AAA: Triple-adaptive Bayesian designs for the identification of optimal dose combinations in dual-agent dose-finding trials

4 authors, including:

Yuan Ji
University of Chicago
152 PUBLICATIONS   **1,911** CITATIONS

Daniel V T Catenacci
University of Chicago
170 PUBLICATIONS   **2,371** CITATIONS

Some of the authors of this publication are also working on these related projects:

Project CRM method View project

Project Bioinformatics View project

AAA: Triple-adaptive Bayesian designs for the identification of optimal dose combinations in dual-agent dose-finding trials

Jiaying Lyu

Department of Biostatistics, School of Public Health, Fudan University, Shanghai, China.

Yuan Ji

Program for Computational Biology & Medicine, NorthShore University Health System, Evanston, USA.

Department of Public Health Sciences, The University of Chicago, Chicago, USA.

E-mail: koaeraser@gmail.com

Naiqing Zhao

Department of Biostatistics, School of Public Health, Fudan University, Shanghai, China.

Daniel V.T. Catenacci

Department of Medicine, Section of Hematology & Oncology, Chicago, USA.

University of Chicago Medical Center, Chicago, USA.

Summary. We propose a flexible design for the identification of optimal dose combinations in dual-agent dose-finding clinical trials. The design is called AAA, standing for three adaptations: adaptive model selection, adaptive dose insertion, and adaptive cohort division. The adaptations highlight the need and opportunity for innovation for dual-agent dose finding, and are supported by the numerical results presented in the proposed simulation studies. To our knowledge, this is the first design that allows for all three adaptations at the same time. We find that AAA improves the statistical inference, enhances the chance of finding the optimal dose combinations, and shortens the trial duration. A clinical trial is being planned to apply the AAA design.

1. Introduction

Dual-agent dose-finding trials are becoming much more popular in oncology as more and new drugs become available. The traditional two-agent dose-finding trials often aim to capture the dose-toxicity relationship for the combinations and identify one or more maximum tolerated dose combination (MTDC) of two agents. The MTDC is defined as the dose combination at which the probability that a patient experiences dose limiting toxicity (DLT) is less than a prespecified target rate p_T, which is usually determined by physicians or clinical teams. A large number of designs have been proposed to find

the MTDC for trials with cytotoxic agents. For example, Yin and Yuan (2009) introduced Bayesian dose-finding approaches using copula regression models. Braun and Wang (2010) developed a novel hierarchical Bayesian design accounting for patients heterogeneity, and Hirakawa et al. (2013) developed a likelihood-based dose-finding method using a shrinkage logistic model. Conaway et al. (2004) estimated the MTDC by determining the complete and partial orders of the toxicity probabilities using nodal and non-nodal parameters. Later, Wages et al. (2011) applied model selection to estimate possible complete orderings associated with the partial order based on the continual reassessment method (CRM). As a review, Hirakawa et al. (2015) compared these five model-based dose-finding designs. They found that the performance of designs varied depending on the dose matrix and the location and number of true MTDCs. More recently, Mander and Sweeting (2015) published a curve-free method that relied on the product of independent beta probabilities. Lin and Yin (2016) developed a Bayesian optimal interval design for dual agents. Sun and Braun (2015) proposed a two-stage adaptive algorithm based on modified biased coin design, and Wages (2016) extended the CRM to identify an MTDC contour for dual agents.

A key assumption to all the works above is the monotonicity of the dose-toxicity response and the dose-efficacy response, which is true in the case of cytotoxic agents (Le Tourneau et al., 2009). As for many new cancer biological or immunological agents, such as the chimeric antigen receptors T-cell (CAR-T) therapies, such monotonic relationship may not be true, especially for dose-efficacy relationship (Li et al., 2016). For example, the dose-efficacy curves may follow a non-monotonic pattern, and efficacy may even decrease at higher dose levels (Hoff and Ellis, 2007). Therefore, traditional dose-finding designs with a focus on finding the MTDC are not suitable for trials of non-cytotoxic agents. In contrast to the various literature for dual cytotoxic agents dose finding, there is a scarcity of designs for non-cytotoxic agents dose finding. Instead of identifying the MTDC, one could consider the biologically optimal dose combination (BODC) for biological agents, the definition of which takes into account both efficacy and toxicity. Wages and Conaway (2014) provided a phase I/II adaptive design to find a single dose combination with an acceptable level of toxicity that maximized efficacious response. However, they assumed that the dose-toxicity and dose-efficacy relationships are monotonic among doses of one agent when the dose of other agent is fixed. Cai et al. (2014) proposed a novel dose-finding algorithm to encourage sufficient exploration of untried dose combinations in the two-dimensional space. Guo and Li (2015) used isotonic regression to estimate partially stochastically ordered marginal posterior distributions of the efficacy and toxicity probabilities to estimate the BODC.

For dual-agent trials, due to the challenges in capturing the proper therapeutic range for the dose levels of both agents, the BODC might locate outside the candidate dose range or sandwiched by existing dose combinations. Therefore, a design capable of extrapolating or interpolating a new dose combination when candidate dose combinations are deemed suboptimal can drastically improve one's chance to identify better dose combinations. To this end, Hu et al. (2013) considered an adaptive dose insertion scheme to allow new doses to be inserted during the course of a dose-finding trial. Later, Chu et al. (2016) introduced an extended version. Both methods only consider toxicity outcomes. Guo et al. (2015) proposed a toxicity- and efficacy- based dose insertion

design with adaptive model selection (TEAMS) for single-agent trials, and illustrated the importance of correct model specification for dose insertion. They showed that in order to insert the right doses, the dose-efficacy relationship must be properly identified.

In this paper, we extend the idea of TEAMS to dual agents and propose the AAA (triple A) design. The AAA design is named after three adaptive features. First, to describe the appropriate dose-efficacy curve, we present an adaptive Bayesian model selection procedure based on median posterior probability models (Barbieri and Berger, 2004) that allows the dose-efficacy model to vary between the monotone pattern and non-monotone pattern. Second, with the correct models being selected, we propose adaptive dose insertion allowing new dose combinations to be extrapolated or interpolated during the course of the trial. Last, importantly and innovatively, we consider adaptive cohort division (ACD) and allow multiple cohorts of patients to be enrolled simultaneously during the course of the trial. We show that ACD accelerates trial conduct and shortens trial duration.

We consider a conceived clinical study at The University of Chicago involving a MEK inhibitor and a PIK3CA inhibitor, both with four doses at their regular monotherapy dose, two lower doses and one higher dose. This Phase I dose finding study will enroll late-stage caner patients with a primary endpoint aiming to improve efficacy rate from 5% to 30% with the optimal tolerated dose combination. We will use this study as the basis for our numerical studies later.

The remainder of this article is organized as follows. In Sections 2 and 3, we describe the probability model and the AAA design. In Section 4, using the phase I trial we examine the operating characteristics of the AAA design through simulation studies. To evaluate the time reduction by using ACD, we examine the duration of the trial in the simulated trials. We conclude with a brief discussion in Section 5.

2. Methods

2.1. Dose-response models

Consider a trial combining J doses of agent A, denoted by $x_{a,1} < \cdots < x_{a,J}$, and K doses of agent B, denoted by $x_{b,1} < \cdots < x_{b,K}$, for dose finding. Without loss of generality, we assume $J \geq K$ and that the dosage values of the $x_{a,j}$'s and $x_{b,k}$'s have been standardized to have mean 0 and standard deviation 0.5. Let $\boldsymbol{x}_{jk} = (x_{a,j}, x_{b,k})$ denote the combination of dose levels j and k, and let $p(\boldsymbol{x}_{jk})$ and $q(\boldsymbol{x}_{jk})$ denote probabilities of the toxicity event and efficacy event for the dose combination $(x_{a,j}, x_{b,k})$, respectively, for $j = 1, 2, \ldots, J$, and $k = 1, 2, \ldots, K$.

Assume that $p(\boldsymbol{x})$ follows a linear logistic model and $q(\boldsymbol{x})$ follows a quadratic logistic model in order to incorporate a non-monotone pattern in the dose-efficacy model

$$\text{logit}\{p(\boldsymbol{x})\} = \alpha_0 + \alpha_1 x_a + \alpha_2 x_b, \tag{1}$$

$$\text{logit}\{q(\boldsymbol{x})\} = \beta_0 + \beta_1 x_a + \beta_2 x_b + \beta_3 x_a^2 + \beta_4 x_b^2, \tag{2}$$

where $\alpha_1 > 0$, $\alpha_2 > 0$, and $\boldsymbol{x} = (x_a, x_b)$ is the vector of the dose combination. Later we briefly discuss adding an interaction term $\beta_5 x_a x_b$ in the last section. Denote $\boldsymbol{\alpha} = (\alpha_0, \alpha_1, \alpha_2)'$ and $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4)'$, the vector of regression parameters in the dose-toxicity model (1) and dose-efficacy model (2), respectively. Here, we use a working

model and assume that the binary outcomes of toxicity and efficacy are independent. This working independence between efficacy and toxicity outcome in dose-finding designs has been extensively discussed in the literature (Cai et al., 2014; Ivanova et al., 2009). We also assume that toxicity is monotone with the dose as a conservative choice of model. In other words, $\alpha_1 > 0$ and $\alpha_2 > 0$ in model (1).

2.2.   Utility function and definition of BODC

Utility-based decision criteria have been adopted frequently in recent dose-finding trials (Thall and Nguyen, 2012; Lee et al., 2015; Quintana et al., 2016; Li et al., 2016). In this paper, we construct utility functions for dose safety and efficacy evaluation. Denote $U_T(p(\boldsymbol{x}), \eta_0)$ and $U_E(q(\boldsymbol{x}), \tilde{\boldsymbol{\eta}})$ the utility for safety and efficacy at dose combination $\boldsymbol{x} = (x_a, x_b)$, respectively, and define

$$U_T(p(\boldsymbol{x}), \eta_0) = \begin{cases} 1 - \frac{1-\eta_0}{p_T} p(\boldsymbol{x}), & p(\boldsymbol{x}) \in (0, p_T], \\ 0, & p(\boldsymbol{x}) \in (p_T, 1), \end{cases} \tag{3}$$

$$U_E(q(\boldsymbol{x}), \tilde{\boldsymbol{\eta}}) = \eta_1 \cdot \exp\{\eta_2 \cdot q(\boldsymbol{x})\} + \eta_3, \qquad \eta_2 > 0, \tag{4}$$

where $\tilde{\boldsymbol{\eta}} = (\eta_1, \eta_2, \eta_3)'$. See Figure 1 for an illustration. Here, the utility for safety U_T in (3) is a truncated linear decreasing function with $p(\boldsymbol{x})$; we assume that the toxicity utility U_T decreases with toxicity probability and drops to 0 if $p(\boldsymbol{x}) > p_T$, i.e., there is no utility when toxicity probability $p(\boldsymbol{x})$ is larger than p_T. Usually p_T is around 0.3 for oncology trials. The utility for efficacy U_E in (4) follows an exponential function with parameters η_1, η_2 and η_3, where $(\eta_1 + \eta_3)$ decides the utility value when there is no efficacy and (η_1, η_2) decide how fast utility increases when efficacy probability $q(\boldsymbol{x})$ increases. Combining the utilities for toxicity and efficacy, we define the overall utility score as

$$U(\boldsymbol{x}, \boldsymbol{\eta}) = U_T(p(\boldsymbol{x}), \eta_0) \cdot U_E(q(\boldsymbol{x}), \tilde{\boldsymbol{\eta}}), \tag{5}$$

where $\boldsymbol{\eta} = (\eta_0, \eta_1, \eta_2, \eta_3)'$.

Then the biologically optimal dose combination (BODC), $\boldsymbol{x}_{opt} = (x_{a,opt}, x_{b,opt})$, is defined as the dose combination that maximizes the utility function, i.e.

$$\boldsymbol{x}_{opt}(\boldsymbol{\theta}) = \underset{\boldsymbol{x}}{\operatorname{argmax}}\, U(\boldsymbol{x}, \boldsymbol{\theta}).$$

To specify the unknown values $(\eta_0, \eta_1, \eta_2, \eta_3)$, we follow a procedure suggested by Thall and Cook (2004). We elicit with physicians two pairs of toxicity-efficacy trade-off, $(0, q_1^*)$ and (p_T, q_2^*), that have the same utility value U^*, say $U^* = 0.3$. For example, $q_1^* = 0.1$, $q_2^* = 0.3$. This gives the two equations: $U_T(0, \eta_0) \cdot U_E(q_1^*, \tilde{\boldsymbol{\eta}}) = U^*$ and $U_T(p_T, \eta_0) \cdot U_E(q_2^*, \tilde{\boldsymbol{\eta}}) = U^*$. In addition, U_T and U_E must have the same scale $(0, 1)$, which implies that 1) $U_E(q(\boldsymbol{x}) = 0, \tilde{\boldsymbol{\eta}}) = \eta_1 + \eta_3 = 0$; and 2) $U_E(q(\boldsymbol{x}) = 1, \tilde{\boldsymbol{\eta}}) = \eta_1 \cdot \exp(\eta_2) + \eta_3 = 1$. Therefore, we have a set of four nonlinear equations and four unknown parameters. The numerical solution of $\hat{\boldsymbol{\eta}}$ can be easily solved numerically.

2.3.   Adaptive model selection for the dose-efficacy model

Efficacy could be either monotone or non-monotone with dose combination, depending on many factors such as the pharmacology and mechanism of action of the drug. Proposing

adaptive model selection, we allow adaptation in the model choice throughout the trial. Briefly, when the number of explored dose combinations or the sample size is small, a simpler model, such as a linear logistic model, may fit the data better to avoid the wrong estimation of dose-response curve due to the model mis-specification. As the trial proceeds, more dose combinations are explored, more patient data are accumulated, and more complex models such as a non-monotone quadratic logit model might be beneficial to obtain better estimates (Guo et al., 2015).

Consider a model selection framework for the efficacy regression coefficients in (2) as follows:

$$M_1 : \beta_3 = \beta_4 = 0; \quad M_2 : \beta_3 \neq 0, \beta_4 = 0; \quad M_3 : \beta_3 = 0, \beta_4 \neq 0; \quad M_4 : \beta_3 \neq 0, \beta_4 \neq 0.$$

Similar to Guo et al. (2015), we adopt the inverse moment priors (iMOM) (Johnson and Rossell, 2010) on β_3 under M_2 and M_4, β_4 under M_3 and M_4, in which cases either or both of them are assumed to be non-zero. The iMOM prior has no probability mass at the null point ($\beta_i = 0, i = 3, 4$) and takes the form

$$\pi(\beta_i|M_l) = \frac{k\tau^{\nu/2}}{\Gamma(\nu/2k)}|\beta_i|^{-(\nu+1)} \exp\left(\frac{\tau^k}{\beta_i^{2k}}\right), \qquad (i, l) \in \{(3, 2), (3, 4), (4, 3), (4, 4)\}$$

for $k, \nu, \tau > 0$. The choice of iMOM prior is discussed in detail in Appendix A. The prior when β_3 and β_4 equals zero is simply a point mass at zero, i.e.

$$\pi(\beta_i|M_l) = \mathbf{1}\{\beta_i = 0\}, \qquad (i, l) \in \{(3, 1), (3, 3), (4, 1), (4, 2)\},$$

where $\mathbf{1}\{\cdot\}$ is the indicator function. With no evidence favoring any of the hypotheses over the others a priori, we take $P(M_1) = P(M_2) = P(M_3) = P(M_4) = 1/4$.

In the model selection, we compute $P(M_l|Data)$, the posterior probability of each model and select the *median probability model (MPM)* to be the dose response model. MPM is defined as the model consisting of those variables which have overall posterior probability greater than or equal to 1/2 (Barbieri and Berger, 2004). Specifically, in this paper, denote p_3 and p_4 the posterior inclusion probability for the quadratic terms x_a^2 and x_b^2, respectively, and define

$$p_3 = P(M_2|Data) + P(M_4|Data), \tag{6}$$
$$p_4 = P(M_3|Data) + P(M_4|Data), \tag{7}$$

which are also the overall posterior probability that $\beta_3 \neq 0$ and $\beta_4 \neq 0$, respectively. The posterior probability of model M_l, $P(M_l|Data)$ in (6) and (7), $l = 2, 3, 4$, is given by

$$P(M_l|Data) = \frac{P(Data|M_l)P(M_l)}{\sum_{l=1}^{4} P(Data|M_l)P(M_l)},$$

where $P(Data|M_l)$ is the marginal distribution of the data under the prior of model M_l, given by

$$P(Data|M_l) = \int \pi(\boldsymbol{\alpha}, \boldsymbol{\beta}_l|M_l)\mathcal{L}(Data|\boldsymbol{\alpha}, \boldsymbol{\beta}_l, M_l)d\boldsymbol{\alpha}d\boldsymbol{\beta}_l.$$

Here $\mathcal{L}(Data|\boldsymbol{\alpha}, \boldsymbol{\beta}_l, M_l)$ is the likelihood function under model M_l, $l = 1, 2, 3, 4$, $\boldsymbol{\beta}_1 = \{\beta_0, \beta_1, \beta_2, \beta_3 = \beta_4 = 0\}$, $\boldsymbol{\beta}_2 = \{\beta_0, \beta_1, \beta_2, \beta_3, \beta_4 = 0\}$, $\boldsymbol{\beta}_3 = \{\beta_0, \beta_1, \beta_2, \beta_3 = 0, \beta_4\}$ and $\boldsymbol{\beta}_4 = \{\beta_0, \beta_1, \beta_2, \beta_3, \beta_4\}$. Since the integral does not have a closed form, numerical integration such as Monte Carlo integration is applied. Specifically, we use the harmonic mean of likelihood values here (Kass and Raftery, 1995). Let $\boldsymbol{\beta}_l^{(1)}, \boldsymbol{\beta}_l^{(2)}, \ldots, \boldsymbol{\beta}_l^{(B)}$ be a Markov chain Monte Carlo (MCMC) sample from the posterior distribution of $\boldsymbol{\beta}_l$ under model M_l; suppressing terms related to $\boldsymbol{\alpha}$ it can be shown that

$$\int \pi(\boldsymbol{\beta_l}|M_l)\mathcal{L}(Data|\boldsymbol{\beta}_l, M_l)d\boldsymbol{\beta}_l \approx \left\{ \frac{1}{B} \sum_{b=1}^{B} \mathcal{L}(Data|\boldsymbol{\beta}_l^{(b)}, M_l)^{-1} \right\}^{-1},$$

for $l = 1, 2, 3, 4$. Kass and Raftery (1995) showed that the harmonic mean approach is more efficient than directly sampling from the prior, especially when the likelihood function is highly concentrated in an area with low prior probabilities.

We perform model selection based on posterior inclusion probability, p_3 and p_4 (Table 1). For instance, If $p_3 \geq 1/2$ and $p_4 \geq 1/2$, the quadratic terms of both agents, x_a^2 and x_b^2, are included in the model (2), i.e. $\beta_3, \beta_4 \neq 0$. Therefore, we select the model M_4.

2.4. Adaptive dose combination insertion
The therapeutic window of two different drugs is often complex and difficult to delineate. In a trial that prespecifies a set of dose combinations for investigation, a new dose combination should be inserted when the BODC, $\boldsymbol{x}_{opt}$, is distant to all the existing dose combinations in the trial. This is our second proposed adaptation. Mathematically, we propose an activation rule for triggering the dose-insertion procedure. Let $\mathcal{R}_C(\boldsymbol{x}_{opt})$ represent the $C\%$ (e.g., $C = 90$) posterior credible circular region of $\boldsymbol{x}_{opt}$, defined as follows:

$$\mathcal{R}_C(\boldsymbol{x}_{opt}, r) = \{(x_{a,opt}, x_{b,opt}) : Pr\{(x_{a,opt} - x_{a,0})^2 + (x_{b,opt} - x_{b,0})^2 \leq r \mid Data\} = C\%\},$$

where $(x_{a,0}, x_{b,0})$ and r are the center and the radius of the circular region, respectively. Define A as the indicator of dose insertion,

$$A = \begin{cases} 1, & \text{if} \quad \mathcal{R}_C(\boldsymbol{x}_{opt}, r) \bigcap \{(x_{a,j}, x_{b,k}) : j = 1, \ldots, J, k = 1, \ldots, K\} = \emptyset, \\ 0, & \text{if} \quad \mathcal{R}_C(\boldsymbol{x}_{opt}, r) \bigcap \{(x_{a,j}, x_{b,k}) : j = 1, \ldots, J, k = 1, \ldots, K\} \neq \emptyset, \end{cases} \tag{8}$$

where $\emptyset$ denotes the empty set. When $A = 1$, the credible region does not cover any existing dose combinations, and the dose-insertion procedure is activated. Otherwise, the trial proceeds by treating the next cohort at one of the existing dose combinations.

2.5. Adaptive cohort division
Adaptive cohort division (ACD) is the third and an innovative adaptation. When two or more doses are considered similarly desirable for the next cohort of patients based on the collected data, the proposed AAA design allows patients to be enrolled simultaneously in parallel cohorts.

The main idea is as follows. When we encounter a toxic dose combination during the trial, a de-escalation is needed that decreases the dose level of either drug. To increase efficiency, we propose to de-escalate to two untried lower dose combinations with parallel patient enrollment at both dose combinations. That is, we open two cohorts concurrently in this case. Cohorts are collapsed if a new dose combination is inserted, in which case the new single cohort will be enrolled at the inserted dose, or the multiple cohorts all point to the same dose combination for future patients.

Because of the ACD procedure, multiple cohorts can be enrolled at the same time. Some cohorts might finish enrollment and follow up faster than others. When a cohort finishes follow up, the efficacy and toxicity response data of the patients in the cohort are observed. At this point, a decision must be made as to the next dose combination for future patients. However, at that moment there might be other cohorts still enrolling, in which case some patients might still be followed without response data while others might have completed follow up with outcomes. To fully use the existing information, we include the patients with complete data in all cohorts in the inference and decision making. In other words, we make a decision on the next dose combination based on the response data from all completers from all cohorts. This achieves faster enrollment and exploration of the new dose combinations, thereby shortening trial duration.

2.6.   Likelihood and prior specification

Let y_{jk}, z_{jk} and n_{jk} be the numbers of toxicity responses, efficacy responses and total patients treated at dose combination $(x_{a,j}, x_{b,k})$ when a cohort completes follow up during the trial, for $j = 1, 2, \cdots, J$, and $k = 1, 2, \cdots, K$. Note these numbers include all completers in all cohorts. For the observed $Data \equiv \{(y_{jk}, z_{jk}, n_{jk}), j = 1, 2, \cdots, J, k = 1, 2, \cdots, K\}$, the likelihood function under model M_l is the product of the binomial densities, i.e.

$$\mathcal{L}(Data|\boldsymbol{\alpha}, \boldsymbol{\beta}_l, M_l) \propto \prod_{j=1}^{J} \prod_{k=1}^{K} p(\boldsymbol{x}_{jk}|\boldsymbol{\alpha})^{y_{jk}} \{1 - p(\boldsymbol{x}_{jk}|\boldsymbol{\alpha})\}^{n_{jk}-y_{jk}} q(\boldsymbol{x}_{jk}|\boldsymbol{\beta}_l)^{z_{jk}} \{1 - q(\boldsymbol{x}_{jk}|\boldsymbol{\beta}_l)\}^{n_{jk}-z_{jk}}$$

where $l = 1, 2, 3, 4$ index four different models. Denote $\pi_E(\boldsymbol{\beta}_l|M_l)$ and $\pi_T(\boldsymbol{\alpha})$ the priors for $\boldsymbol{\beta}_l$ and $\boldsymbol{\alpha}$. Assuming the prior independence between $\boldsymbol{\beta}_l$ and $\boldsymbol{\alpha}$, the joint conditional posterior under M_L is given by

$$\pi(\boldsymbol{\theta}_l|Data, M_l) \propto \mathcal{L}(Data|\boldsymbol{\alpha}, \boldsymbol{\beta}_l, M_l) \pi_E(\boldsymbol{\beta}_l|M_l) \pi_T(\boldsymbol{\alpha}),$$

where $\boldsymbol{\theta}_l = (\boldsymbol{\alpha}', \boldsymbol{\beta}_l')'$.

For the prior specification of parameters in the efficacy model (2) other than β_3 and β_4, we use a weakly informative default prior for β_0, β_1 and β_2, recommended by Gelman et al. (2008). That is, $\beta_0 \sim Cauchy(0, 10)$, and $\beta_1, \beta_2 \sim Cauchy(0, 2.5)$, where $Cauchy(c, d)$ denotes a Cauchy distribution with the center parameter c and the scale parameter d. These weakly informative and appropriately regularized priors improves the estimation stability and still ensures that the data are able to dominate the priors (Gelman et al., 2008). For the iMOM priors for β_3 and β_4, we use the default values for k and ν: $k = \nu = 1$, recommended by Johnson and Rossell (2010). With respect to the

choice of parameter τ, we provide a function of the dose level $\tau = f(x_1, x_2, \cdots, x_J)$ to set τ for different scenarios. See detail in Web Appendix A.

For the toxicity model (1), we also adopt the weakly informative default prior $Cauchy(0, 10)$ for intercept α_0. We assign α_1 and α_2 independent gamma distributions with the shape parameter of 0.5 and the rate parameter of 0.5. This gives mean 1 and variance 2.

3. Trial design

3.1. Overview

The proposed dose-finding design consists of two stages. Stage I is a run-in period, in which we escalate the dose along the diagonal of the dose combination matrix in order to explore the dose-combination space quickly and collect preliminary data for stage II. This is similar to Cai et al. (2014). Stage I stops if we reach the highest dose combination $(x_{a,J}, x_{b,K})$ or if we encounter a dose combination that violates the safety requirement, i.e., $Pr\{p(\boldsymbol{x}_{jk}) > p_T | Data\} > \xi$, where ξ is a prespecified safety cutoff probability and often set to be close to 1, e.g. $\xi = 0.95$. If the dose matrix is not square (i.e., $J > K$), after first escalating the dose along the diagonal to the $(x_{a,K}, x_{b,K})$, we escalate the dose by holding the dose level of agent B at K and increasing the dose level of agent A from $(x_{a,K}, x_{b,K})$ to $(x_{a,K+1}, x_{b,K})$ and so on. After stage I, the trials enters stage II, adaptive dose finding.

In stage II, we apply the toxicity and efficacy probability models for inference, the utility function for dose assessment, and the three adaptive procedures (model selection, dose insertion and cohort division) in Section 2 for adaptive dose finding. A simple flow chart in Figure 3 depicts the flow of stage II in the AAA design. Specifically, once a cohort of patients completes follow up in the trial, we update the recorded outcome data from existing doses and enrolled patients, generate MCMC posterior samples of the parameters under M_1, M_2, M_3 and M_4 respectively, denoted by $\{\boldsymbol{\theta}_l^{(b)}, b = 1, 2, \cdots, B\}$, $l = 1, 2, 3, 4$, and carry out the adaptive model selection based on the MPM using the updated data. Suppose M_{l^*} is selected, we obtain an MCMC posterior sample of $\boldsymbol{\theta}_{l^*}$ under the selected M_{l^*}. For each simulated values $\boldsymbol{\theta}_{l^*}^{(b)}$ from the b-th MCMC iteration, $b = 1, 2, \cdots, R$, we maximize the utility function $U(\boldsymbol{x}, \boldsymbol{\theta}_{l^*}^{(b)})$ with respect to dose combination $\boldsymbol{x}$, to obtain a posterior sample of BODC, i.e.

$$\hat{\boldsymbol{x}}_{opt,l^*}^{(b)} = (\hat{x}_{a,opt,l^*}^{(b)}, \hat{x}_{b,opt,l^*}^{(b)}) = \operatorname*{argmax}_{\boldsymbol{x}} U(\boldsymbol{x}, \boldsymbol{\theta}_{l^*}^{(b)}), \quad l^* \in \{1, 2, 3, 4\}, \qquad (9)$$

for which the posterior mean of BODC is estimated to be

$$\hat{\boldsymbol{x}}_{opt} = (\hat{x}_{a,opt}, \hat{x}_{b,opt}) = \frac{\sum_{b=1}^{B} \hat{\boldsymbol{x}}_{opt,l^*}^{(b)}}{B}. \qquad (10)$$

3.2. Deciding the next dose combination

We first compute the dose-insertion activator A using the posterior sample of BODC $\{\hat{\boldsymbol{x}}_{opt,l^*}^{(b)} : b = 1, 2, \cdots, B\}$, and let

$$
\begin{aligned}
\hat{A} &= \mathbf{1}\left\{Pr\left\{(x_{a,opt,l^*} - \hat{x}_{a,opt})^2 + (x_{b,opt,l^*} - \hat{x}_{b,opt})^2 \leq \hat{r}\right\} > C\%\right\} \\
&\approx \mathbf{1}\left\{\frac{1}{B}\sum_b \mathbf{1}\left\{x_{a,opt,l^*}^{(b)} - \hat{x}_{a,opt})^2 + (x_{b,opt,l^*}^{(b)} - \hat{x}_{b,opt})^2 \leq \hat{r}\right\} > C\%\right\}, \quad (11)
\end{aligned}
$$

where $\mathbf{1}\{\cdot\}$ is the indicator function and $\hat{r}$ is the minimum Euclidean distance among the distances between the center and the existing dose combinations, denoted by

$$
\hat{r} = \min_{j,k}\left\{\sqrt{(x_{a,j} - \hat{x}_{a,opt})^2 + (x_{b,k} - \hat{x}_{b,opt})^2}\right\}.
$$

We can easily see that $\hat{A}$ in (11) is a posterior estimate of (8).

   If dose insertion is needed, i.e., $\hat{A} = 1$, we insert the new dose combination $(\hat{x}_{a,opt}, \hat{x}_{b,opt})$ and two sets of new dose combinations

$$
\{(x_{a,1}, \hat{x}_{b,opt}), \cdots, (x_{a,J}, \hat{x}_{b,opt})\} \quad \text{and} \quad \{(\hat{x}_{a,opt}, x_{b,1}), \cdots, (\hat{x}_{a,opt}, x_{b,K})\}
$$

into the dose combination matrix, as shown in Figure 2, and assign the next cohort to the new dose combination $(\hat{x}_{a,opt}, \hat{x}_{b,opt})$.

   If dose insertion is not needed, i.e., $\hat{A} = 0$, we assign the next cohort of patients according to the utility of the existing dose combinations. Let N denote the prespecified maximum sample size, N_1 denote the number of patients in stage I, and $N_2 = N - N_1$ be the total number of patients available for stage II. Given the current dose combination $\boldsymbol{x}_{jk} = (x_{a,j}, x_{b,k})$, we define 1-degree admissible dose set, denoted by $\mathcal{A}_1$, as dose combination $\boldsymbol{x}_{j'k'}$, whose dose levels are different from $\boldsymbol{x}_{jk}$ no more than 1 level and satisfy the safety requirement, i.e., $\mathcal{A}_1 = \{\boldsymbol{x}_{j'k'} : |j' - j| \leq 1, |k' - k| \leq 1, Pr\{p(\boldsymbol{x}_{j'k'}) > p_T | Data\} \leq \xi\}$, where ξ is close to 1. Then the next dose combination is decided based on the following algorithm:

If $(x_{a,j}, x_{b,k})$ is considered safe, i.e., $Pr\{p(\boldsymbol{x}_{jk}) > p_T | Data\} \leq \xi$, where ξ is close to 1, assign patients as follows:

   [a.] Based on the accumulated trial data, determine dose set $\mathcal{A}_1$.

   [b.] Among the dose combinations in $\mathcal{A}_1$, compute the posterior mean utility for each combination, i.e., $\bar{U}(\boldsymbol{x}_{j'k'}, \boldsymbol{\theta}_{l^*}) = \frac{1}{B}\sum_{b=1}^{B} U(\boldsymbol{x}_{j'k'}, \boldsymbol{\theta}_{l^*}^{(b)})$, $\boldsymbol{x}_{j'k'} \in \mathcal{A}_1$ and identify the dose combination $\boldsymbol{x}_{j^*k^*} = (x_{a,j^*}, x_{b,k^*})$ with the highest posterior mean utility under the safety constrain $(j^* - j) + (k^* - k) \leq 1$, i.e.,

$$
\boldsymbol{x}_{j^*k^*} = \underset{\boldsymbol{x}_{j'k'}}{\arg\max}\{\bar{U}(\boldsymbol{x}_{j'k'}, \boldsymbol{\theta}_{l^*})\} \text{ subject to } \boldsymbol{x}_{j'k'} \in \mathcal{A}_1, \text{ and } (j' - j) + (k' - k) \leq 1 \quad (12)
$$

   [c.] If dose combination $\boldsymbol{x}_{j^*k^*}$ has not been used to treat any patient thus far, or all doses in $\mathcal{A}_1$ have been used to treat patients, we assign the next cohort of patients to

$\boldsymbol{x}_{j^*k^*}$. However, if $\boldsymbol{x}_{j^*k^*}$ has been used and there are some untried dose combinations in $\mathcal{A}_1$, we assign the next cohort of patients to $\boldsymbol{x}_{j^*k^*}$ only if

$$\hat{Pr}\{U(\boldsymbol{x}_{j^*,k^*}, \boldsymbol{\theta}_{l^*}) > U_0 \mid Data\} \approx \frac{1}{B} \sum_{b=1}^{B} \mathbf{1}\{U(\boldsymbol{x}_{j^*k^*}, \boldsymbol{\theta}_{l^*}^{(b)}) > U_0\} > \left(\frac{N_2 - n_2}{N_2}\right)^{\omega}$$

where U_0 is the lowest acceptable utility value, n_2 is the total number of patients that have been treated in stage II and ω is a known tuning parameter controlling how stringent the threshold is. Otherwise, exclude $\boldsymbol{x}_{j^*k^*}$ from $\mathcal{A}_1$ and return to step b.

If $(x_{a,j}, x_{b,k})$ is deemed toxic, i.e., $Pr\{p(\boldsymbol{x}_{jk}) > p_T|Data\} > \xi$, de-escalate to the untried 1-degree lower doses, $\boldsymbol{x}_{j-1,k} = (x_{a,j-1}, x_{b,k})$ or $\boldsymbol{x}_{j,k-1} = (x_{a,j}, x_{b,k-1})$ or both. That is, if both dose combinations exist and have not been used, two cohorts of patients are recruited and assigned to $(x_{a,j-1}, x_{b,k})$ and $(x_{a,j}, x_{b,k-1})$, simultaneously. If only one dose exists and has not been used, assign the next cohort of patients to this dose. If both doses exist but both have been used, terminate this cohort and do not recommend any dose for next cohort until there is a cohort newly completed.

As seen above, we adopt a concept of 1-degree admissible neighbor $\mathcal{A}_1$ in assigning the next dose combinations when dose insertion is not needed. Cai et al. (2014) demonstrate that this admissible neighbor and adaptive rule (step c) not only encourage the exploration of untried dose combinations to avoid the problem of trapping in suboptimal doses, but also restrict the dose escalation/de-escalation within the neighbors of the current dose, avoiding dramatic dose changes and improving the reliability of the dose-finding.

3.3. Dose-finding Algorithm

The AAA design is summarized in Box 1. Additional rules listed in Box 2 are for ethics and stability concern. For brevity, we use "dose "to denote a dose combination hereinafter.

Box 1: The AAA design for phase I/II dose finding trials

The trial starts with the treatment of the first cohort of patients at the lowest dose $(x_{a,1}, x_{b,1})$. Suppose that patients are being treated at dose $\boldsymbol{x}_{jk} = (x_{a,j}, x_{b,k})$. A dose is deemed toxic if $Pr\{p(\boldsymbol{x}_{jk}) > p_T | Data\} > \xi$, where ξ is close to 1; otherwise, the dose is safe. Let N_1 and N_2 denote the maximum sample size of stage I and II, respectively. Let n_2 be the number of patients currently enrolled in stage II.

Stage I Run in Period

I1 If dose $(x_{a,j}, x_{b,k})$ is safe, escalate diagonally and treat the next cohort at $(x_{a,j+1}, x_{b,k+1})$. If $j = k = K$, escalate to $(x_{a,j+1}, x_{b,K})$.

I2 Stage I is complete when either dose $(x_{a,j}, x_{b,k})$ is deemed toxic or the highest dose combination $(x_{a,J}, x_{b,K})$ is reached. Stage II starts.

Stage II Adaptive Dose Finding

II1 Once a cohort completes follow up, collect their efficacy and toxicity outcomes.

II2 Using the accumulated trial data on all the completers, generate MCMC posterior samples of parameters under models M_1, M_2, M_3 and M_4, respectively, and carry out adaptive model selection. Suppose model l^* is selected, denote the MCMC posterior sample $\{\boldsymbol{\theta}_{l^*}^{(b)}, b = 1, \ldots, B\}$ under the selected dose-efficacy model M_{l^*}, $l^* \in \{1, 2, 3, 4\}$.

II3 Obtain a posterior sample $\boldsymbol{x}_{opt}$ under M_{l^*}, i.e. $\{\hat{\boldsymbol{x}}_{opt,l^*}^{(b)}, b = 1, 2, \cdots, B\}$, from (9). Then compute the posterior mean BODC, $\hat{\boldsymbol{x}}_{opt} = (\hat{x}_{a,opt}, \hat{x}_{b,opt})$, from (10), and the decision indicator $\hat{A}$ in (11).

(a) If $\hat{A} = 1$, the new dose $\hat{\boldsymbol{x}}_{opt}$ is inserted in the trial and assigned to the next cohort. In addition, two sets of doses, $\{(x_{a,1}, \hat{x}_{b,opt}), \cdots, (x_{a,J}, \hat{x}_{b,opt})\}$ and $\{(\hat{x}_{a,opt}, x_{b,1}), \cdots, (\hat{x}_{a,opt}, x_{b,K})\}$, are inserted as well. Then go to step II1 and wait for the completion of this cohort.

(b) If $\hat{A} = 0$ and $\boldsymbol{x}_{jk}$ is safe,

  b1) identify $\mathcal{A}_1$ as the set of safe neighbors of $\boldsymbol{x}_{jk}$ with degree 1.

  b2) In $\mathcal{A}_1$, identify the dose $\boldsymbol{x}_{j^*k^*}$ that has the highest posterior mean utility under the safety constraint $(j^* - j) + (k^* - k) \leq 1$, from (12).

  b3) If $n_{j^*k^*} = 0$ or $n_{rs} \neq 0$, $\forall\, \boldsymbol{x}_{r,s} \in \mathcal{A}_1$, treat the next cohort at dose $\boldsymbol{x}_{j^*k^*}$. Otherwise, if $Pr\{U(\boldsymbol{x}_{j^*,k^*}, \boldsymbol{\theta}_{l^*} \mid Data) > U_0\} > \left(\frac{N_2 - n_2}{N_2}\right)^{\omega}$, treat the next cohort at $\boldsymbol{x}_{j^*k^*}$; otherwise, remove $\boldsymbol{x}_{j^*k^*}$ from $\mathcal{A}_1$ and go to step b2.

(c) If $\hat{A} = 0$ and $\boldsymbol{x}_{jk}$ is toxic, de-escalate to the untried 1-degree lower doses allowing cohort division:

  - If $\{j, k \geq 2$ and $n_{j-1,k} = n_{j,k-1} = 0\}$, simultaneously enroll two cohorts of patients at both doses $\boldsymbol{x}_{j-1,k}$ and $\boldsymbol{x}_{j,k-1}$;

  - If $\{j, k \geq 2$ and $n_{j-1,k} = 0$ but $n_{j,k-1} > 0\}$, or $\{j \geq 2, k = 1, \text{and } n_{j-1,k} = 0\}$, assign the next cohort to dose $\boldsymbol{x}_{j-1,k}$;

  - If $\{j, k \geq 2$ and $n_{j,k-1} = 0$ but $n_{j-1,k} > 0\}$, or $\{k \geq 2, j = 1, \text{and } n_{j,k-1} = 0\}$, assign the next cohort to dose $\boldsymbol{x}_{j,k-1}$;

  - Otherwise, terminate this cohort and do not recommend any dose.

(d) If no dose is recommended in (a)-(c), assign the next cohort to the dose $\boldsymbol{x}_{\bar{j}\bar{k}}$ which has the highest posterior mean utility among all the existing safe doses, i.e., $\boldsymbol{x}_{\bar{j}\bar{k}} = \operatorname{argmax}_{\boldsymbol{x}_{j'k'}} \{\bar{U}(\boldsymbol{x}_{j'k'}, \boldsymbol{\theta}_{l^*})\}$

II4 Repeat steps II1-II3 until the maximum sample size $N = N_1 + N_2$ is reached.

II5 Select the dose that has the highest mean utility among all tested safe doses, including the newly inserted dose.

Box 2: Practical rules

- **Rule 1 [Dose extrapolation]** The inserted new dose is not allowed to be more than twice the highest dose or less than half of the lowest dose that has been used in the trial.

- **Rule 2 [Early Termination]** If the lowest dose $\boldsymbol{x}_{11} = (x_{a,1}, x_{b,1})$ is deemed toxic, i.e., $Pr\{p(\boldsymbol{x}_{11}) > p_T | Data\} > \xi$, where ξ is close to 1, and no new dose is inserted, terminate the trial.

- **Rule 3 [Dose Exclusion]** If the dose $\boldsymbol{x}_{jk} = (x_{a,j}, x_{b,k})$ is deemed toxic, i.e., $Pr\{p(\boldsymbol{x}_{jk}) > p_T | Data\} > \xi$, where ξ is close to 1, exclude doses $\{(x_{a,j'}, x_{b,k'}) : j' = j, j+1, \ldots, J, \ k' = k, k+1, \ldots, K\}$, i.e. these doses will never be used in the trial again.

- **Rule 4 [No Skipping Dose]** Restrict the escalation to 1 level increment, i.e. there is no skipping in the escalation. Particularly, if the new dose intended for insertion is higher than any unexplored dose, pause the insertion and go to step II3(b)-II3(d).

---

4. Simulation

4.1. Simulation setup

We consider the motivating trials combining two agents, a MEK inhibitor and a PIK3CA inhibitor, each with 4 dose levels. The maximum sample size is 96 and cohort size is 3. We investigate 10 different scenarios, and all scenarios assume a true linear or quadratic logistic model for both agents in the dose-efficacy relationship, as shown in Figure 4. For each scenario, 1,000 simulated trials are conducted. In the proposed design, we set the MTD toxicity threshold $p_T = 0.3$, the credible level threshold $C\% = 0.90$, the lowest acceptable utility value $U_0 = 0.1$ and the tuning parameter $\omega = 2$. The probability threshold $\xi = 0.95$ for the practical rule and safe requirement. Regarding the utility function, we assume that the two toxicity and efficacy rate pairs, $(0, 0.45)$ and $(0.3, 0.85)$ have the same utility value 0.3. As a result, we obtain the estimated $\hat{\boldsymbol{\eta}} = (0.396, 0.385, 1.280, -0.385)$.

For MCMC computation, we adopt a standard random walk Metropolis-Hasting algorithm. And for each chain, 10,000 MCMC samples are drawn with a burn-in size of 5,000 iterations. The MCMC mixed fast and well with no sign of convergence problems.

For comparison, we apply the design in Cai et al. (2014). For fairness, we slightly modified this design by using the utility function rather than efficacy probability for defining the dose in admissible dose set $\mathcal{A}$. This typically improved the performance of their design based on our experience. There is no dose insertion and model selection in their algorithm. Therefore, we only compare the dose allocation and dose selection. To demonstrate the benefit of adaptive cohort division, we turn off the ACD procedure in AAA and apply a single cohort algorithm, in which we randomly select one dose in step IIc if two doses are available for ACD. See Web Appendix B for details of the simulation scheme for patients enrollment and follow-up.

4.2. Operating characteristics

The operating characteristics of the proposed algorithm for Scenarios 1-6 are summarized in Table 2 with seven sections per scenario. Section 1 gives a brief description of the true dose response and the need for insertion. Sections 2 and 3 provide the mean (standard deviation) of the posterior means across the 1,000 simulated trials for the regression parameters in the dose-toxicity and dose-efficacy models, respectively. Section 4 summarizes the model selection frequencies. Section 5 lists the prespecified dose levels, and the true toxicity and efficacy probabilities of all the doses. Section 6 presents the detailed operating characteristics of the AAA design in terms of patient allocation and dose selection. Lastly, Section 7 presents the results of the modified algorithm from Cai et al. (2014). For brevity, we present simulation results for Scenarios 7-10 in Web Table 2.

In Scenario 1, the efficacy rates firstly increase and decrease later with both agents. The true BODC $x_{opt} = (0.441, 0.476)$ is bracketed by doses $(0.3, 0.35)$, $(0.3, 0.65)$, $(0.6, 0.35)$ and $(0.6, 0.65)$. Therefore, new dose should be inserted (Figure 4). From Table 2, we see that the AAA design inserts new doses with mean $(0.430, 0.448)$ in 59.3% of the simulated trials. The utility of the mean inserted dose combination under the true model is 0.465, which is higher than that of all prespecified dose combinations. Among all the patients, 25.9% are treated at the inserted new dose combinations. At the end of 55.9% of the simulated trials, an inserted dose is claimed to be the BODC. The mean selected BODC at the end of the trial is $(0.456, 0.464)$, close to the true BODC. Also, 90.5% of the trials correctly choose the quadratic logistic regression for both agents for dose-efficacy curve at the end, and the posterior sample means of β are close to the true values. Section 7 shows the selection percentages and patient allocation based on Cai et al. (2014). The results are reasonable as most patients are allocated to the four doses that surround the true BODC. However, since their design dose not allow dose insertion, it cannot correctly identify the true BODC.

Scenario 2 has a similar pattern of dose-toxicity and dose-efficacy curves with Scenario 1. From Figure 4, there is no need for insertion. The AAA design only inserts new doses in 15.5% of trials. The mean selected dose combination is $(0.454, 0.466)$, close to the true BODC $(0.442, 0.477)$.

Scenario 3 reflects a setting where only a few doses are tolerable while others are overly toxic (Figure 4). Because the utility of the existing doses $x_{31} = (0.55, 0.2)$ and $x_{22} = (0.25, 0.5)$ are 0.240 and 0.217, respectively, which are not much different from the utility of true BODC $(0.41, 0.332)$, 0.297, only 29.6% of the trials insert new doses. Note that in this scenario, AAA selects the true quadratic logistic model for the dose-efficacy 73% of the times, a smaller percentage than Scenarios 1 and 2. This is because only a few patients (22.7% in total) are assigned to the doses at the upper right corner of the dose matrix (shown in Figure 4) due to their high toxicity. Therefore, there is insufficient amount of information for better model selection in this scenario.

Scenario 4 is a situation when all combinations are higher than MTD, and hence 8.6% of the trials are terminated at an early stage according to Practical Rule 2. Among completed trials, the proportion of selecting the true quadratic model for efficacy is 67.8%. New dose combinations with a mean $(0.354, 0.354)$ are inserted among 92.5% of the completed trials, and 96.7% end with selecting new dose combinations as BODC.

The mean selected dose combination is $(0.242, 0.245)$, which is lower than the MTD and close to the true BODC $(0.304, 0.262)$. This scenario demonstrates the important safety feature of the proposed design. Also, AAA not only stops the trial early due to toxicity, it performs desirable dose insertion below all the prespecified doses and identifies the correct BODC.

Unlike the previous four scenarios, Scenario 5 presents a situation where the pre-specified dose matrix only covers the bottom-left corner of the quadratic dose-efficacy curve. Thus, the efficacy grows with both agents and the true BODC $(0.474, 0.472)$ locates at the upper-right corner beyond the dose matrix. The insertion rate is 11.1% and the mean selected BODC is $(0.358, 0.345)$, which is poorly estimated. And all four dose-efficacy models are selected at similar rates. There are two reasons. First, because the prespecified doses do not cover a wide range of dose response surface, data on these doses could not provide a good estimate of the entire curve. Second, with a relatively small sample size $N = 96$, the simulated trials often run out of patients before the BODC combination is reached. If a large N is allowed, the dose matrix could be extrapolated well and the dose insertion algorithm would perform better (results not shown).

Scenario 6 assumes that the true dose-efficacy is a linear logistic model. About 32.2% of the simulated trials insert new doses, which is as expected, and the linear model is selected at a relatively high rate, 70.3%. A total of 28.7% of trials select the new doses as the BODC with the mean $(1.002, 0.341)$.

In Scenarios 7-10 (see Web Table 2), the dose-efficacy curves increase first and fall later with one agent, but are monotone with the other agent. In addition, the true BODC are all located outside the prespecified dose matrix at the right-side, top-side, left-side and bottom-side regions, respectively. In all four scenarios, insertion is needed, and the average insertion rates are 44.2%, 29.7%, 88.2% and 79.9%, respectively. The mean selected doses are close to the true BODCs in all four scenarios.

Comparing to the results of Cai et al. (2014) in Section 7 of all scenarios, we can see that the doses close to the true BODC are assigned more patients and selected more frequently. This shows that their design performs well in general. However, when encountering situations where dose insertion is needed, their design is unable to select the true BODC due to the lack of dose insertion.

4.3.   Time duration

Table 3 demonstrates the benefit of ACD in shortening trial duration. In particular, step IIc in Box 1 is expected to speed up the trial process and reduce the time duration. It can be seen from Table 3 that about 100 days or more can be saved across most scenarios with the ACD procedure. The trial is never longer with ACD than without ACD. Scenarios with more toxic doses result in more reduction of trial time. For example, for Scenario 3, the trial duration is reduced by about 400 days. However, if all doses are overly toxic, the reduction of time duration is negligible, since cohort division is not allowed for inserted doses. The performance of multiple cohorts and single cohorts are almost the same in terms of the mean selected BODC, patients allocation, and the percentage of being selected as the BODC.

4.4. Sensitivity to sample size

Lastly, to evaluate the effect of sample size, we apply the algorithm with a smaller sample size of 66 on Scenarios 1, 4 and 5. Results are summarized in Web Table 3. We find the that the reduction of sample size has larger influence on Scenario 1 than Scenarios 4 and 5. Specifically, for Scenario 1, the insertion rate is reduced to 42% (from 55.9%) and the proportion of selecting the correct quadratic model is declined to 87.5% (from 90.5%), although the mean inserted dose combination $(0.44, 0.455)$, and the mean selected BODC $(0.460, 0.466)$, are still close to the truth. Overall, the AAA perform reasonably well with a significant reduction of sample size.

5. Discussion

We propose a new Bayesian adaptive dose insertion design for dual-agents phase I/II oncology trials. The dose-insertion procedure based on both efficacy and toxicity enables us to locate more desirable dose combinations. Bayesian model selection during the trial allows the dose-efficacy relationship to be adapted between linear and quadratic logistic models. The model selection has been shown be important (Guo et al., 2015) in dose insertion and maintaining a high efficiency of the dose-finding trial. The adaptive cohort division speeds up the trial process and shortens the time duration in most scenarios. Simulation results show that the proposed design has superior operating characteristics.

The AAA design is a utility-based method. Clearly, the performance heavily depends on the definition of utility. Although we choose utility as a multiplication of a linear truncated function (safety utility) and an exponential function (efficacy) in this paper, one can use other reasonable alternative utilities. Choosing an appropriate utility function has to be done for individual trials and through discussion between clinicians and statisticians. For different diseases and drugs, different trade-off between efficacy and toxicity might be allowed. Nevertheless, changes on the utility function is a separate topic and does not affect the overall statistical design illustrated in Boxes 1 and 2. In other words, the algorithm presented therein is expected to find the optimal dose combination with high likelihood based on the defined utility, however it is defined.

In our models, we do not include an interaction effect $\beta_5 x_a x_b$ for the two agents for a couple of reasons. First, our current clinical trial cannot afford to recruit a large number patients, say a few hundreds, which is typically needed for estimating interaction terms in a regression model reliably. Second, it has been demonstrated that for the purpose of dose finding, a local fit with a working model of the response surface does not affect much the efficiency of dose finding. That is, even when models are mis-specified (by not including a true interaction term), the dose-finding decisions are not affected severely and still lead to reasonable performance. This has been shown in Cai et al. (2014) and Wang and Ivanova (2005) in the context of drug combination trials. To see this, we simulated two scenarios consisting of the interaction term $\beta_5 x_a x_b$ (Web Figure 2) and summarize results in Web Table 4. We can see that the AAA design performs well in finding the optimal dose combination and allocating patients to desirable dose combinations. The inference on the the parameter of the interactive term β_5 is biased. Typically, accurate estimation of β_5 requires a large sample size and well-placed dose combinations across the dose-response surface. In our current example, we only have 4×4 dose levels for two

agents, which leads to the difficulty in capturing the rotation of the efficacy curve, i.e., β_5. We found that when we increased to a combination of 6×6 and 7×7 dose levels, the interaction can be well estimated (results are not shown).

Supplementary Materials

Appendices, Tables, and Figures, referenced in Sections 2.3, 2.6, 4.1, 4.2 and 5, are available with this paper at the online website.

Acknowledgements

References

Barbieri, M. M. and Berger, J. O. (2004) Optimal predictive model selection. *Annals of Statistics*, 870–897.

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.

Cai, C., Ji, Y. and Yuan, Y. (2014) A bayesian dose-finding design for oncology clinical trials of combinational biological agents. *Journal of the Royal Statistical Society Series C*, **63**, 159–173.

Chu, Y., Pan, H. and Yuan, Y. (2016) Adaptive dose modification for phase i clinical trials. *Statistics in medicine*, **35**, 3497–3508.

Conaway, M. R., Dunbar, S. and Peddada, S. D. (2004) Designs for single-or multiple-agent phase i trials. *Biometrics*, **60**, 661–669.

Gelman, A., Jakulin, A., Pittau, M. G. and Su, Y.-S. (2008) A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 1360–1383.

Guo, B. and Li, Y. (2015) Bayesian dose-finding designs for combination of molecularly targeted agents assuming partial stochastic ordering. *Statistics in medicine*, **34**, 859–875.

Guo, W., Ni, Y. and Ji, Y. (2015) Teams: Toxicity-and efficacy-based dose-insertion design with adaptive model selection for phase i/ii dose-escalation trials in oncology. *Statistics in biosciences*, **7**, 432–459.

Hirakawa, A., Hamada, C. and Matsui, S. (2013) A dose-finding approach based on shrunken predictive probability for combinations of two agents in phase i trials. *Statistics in medicine*, **32**, 4515–4525.

Hirakawa, A., Wages, N. A., Sato, H. and Matsui, S. (2015) A comparative study of adaptive dose-finding designs for phase i oncology trials of combination therapies. *Statistics in medicine*, **34**, 3194–3213.

Hoff, P. M. and Ellis, L. M. (2007) Targeted therapy trials: approval strategies, target validation, or helping patients? *Journal of Clinical Oncology*, **25**, 1639–1641.

Hu, B., Bekele, B. N. and Ji, Y. (2013) Adaptive dose insertion in early phase clinical trials. *Clinical Trials*, **10**, 216–224.

Ivanova, A., Liu, K., Snyder, E. and Snavely, D. (2009) An adaptive design for identifying the dose with the best efficacy/tolerability profile with application to a crossover dose-finding study. *Statistics in Medicine*, **28**, 2941–2951.

Johnson, V. E. and Rossell, D. (2010) On the use of non-local prior densities in bayesian hypothesis tests. *Journal of the Royal Statistical Society Series B*, **72**, 143–170.

Kass, R. E. and Raftery, A. E. (1995) Bayes factors. *Journal of the American Statistical Association*, **90**, 773–795.

Le Tourneau, C., Lee, J. J. and Siu, L. L. (2009) Dose escalation methods in phase i cancer clinical trial. *Journal Of the National Cancer Institute*, **101**, 708–720.

Lee, J., Thall, P. F., Ji, Y. and Müller, P. (2015) Bayesian dose-finding in two treatment cycles based on the joint utility of efficacy and toxicity. *Journal of the American Statistical Association*, **110**, 711–722.

Li, D. H., Whitmore, J. B., Guo, W. and Ji, Y. (2016) Toxicity and efficacy probability interval design for phase i adoptive cell therapy dose-finding clinical trials. *Clinical Cancer Research*. URL: `http://clincancerres.aacrjournals.org/content/early/2016/12/12/1078-0432.CCR-16-1125`.

Lin, R. and Yin, G. (2016) Bayesian optimal interval design for dose finding in drug-combination trials. *Statistical methods in medical research*, 0962280215594494.

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

Quintana, M., Li, D. H., Albertson, T. M. and Connor, J. T. (2016) A bayesian adaptive phase 1 design to determine the optimal dose and schedule of an adoptive t-cell therapy in a mixed patient population. *Contemporary clinical trials*, **48**, 153–165.

Sun, Z. and Braun, T. M. (2015) A two-dimensional biased coin design for dual-agent dose-finding trials. *Clinical Trials*, **12**, 596–607.

Thall, P. F. and Cook, J. D. (2004) Dose-finding based on efficacy–toxicity trade-offs. *Biometrics*, **60**, 684–693.

Thall, P. F. and Nguyen, H. Q. (2012) Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *Journal of biopharmaceutical statistics*, **22**, 785–801.

Wages, N. A. (2016) Identifying a maximum tolerated contour in two-dimensional dose finding. *Statistics in medicine*.

Wages, N. A. and Conaway, M. R. (2014) Phase i/ii adaptive design for drug combination oncology trials. *Statistics in medicine*, **33**, 1990–2003.

Wages, N. A., Conaway, M. R. and O'Quigley, J. (2011) Dose-finding design for multi-drug combinations. *Clinical Trials*, **8**, 380–389.

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

Yin, G. and Yuan, Y. (2009) Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **58**, 211–224. URL: `http://dx.doi.org/10.1111/j.1467-9876.2009.00649.x`.
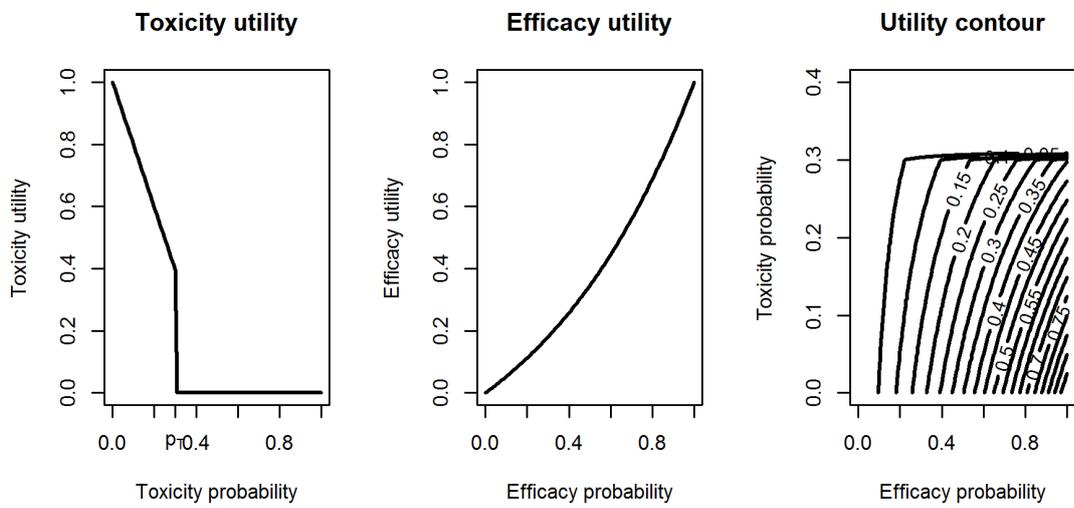
Fig. 1: Utility functions: From left to right are utility for safety (truncated at p_T and sharply decreases to 0), utility for efficacy and the overall utility contours.
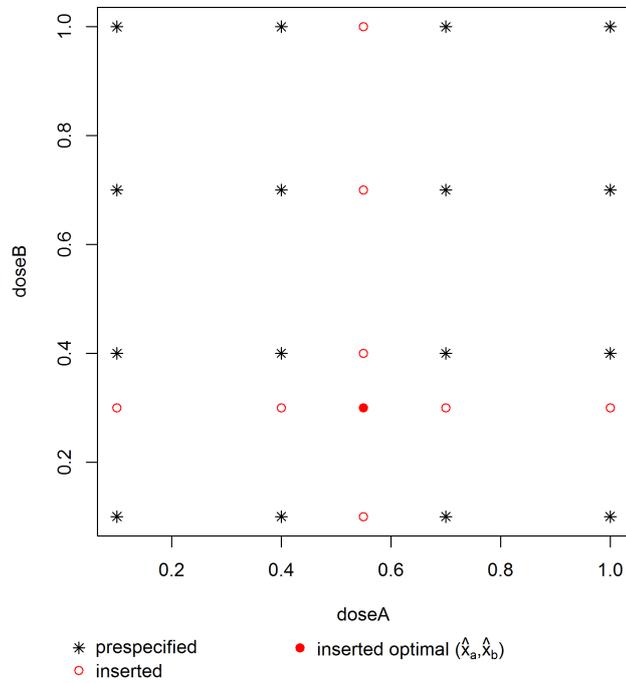
Fig. 2: The procedure of dose insertion. Black dots: prespecified dose combinations, red dots: inserted dose combinations.
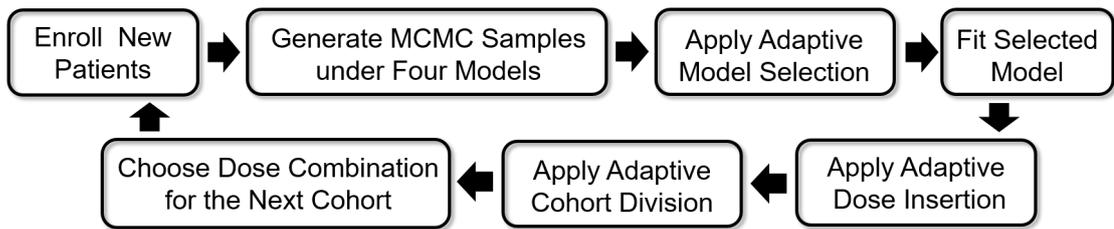
Fig. 3: A simple flow chart for stage II in the AAA design.
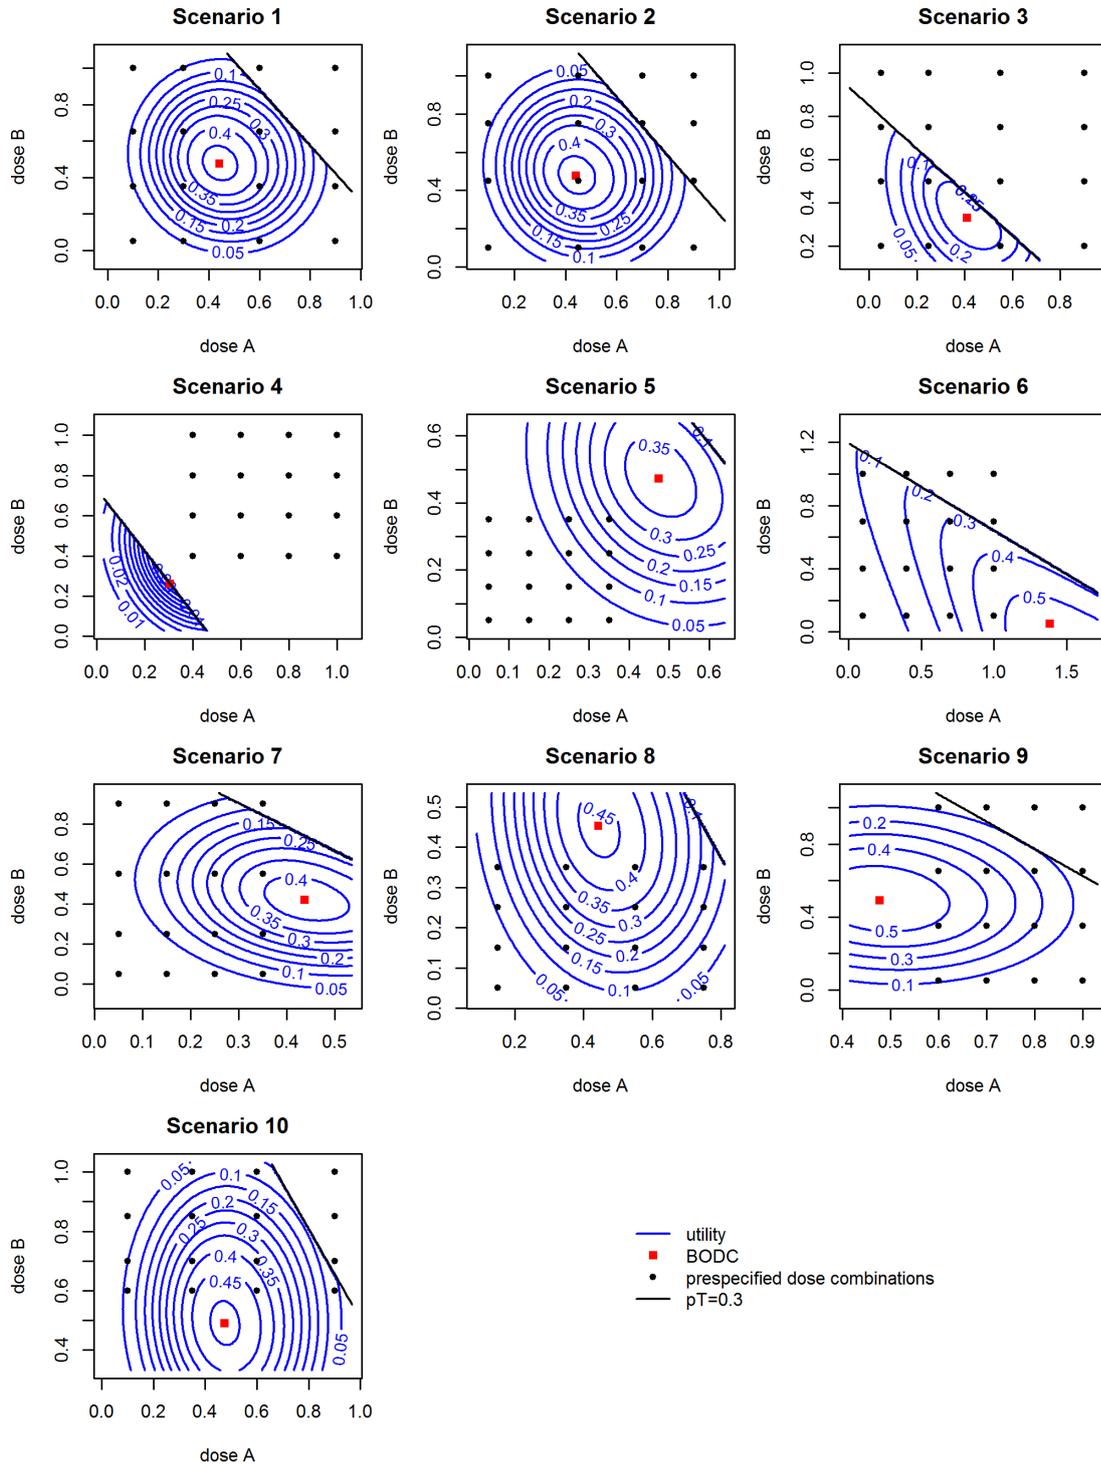
Fig. 4: Ten scenarios in the simulation study. Black dots represent the prespecified dose combinations. Red dots marks the true BODC. Blue contour: dose-utility contour. Black line: dose combinations with toxicity probability p_T. Any dose combinations beyond the p_T line with higher toxicity probability are of 0 utility.

Table 1: Median probability model selection rules.

p_3	p_4	β_3	β_4	The Selected Dose-Efficacy Model
$< 1/2$	$< 1/2$	$= 0$	$= 0$	linear logistic model, M_1
$\geq 1/2$	$< 1/2$	$\neq 0$	$= 0$	Quadratic for agent A, M_2
$< 1/2$	$\geq 1/2$	$= 0$	$\neq 0$	Quadratic for agent B, M_3
$\geq 1/2$	$\geq 1/2$	$\neq 0$	$\neq 0$	Quadratic for both agents A and B, M_4

Table 2: Simulation results for scenarios 1-6. For each scenario, 1,000 trials are conducted on computer and the operating characteristics are summarized in seven sections. Section 1 gives a brief description of the true dose response and the need for insertion or not. Section 2 and 3 provide the average (standard deviation) of the posterior means across 1,000 simulated trials for the regression parameters. Section 4 summarized the model selection. Section 5 presents the true toxicity and efficacy probability at each dose combination and Section 6 presents the detailed operating characteristics in terms of patient allocation and dose selection. Section 7 demonstrates the results of the design in Cai et al. (2014).

Scenario 1 % early stop: 0					
1. Quadratic; Efficacy peaks in the left side of MTD; Insertion is needed.					

2. Toxicity Parameter Estimation, True $(\alpha_0, \alpha_1, \alpha_2) = (-1.58, 1.4, 1.06)$

			$\hat{\alpha_0}$	$\hat{\alpha_1}$	$\hat{\alpha_2}$
			-1.546(0.315)	1.161(0.874)	1.041(0.823)

3. Efficacy Parameter Estimation,True $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1.45, 0.7, 0.59, -9.8, -7.29)$

	$\hat{\beta_0}$	$\hat{\beta_1}$	$\hat{\beta_2}$	$\hat{\beta_3}$	$\hat{\beta_4}$
Linear	NaN(NaN)	NaN(NaN)	NaN(NaN)	.	.
Quadratic of agent A	1.475(0.549)	0.464(1.661)	0.917(1.278)	-22.318(8.56)	.
Quadratic of agent B	1.195(0.439)	1.459(1.368)	-0.256(1.589)	.	-23.177(8.323)
Quadratic of both agents	1.535(0.523)	0.527(1.56)	0.437(1.544)	-13.758(6.928)	-11.861(6.021)

4. Model Selection percentages

Linear	Quadratic of agent A	Quadratic of agent B	Quadratic of both agents
0%	5.8%	3.7%	90.5%

5. True Toxicity and Efficacy Probability: Tox/Eff

Dose Level (A\B)	0.05	0.35	0.65	1
0.1	0.05/0.01	0.07/0.1	0.1/0.14	0.16/0.02
0.3	0.07/0.12	0.11/0.56	0.15/0.64	0.21/0.17
0.6	0.13/0.19	0.18/0.7	0.24/0.76	0.33/0.27
0.9	0.21/0.01	0.28/0.1	0.37/0.14	0.48/0.02

6. Dose Allocation and Selection Percentages %pts/%sel
BODC=(0.441,0.476), BMS:%ins=59.3%
Dose inserted: 0.43(0.111), 0.448(0.142); Dose selected: 0.456(0.109), 0.464(0.118)

Dose Level(A\B)	0.05	0.35	0.65	1	Dose inserted	Dose selected
0.1	0.031/0	0.011/0.003	0.009/0.001	0.001/0	0.257/0.559	0.43/
0.3	0.01/0.001	0.162/0.086	0.115/0.098	0.006/0.001		
0.6	0.017/0.01	0.189/0.184	0.09/0.054	0.02/0.002		
0.9	0.001/0	0.004/0.001	0.019/0	0.058/0		

7. Dose Allocation and Selection Percentages for Cai's Algorithm (% early stop: 0): %pts/%sel

Dose Level(A\B)	0.05	0.35	0.65	1
0.1	0.031/0	0.008/0	0.009/0.004	0.001/0
0.3	0.009/0	0.219/0.185	0.225/0.271	0.004/0.001
0.6	0.013/0.002	0.299/0.399	0.116/0.136	0.016/0.002
0.9	0.001/0	0.002/0	0.016/0	0.031/0

Scenario 2 % early stop: 0					
1. Quadratic; Efficacy peaks in the left side of MTD; Insertion is not needed					

2. Toxicity Parameter Estimation, True $(\alpha_0, \alpha_1, \alpha_2) = (-1.38, 1.38, 1.01)$

			$\hat{\alpha_0}$	$\hat{\alpha_1}$	$\hat{\alpha_2}$
			-1.392(0.285)	1.13(0.913)	1.031(0.839)

3. Efficacy Parameter Estimation,True $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1.44, -1.03, -0.5, -9.52, -6.64)$

	$\hat{\beta_0}$	$\hat{\beta_1}$	$\hat{\beta_2}$	$\hat{\beta_3}$	$\hat{\beta_4}$
Linear	-0.952(0.53)	0.994(1.335)	-0.432(0.896)	.	.
Quadratic of agent A	1.794(0.57)	-0.88(2.51)	1.359(1.516)	-34.834(13.359)	.
Quadratic of agent B	2.413(0.674)	1.46(1.583)	-0.688(2.026)	.	-40.72(14.957)

Table 2 (continued)

Quadratic of both agents	1.776(0.509)	-0.332(1.928)	-0.275(1.686)	-15.634(8.089)	-11.464(5.703)

4. Model Selection percentages

Linear	Quadratic of agent A	Quadratic of agent B	Quadratic of both agents
0.1%	7.4%	3.2%	89.3%

5. True Toxicity and Efficacy Probability: Tox/Eff

Dose Level (A\B)	0.1	0.45	0.75	1
0.1	0.05/0.02	0.08/0.14	0.12/0.1	0.15/0.02
0.45	0.1/0.32	0.15/0.79	0.21/0.72	0.27/0.3
0.7	0.16/0.18	0.23/0.64	0.3/0.55	0.38/0.17
0.9	0.22/0.02	0.31/0.14	0.39/0.1	0.47/0.02

6. Dose Allocation and Selection Percentages %pts/%sel

BODC=(0.442,0.477), BMS:%ins=15.6%

Dose inserted: 0.533(0.146), 0.575(0.142); Dose selected: 0.454(0.051), 0.466(0.093)

Dose Level(A\B)	0.1	0.45	0.75	1	Dose inserted	Dose selected
0.1	0.031/0.001	0.022/0.006	0.003/0.002	0.001/0	0.042/0.063	0.611/
0.45	0.046/0.021	0.633/0.835	0.051/0.054	0.004/0		
0.7	0.003/0	0.032/0.018	0.032/0	0.019/0		
0.9	0/0	0.003/0	0.021/0	0.057/0		

7. Dose Allocation and Selection Percentages for Cai's Algorithm (% early stop: 0): %pts/%sel

Dose Level(A\B)	0.1	0.45	0.75	1
0.1	0.032/0.001	0.017/0.002	0.005/0.001	0.002/0
0.45	0.042/0.023	0.706/0.91	0.053/0.05	0.003/0
0.7	0.007/0	0.034/0.012	0.032/0.001	0.015/0
0.9	0.002/0	0.003/0	0.017/0	0.03/0

Scenario 3 % early stop: 0.001

1. Quadratic; Efficacy peaks in the right side of MTD; Only a few dose combos are tolerable; Insertion is needed.

2. Toxicity Parameter Estimation, True $(\alpha_0, \alpha_1, \alpha_2) = (-0.25, 2.22, 2.05)$

	$\hat{\alpha_0}$	$\hat{\alpha_1}$	$\hat{\alpha_2}$
	-0.208(0.346)	1.994(0.956)	1.792(0.794)

3. Efficacy Parameter Estimation, True $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1.69, 2.87, 0.14, -8.24, -3.75)$

	$\hat{\beta_0}$	$\hat{\beta_1}$	$\hat{\beta_2}$	$\hat{\beta_3}$	$\hat{\beta_4}$
Linear	1.025(0.797)	6.108(2.633)	0.137(0.842)	.	.
Quadratic of agent A	1.126(0.538)	2.209(1.189)	1.122(0.856)	-9.821(4.077)	.
Quadratic of agent B	1.394(0.635)	4.565(1.637)	-0.177(1.244)	.	-8.127(3.87)
Quadratic of both agents	1.891(0.62)	2.821(1.347)	0.212(0.966)	-9.364(3.768)	-5.297(2.104)

4. Model Selection percentages

Linear	Quadratic of agent A	Quadratic of agent B	Quadratic of both agents
0.7%	24%	2.2%	73%

5. True Toxicity and Efficacy Probability: Tox/Eff

Dose Level (A\B)	0.2	0.5	0.75	1
0.05	0.07/0.03	0.15/0.1	0.27/0.1	0.44/0.04
0.25	0.11/0.27	0.24/0.58	0.4/0.58	0.59/0.34
0.55	0.24/0.62	0.44/0.86	0.62/0.86	0.78/0.69
0.9	0.48/0.24	0.69/0.54	0.82/0.54	0.91/0.3

6. Dose Allocation and Selection Percentages %pts/%sel

BODC=(0.41,0.332), BMS:%ins=29.6%

Dose inserted: 0.323(0.184), 0.286(0.148); Dose selected: 0.38(0.145), 0.336(0.151)

Dose Level(A\B)	0.2	0.5	0.75	1	Dose inserted	Dose selected
0.05	0.032/0.005	0.034/0.01	0.025/0.002	0.011/0	0.075/0.189	0.408/
0.25	0.115/0.085	0.27/0.356	0.034/0.011	0.018/0.001		
0.55	0.222/0.323	0.029/0.016	0.031/0	0.023/0		
0.9	0.013/0.002	0.016/0	0.023/0	0.029/0		

7. Dose Allocation and Selection Percentages for Cai's Algorithm (% early stop: 0): %pts/%sel

Dose Level(A\B)	0.2	0.5	0.75	1
0.05	0.032/0.008	0.034/0.005	0.023/0.008	0.012/0

Table 2 (continued)

0.25	0.117/0.079	0.285/0.367	0.037/0.023	0.009/0
0.55	0.32/0.499	0.033/0.01	0.031/0	0.018/0
0.9	0.013/0.001	0.007/0	0.009/0	0.023/0

Scenario 4 % early stop: 0.086

1. Quadratic; Efficacy peaks in the right side of MTD; All dose are overly toxic; Insertion is needed.

2. Toxicity Parameter Estimation, True $(\alpha_0, \alpha_1, \alpha_2) = (0.51, 1.03, 0.67)$

	$\hat{\alpha_0}$	$\hat{\alpha_1}$	$\hat{\alpha_2}$
	0.639(0.475)	1.039(0.592)	0.701(0.488)

3. Efficacy Parameter Estimation,True $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (0.69, -3.2, -1.76, -4.8, -2.93)$

	$\hat{\beta_0}$	$\hat{\beta_1}$	$\hat{\beta_2}$	$\hat{\beta_3}$	$\hat{\beta_4}$
Linear	1.209(0.825)	1.251(1.059)	0.453(0.911)	.	.
Quadratic of agent A	0.487(0.622)	-3.273(1.669)	0.286(0.786)	-5.286(1.899)	.
Quadratic of agent B	0.857(0.699)	0.997(0.911)	-2.232(1.488)	.	-3.727(1.782)
Quadratic of both agents	0.724(0.654)	-2.92(1.58)	-1.58(1.289)	-4.749(1.756)	-2.704(1.436)

4. Model Selection percentages

Linear	Quadratic of agent A	Quadratic of agent B	Quadratic of both agents
1.5%	18.3%	6.8%	68.7%

5. True Toxicity and Efficacy Probability: Tox/Eff

Dose Level (A\B)	0.4	0.6	0.8	1
0.4	0.38/0.72	0.45/0.76	0.51/0.62	0.57/0.25
0.6	0.48/0.76	0.55/0.8	0.61/0.66	0.67/0.29
0.8	0.58/0.48	0.64/0.53	0.7/0.36	0.75/0.11
1	0.67/0.06	0.73/0.07	0.78/0.04	0.82/0.01

6. Dose Allocation and Selection Percentages %pts/%sel

BODC=(0.304,0.262), BMS:%ins=92.5%

Dose inserted: 0.354(0.219), 0.354(0.231); Dose selected: 0.242(0.073), 0.245(0.09)

Dose Level(A\B)	0.4	0.6	0.8	1	Dose inserted	Dose selected
0.4	0.081/0.033	0.025/0	0.018/0	0.008/0	0.678/0.967	0.208/
0.6	0.023/0	0.032/0	0.015/0	0.005/0		
0.8	0.014/0	0.015/0	0.026/0	0.013/0		
1	0.006/0	0.004/0	0.013/0	0.024/0		

7. Dose Allocation and Selection Percentages for Cai's Algorithm (% early stop: 0.603): %pts/%sel

Dose Level(A\B)	0.4	0.6	0.8	1
0.4	0.568/1	0.066/0	0.033/0	0.018/0
0.6	0.064/0	0.04/0	0.03/0	0.009/0
0.8	0.034/0	0.028/0	0.034/0	0.016/0
1	0.018/0	0.007/0	0.011/0	0.023/0

Scenario 5 % early stop: 0

1. Quadratic; All dose are safe; The dose range only covers the bottom-left corner of the efficacy-dose curve. Insertion is needed.

2. Toxicity Parameter Estimation, True $(\alpha_0, \alpha_1, \alpha_2) = (-2.8, 0.77, 0.52)$

	$\hat{\alpha_0}$	$\hat{\alpha_1}$	$\hat{\alpha_2}$
	-3.115(0.808)	0.924(0.976)	0.869(0.932)

3. Efficacy Parameter Estimation,True $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (-2.86, 3.56, 2.22, -1.2, -0.73)$

	$\hat{\beta_0}$	$\hat{\beta_1}$	$\hat{\beta_2}$	$\hat{\beta_3}$	$\hat{\beta_4}$
Linear	-3.671(1.283)	3.062(1.935)	2.93(1.762)	.	.
Quadratic of agent A	-4.262(1.622)	3.814(2.555)	3.031(1.847)	0.318(2.519)	.
Quadratic of agent B	-4.078(1.506)	3.599(2.145)	3.006(2.131)	.	-0.05(2.547)
Quadratic of both agents	-4.597(1.578)	4.238(2.504)	3.108(2.04)	0.372(2.612)	-0.363(2.384)

4. Model Selection percentages

Linear	Quadratic of agent A	Quadratic of agent B	Quadratic of both agents
25.4%	23.9%	23.2%	27.5%

5. True Toxicity and Efficacy Probability: Tox/Eff

Dose Level (A\B)	0.05	0.15	0.25	0.35

Table 2 (continued)

0.05	0.03/0	0.03/0	0.04/0.01	0.05/0.01
0.15	0.04/0.01	0.05/0.02	0.05/0.04	0.07/0.07
0.25	0.05/0.02	0.06/0.06	0.07/0.14	0.09/0.24
0.35	0.07/0.06	0.08/0.16	0.1/0.31	0.11/0.46

6. Dose Allocation and Selection Percentages %pts/%sel

BODC=(0.474,0.472), BMS:%ins=11.1%

Dose inserted: 0.443(0.205), 0.243(0.166); Dose selected: 0.358(0.053), 0.345(0.041)

Dose Level(A\B)	0.05	0.15	0.25	0.35	Dose inserted	Dose selected
0.05	0.031/0	0.007/0	0.008/0	0.016/0.001	0.025/0.051	0.625/
0.15	0.007/0	0.031/0	0.016/0	0.021/0.002		
0.25	0.008/0	0.016/0	0.033/0.001	0.045/0.009		
0.35	0.022/0	0.028/0.002	0.068/0.029	0.618/0.905		

7. Dose Allocation and Selection Percentages for Cai's Algorithm (% early stop: 0): %pts/%sel

Dose Level(A\B)	0.05	0.15	0.25	0.35
0.05	0.032/0.001	0.014/0	0.016/0.001	0.024/0.001
0.15	0.013/0	0.032/0	0.024/0	0.028/0.003
0.25	0.017/0.001	0.025/0	0.04/0.006	0.071/0.051
0.35	0.035/0.006	0.045/0.023	0.139/0.162	0.447/0.745

Scenario 6 % early stop: 0.001

1. Description

2. Toxicity Parameter Estimation, True $(\alpha_0, \alpha_1, \alpha_2) = (-1.46, 0.77, 1.39)$

	$\hat{\alpha_0}$	$\hat{\alpha_1}$	$\hat{\alpha_2}$
	-1.427(0.379)	0.86(0.703)	1.164(0.71)

3. Efficacy Parameter Estimation,True $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (0.3, 3.1, 1.55, ., .)$

	$\hat{\beta_0}$	$\hat{\beta_1}$	$\hat{\beta_2}$	$\hat{\beta_3}$	$\hat{\beta_4}$
Linear	0.237(0.434)	3.226(1.123)	1.696(1.033)	.	.
Quadratic of agent A	-0.02(0.683)	7.546(2.633)	1.894(1.172)	-8.485(3.304)	.
Quadratic of agent B	0.41(0.581)	3.936(1.426)	3.709(1.765)	.	-7.52(2.633)
Quadratic of both agents	-0.294(0.827)	10.201(3.523)	5.628(2.276)	-10.653(4.515)	-8.081(3.232)

4. Model Selection percentages

Linear	Quadratic of agent A	Quadratic of agent B	Quadratic of both agents
70.3%	13.3%	12.6%	3.7%

5. True Toxicity and Efficacy Probability: Tox/Eff

Dose Level (A\B)	0.1	0.4	0.7	1
0.1	0.06/0.08	0.1/0.14	0.16/0.23	0.25/0.35
0.4	0.08/0.23	0.13/0.35	0.21/0.5	0.31/0.65
0.7	0.11/0.5	0.17/0.65	0.26/0.77	0.38/0.86
1	0.14/0.77	0.22/0.86	0.32/0.92	0.45/0.95

6. Dose Allocation and Selection Percentages %pts/%sel

BODC=(1.384,0.05), BMS:%ins=32.2%

Dose inserted: 1.174(0.511), 0.256(0.25); Dose selected: 1.002(0.367), 0.341(0.316)

Dose Level(A\B)	0.1	0.4	0.7	1	Dose inserted	Dose selected
0.1	0.031/0.003	0.002/0	0.009/0.002	0.015/0.012	0.079/0.202	0.287/
0.4	0.003/0.002	0.04/0.001	0.06/0.02	0.038/0.053		
0.7	0.036/0.023	0.118/0.052	0.137/0.064	0.028/0.034		
1	0.184/0.32	0.128/0.168	0.033/0.021	0.059/0.023		

Table 3: Time duration comparison between the adaptive cohort division algorithm and the single cohort under the AAA. Entries are simulated trial duration is days for two-agent trials of sample size 96 patients.

Scenario	Algorithm	Mean(Sd)	Min~Max	p-value
Scenario 1	cohort divison	927.53(185.32)	$346 \sim 1076$	$< 0.0001^{**}$
	single cohort	1037.89(59.9)	$558 \sim 1077$	
Scenario 2	cohort divison	914.34(203.19)	$367 \sim 1075$	$< 0.0001^{**}$
	single cohort	1030.54(76.2)	$495 \sim 1074$	
Scenario 3	cohort divison	622.82(172.78)	$26 \sim 1062$	$< 0.0001^{**}$
	single cohort	1017.16(84.52)	$26 \sim 1071$	
Scenario 4	cohort divison	898.55(234.54)	$6 \sim 1067$	$< 0.0001^{**}$
	single cohort	944.86(249.64)	$6 \sim 1073$	
Scenario 5	cohort divison	1045.27(61.43)	$522 \sim 1081$	0.2376
	single cohort	1050.79(46.31)	$599 \sim 1077$	
Scenario 6	cohort divison	910.56(205.81)	$28 \sim 1073$	$< 0.0001^{**}$
	single cohort	1025.39(80.2)	$28 \sim 1072$	
Scenario 7	cohort divison	986.98(151.56)	$357 \sim 1076$	$< 0.0001^{**}$
	single cohort	1036.87(68.52)	$522 \sim 1077$	
Scenario 8	cohort divison	1011.31(127.67)	$362 \sim 1078$	0.03^{*}
	single cohort	1034.41(77.74)	$523 \sim 1079$	
Scenario 9	cohort divison	962.51(144.38)	$350 \sim 1079$	$< 0.0001^{**}$
	single cohort	1043.68(48.29)	$534 \sim 1076$	
Scenario 10	cohort divison	977.7(145.86)	$23 \sim 1075$	$< 0.0001^{**}$
	single cohort	1037.96(68.01)	$23 \sim 1077$	