher, L. (1990). “Continual reassessment method: a practical design for phase 1 clinical trials in cancer.” *Biometrics*, 46: 33–48. [1036](#)
- Piantadosi, S., Fisher, J., and Grossman, S. (1998). “Practical implementation of a modified continual reassessment method for dose finding trials.” *Cancer Chemotherapy and Pharmacology*, 41: 429–436. [1036](#)
- Storer, B. E. (1989). “Design and analysis of phase I clinical trials.” *Biometrics*, 45: 925–937. [1036](#)
- Thall, P. F., Millikan, R. E., Mueller, P. and Lee, S.-J. (2003). “Dose-finding with two agents in phase I oncology trials.” *Biometrics*, 59: 487–496. [1036](#), [1037](#)
- Wages, N. A., Conaway, M. R., and O’Quigley, J. (2011). “Continual reassessment method for partial ordering.” *Biometrics*, 67: 1555–1563. [1037](#)
- Wang, K. and Ivanova, A. (2005). “Two-dimensional dose-finding in discrete dose space.” *Biometrics*, 61: 217–222. [1036](#)
- Yin, G. (2012). *Clinical Trial Design: Bayesian and Frequentist Adaptive Methods*. New Jersey: John Wiley & Sons Inc. [1036](#)
- Yin, G. and Yuan, Y. (2009a). “Bayesian model averaging continual reassessment method in phase I clinical trials.” *Journal of the American Statistical Association*, 104: 954–968. [1036](#), [1041](#)
- Yin, G. and Yuan, Y. (2009b). “Bayesian dose-finding in oncology for drug combinations by copula regression.” *Journal of the Royal Statistical Society C*, 58: 211–224. [1037](#)
- Yuan, Z., Chappell, R., and Bailey, H. (2007). “The continual reassessment method for multiple toxicity grades: a Bayesian quasi-likelihood approach.” *Biometrics*, 63: 173–179. [1036](#)
- Yuan, Y. and Yin, G. (2008). “Sequential continual reassessment method for two-dimensional dose-finding.” *Statistics in Medicine*, 27: 5664–5678. [1037](#)

Acknowledgments

We would like to thank the two referees and the Associate Editor for many insightful suggestions which immensely strengthened the work. This research is a component of Lin Huo’s doctoral dissertation at The University of Texas School of Public Health. Yuan was supported in part by a National Cancer Institute (USA) grant R01CA154591-01A1, and Yin by a grant from the Research Grants Council of Hong Kong.