Dose Finding for Drug Combination in Early Cancer Phase I Trials using Conditional Escalation with Overdose Control

Mourad Tighiouart, André Rogatko, Steven Piantadosi

Biostatistics and Bioinformatics Research Center, Cedars-Sinai Medical Center
8700 Beverly Blvd, Los Angeles, CA 9004

Abstract

We present Bayesian adaptive designs for dose finding of a combination of two drugs in cancer phase I clinical trials. The goal is to estimate the maximum tolerated dose (MTD) as a curve in the two-dimensional Cartesian plane. Parametric models are used to describe the relationship between the doses of the two agents and the probability of dose limiting toxicity. We investigate different reparameterizations in terms of parameters clinicians can easily interpret. Trial design proceeds using univariate escalation with overdose control, where at each stage of the trial, we seek a dose of one agent using the current posterior distribution of the MTD of this agent given the current dose of the other agent. At the end of the trial, an estimate of the MTD curve is proposed as a function of Bayes estimates of the model parameters. We evaluate design operating characteristics in terms of safety of the trial design and percent of dose recommendation at dose combination neighborhoods around the true MTD curve. We also examine the performance of the approach under model misspecifications for the true dose-toxicity relationship.

Key Words: Cancer Phase I trials; Maximum tolerated dose curve; Escalation with overdose control; Drug combination; Dose limiting toxicity; Continuous dose.

1. Introduction

Cancer phase I clinical trials are small studies of investigational new agents or combination of existing cytotoxic and/or biologic agents. These trials typically enroll patients with advanced form of cancer who have exhausted all standard therapy [1]. The primary objective of such studies is to estimate a maximum tolerable dose (MTD) of a new drug or combinations of drugs for future efficacy evaluation in phase II/III trials. A cancer phase I trial design is an algorithm of dose assignment to successive cohorts of patients in order to estimate the MTD while minimizing the number of patients experiencing severe dose related side effects. A review of single agent dose-finding designs can be found in [2-4]. Combining several drugs can help reduce tumor resistance to chemotherapy by targeting different signaling pathways simultaneously and improve tumor response when using additive or synergistic drugs [5]. Although the majority of phase I trials use drug combinations of several agents, most of them are designed to estimate the MTD of one drug for fixed dose levels of the other drugs. This approach may provide a single safe dose for the combination but it may be suboptimal in terms of therapeutic effects. Trials where the dose levels of at least two agents are allowed to vary yield more than one MTD, or even an infinite number of MTDs in the case of continuous does levels. Estimating the resulting set of MTDs by designing a safe trial is the main

goal of phase I trials with dose combinations of several agents. Let A_i, $i = 1,\ldots,k$ be k drugs and $S_i \subset R^+$ be the set of all possible doses of drug A_i. Denote by $\boldsymbol{x} = (x_1,\ldots, x_k)$ a dose combination of the k drugs and $S = S_1 \times \ldots \times S_k$. Consider a dose-toxicity model

$$\text{Prob}\left(\text{DLT} \mid \text{dose} = \boldsymbol{x}\right) = F(\boldsymbol{x}, \boldsymbol{\xi}), \tag{1.1}$$

where F is a known link function and $\boldsymbol{\xi} \in R^d$ is an unknown parameter. The MTD is defined as the set C of dose combinations $\boldsymbol{x}$ such that the probability of DLT for a patient given dose $\boldsymbol{x}$ equals to a target probability of DLT θ:

$$C = \left\{\boldsymbol{x} \in S : F(\boldsymbol{x}, \boldsymbol{\xi}) = \theta\right\}. \tag{1.2}$$

For $k = 2$, a number of designs for estimating C have been proposed in the past decade, see [6-14]. By design, these methods do not apply to the case of continuous dose levels of the two agents. Although the method of [6] can be used for continuous agents in the second stage of the trial, the first stage of the design does require a discretization of the dose combinations along a diagonal. Except for [6], the methods recommend a single dose combination as the MTD. Furthermore, it is not clear how these methods perform in the presence of a large number of dose combinations in relation to the sample size in the trial, especially if dose escalation by more than one level in either direction is not allowed. In this manuscript, we extend the design described by Tighiouart et al. [15] by allowing the true MTD curve to lie outside the range of doses available in the trial and introduce a new algorithm for dose escalation by treating cohorts of two patients receiving different dose combinations simultaneously. Doses are determined according to escalation with overdose control (EWOC) principle [16-18].

2. Model

2.1 Dose-Toxicity Model

Consider the problem of identifying a tolerable dose(s) of the combination of two cytotoxic agents *A* and *B*. We consider the dose-toxicity model of the form

$$\text{Prob}(Z = 1 \mid x, y) = F(\mu + \beta x + \gamma y + \eta xy), \tag{2.1}$$

where Z is the indicator of DLT, $Z = 1$ if a patient given the dose combination (x,y) exhibits DLT within one cycle of therapy, and $Z = 0$ otherwise, $x \in [X_{\min}, X_{\max}]$ is the dose level of agent *A*, $y \in [Y_{\min}, Y_{\max}]$ is the dose level of agent *B*, and F is a known cumulative distribution function. Suppose that the doses of agents *A* and *B* are continuous and standardized to be in the interval [0, 1].

We will assume that that the probability of DLT increases with the dose of any one of the agents when the other one is held constant. A sufficient condition for this property to hold is to assume $\beta > 0$ and $\gamma > 0$ and the interaction term η is nonnegative. The MTD is defined as any dose combination (x^*, y^*) such that

$$\text{Prob}(Z = 1 \mid x^*, y^*) = \theta. \tag{2.2}$$

The target probability of DLT θ is set relatively high when the DLT is a reversible or non-fatal condition, and low when it is life threatening. Using (2.1) and (2.2), the MTD is

$$C = \left\{ (x^*, y^*) \in [0,1]^2 : y^* = \frac{F^{-1}(\theta) - \mu - \beta x^*}{\gamma + \eta x^*} \right\}. \quad (2.3)$$

We reparameterize model (2.1) in terms of parameters clinicians can easily interpret. We describe two reparameterizations that accommodate prior or lack of knowledge about each drug when used as single agent.

2.1.1 Reparameterization 1

Model (2.1) is reparameterized in terms of $\Gamma_{A|0}$, the MTD of drug A when the level of drug B is $Y_{\min}$, $\Gamma_{B|0}$, the MTD of drug B when the level of drug A is $X_{\min}$, ρ_{00}, the probability of DLT at the minimum available doses of agents A and B, and the interaction parameter η. If MTD estimates of single agent trials using A and B are available, then these estimates can be used to approximate the prior mean and variances of the parameters $\Gamma_{A|0}$, and $\Gamma_{B|0}$. We will assume that $0 \leq \Gamma_{A|0}, \Gamma_{B|0} \leq 1$, i.e., the MTD of each agent when the other one is held at its minimum available dose in the trial is within the range of available doses in the trial. This is the reparameterization used in Tighiouart et al. [15]. It follows that

$$\begin{cases} \mu = F^{-1}(\rho_{00}) \\ \beta = (F^{-1}(\theta) - F^{-1}(\rho_{00})) / \Gamma_{A|0} \\ \gamma = (F^{-1}(\theta) - F^{-1}(\rho_{00})) / \Gamma_{B|0} \end{cases}. \quad (2.4)$$

The MTD (2.3) can be expressed in terms of these new parameters as

$$C = \left\{ (x^*, y^*) : 0 \leq x^*, y^* \leq 1, y^* = \frac{\left(F^{-1}(\theta) - F^{-1}(\rho_{00})\right)\left(1 - x^* / \Gamma_{A|0}\right)}{\left(\left(F^{-1}(\theta) - F^{-1}(\rho_{00})\right) / \Gamma_{B|0}\right) + \eta x^*} \right\}. \quad (2.5)$$

2.1.2 Reparameterization 2

We reparameterize model (2.1) in terms of ρ_{10}, the probability of DLT when the levels of drugs A and B are 1 and 0, respectively, ρ_{01}, the probability of DLT when the levels of drugs A and B are 0 and 1, respectively, ρ_{00}, the probability of DLT when the levels of drugs A and B are both 0, and the interaction parameter η. This reparametrization relaxes the conditions that the MTDs $\Gamma_{A|0}$ and $\Gamma_{B|0}$ described in Section 2.1.1 be bounded within the range of doses available in the trial. It can be shown that

$$\begin{cases} \mu = F^{-1}(\rho_{00}) \\ \beta = (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00})). \\ \gamma = (F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})) \end{cases} \quad (2.6)$$

The MTD (2.3) becomes

$$C=\left\{(x^*,y^*): y^*=\frac{\left(F^{-1}(\theta)-F^{-1}(\rho_{00})\right)-\left(F^{-1}(\rho_{10})-F^{-1}(\rho_{00})\right)x^*}{\left(F^{-1}(\rho_{01})-F^{-1}(\rho_{00})\right)+\eta x^*}\right\}. \qquad (2.7)$$

Figure 1(a) shows some MTD curves under model reparameterization 1 when $\rho_{00} = 0.05$, $\Gamma_{A|0} = 0.7$, $\Gamma_{B|0} = 0.7$, and three values for the interaction coefficient $\eta = 0$, 20, 60. Figure 1(b) shows MTD curves under reparameterization 2 when $\rho_{00} = 0.05$, $\rho_{01} = \rho_{10} = 3\times10^{-6}$, and three values for the interaction coefficient $\eta = 20$, 40, 60.

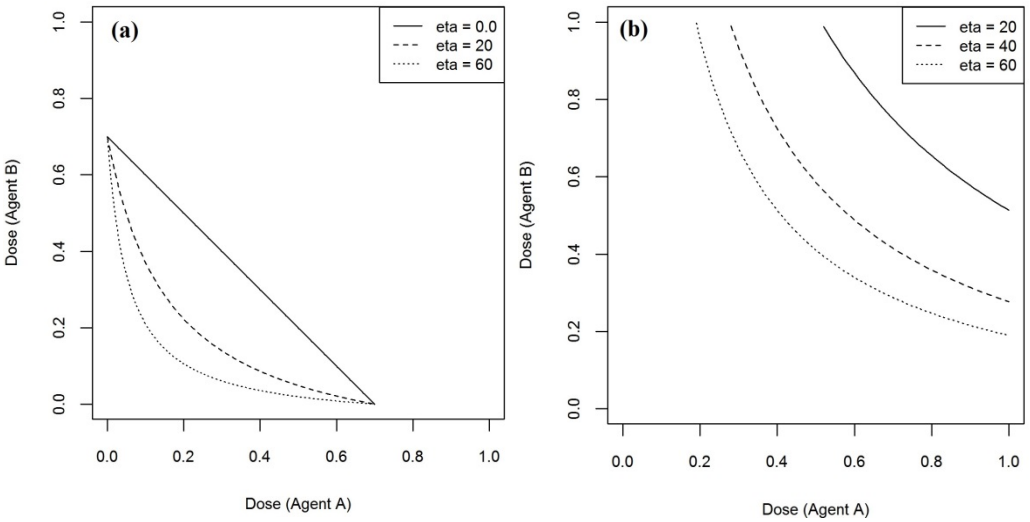

Figure 1. MTD curves for selected values of the interaction coefficient η. Target probability of DLT $\theta = 0.33$ and logistic link $F(u) = (1 + e^{-u})^{-1}$. (a) Reparameterization 1, $\rho_{00} = 0.05$, $\Gamma_{A|0} = 0.7$, $\Gamma_{B|0} = 0.7$, (b) Raparameterization 2, $\rho_{00} = 10^{-7}$, $\rho_{01} = \rho_{10} = 3\times10^{-6}$.

In both cases, the target probability of DLT is $\theta = 0.33$ and the link function is the logistic $F(u) = (1 + e^{-u})^{-1}$. Figure 1(b) illustrates cases where each drug when used as single agent is very safe within the range of doses available in the trial. Let $D_n = \{(x_i,y_i,z_i), i = 1, \ldots, n\}$ be the data after enrolling n patients in the trial. The likelihood function under reparameterization 1 is

$$\begin{aligned} L(\rho_{00},\Gamma_{A|0},\Gamma_{B|0},\eta \mid D_n) = \prod_{i=1}^{n}\left(G(\rho_{00},\Gamma_{A|0},\Gamma_{B|0},\eta;x_i,y_i)\right)^{z_i} \\ \times\left(1-G(\rho_{00},\Gamma_{A|0},\Gamma_{B|0},\eta;x_i,y_i)\right)^{1-z_i}, \end{aligned} \qquad (2.8)$$

where

$$\begin{aligned} &G(\rho_{00},\Gamma_{A|0},\Gamma_{B|0},\eta;x_i,y_i) \\ &= F\left(F^{-1}(\rho_{00})+\frac{F^{-1}(\theta)-F^{-1}(\rho_{00})}{\Gamma_{A|0}}x_i+\frac{F^{-1}(\theta)-F^{-1}(\rho_{00})}{\Gamma_{B|0}}y_i+\eta x_i y_i\right). \end{aligned} \qquad (2.9)$$

Under model reparameterization 2, the likelihood is

$$L(\rho_{00},\rho_{01},\rho_{10},\eta \mid D_n) = \prod_{i=1}^{n} \left(G(\rho_{00},\rho_{01},\rho_{10},\eta;x_i,y_i)\right)^{z_i} \times \left(1-G(\rho_{00},\rho_{01},\rho_{10},\eta;x_i,y_i)\right)^{1-z_i}, \quad (2.10)$$

where

$$G(\rho_{00},\rho_{01},\rho_{10},\eta;x_i,y_i) = F\left(F^{-1}(\rho_{00}) + \left(F^{-1}(\rho_{10}) - F^{-1}(\rho_{00})\right)x_i + \left(F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})\right)y_i + \eta x_i y_i\right). \quad (2.11)$$

2.2 Prior and Posterior Distributions

Under model reparameterization 1, (2.4) implies that $0 < \rho_{00} < \theta$ since $\beta > 0$. We consider the priors $\rho_0/\theta \sim$ beta(a_1, b_1), $\Gamma_{A|0} \sim$ beta(a_2, b_2), $\Gamma_{B|0} \sim$ beta(a_3, b_3), $\eta \sim$ gamma(a, b) with mean $E(\eta) = a / b$ and variance $Var(\eta) = a / b^2$. Vague priors for these parameters are achieved by taking $a_j = b_j = 1$, $j = 1, 2, 3$. A vague prior for η is then achieved by setting $E(\eta) = 8\left(F^{-1}(\theta) - F^{-1}(E(\rho_{00}))\right) \Big/ E(\Gamma_{A|0})E(\Gamma_{B|0})$. A large variance is selected for η, see [15] for the rationale behind this choice. Using Bayes rule, the posterior distribution of the model parameters is proportional to the product of the likelihood and prior distribution

$$\pi(\rho_{00},\Gamma_{A|0},\Gamma_{B|0},\eta \mid D_n) \propto \prod_{i=1}^{n} \left(F(\rho_{00},\Gamma_{A|0},\Gamma_{B|0},\eta;x_i,y_i)\right)^{z_i} \times \left(1-F(\rho_{00},\Gamma_{A|0},\Gamma_{B|0},\eta;x_i,y_i)\right)^{1-z_i} \rho_{00}^{\,a_1-1}\left(1-\rho_{00}\right)^{b_1-1} \times \left(\Gamma_{A|0}\right)^{a_2-1}\left(1-\Gamma_{A|0}\right)^{b_2-1}\left(\Gamma_{B|0}\right)^{a_3-1}\left(1-\Gamma_{B|0}\right)^{a_3-1}\eta^{a-1}e^{-b\eta}. \quad (2.12)$$

Features of this posterior distribution are estimated using WinBUGS[19] and JAGS. For model reparameterization 2, we note that (2.6) implies that $0 < \rho_{00} < \min(\rho_{01}, \rho_{10})$ since β and γ are positive. We assume that ρ_{01}, ρ_{10}, η are independent *a priori* with $\rho_{01} \sim$ beta(a_1, b_1), $\rho_{10} \sim$ beta(a_2, b_2), and conditional on (ρ_{01}, ρ_{10}), $\rho_{00} / \min(\rho_{01}, \rho_{10}) \sim$ beta(a_3, b_3). Vague priors for these parameters are again achieved by taking $a_j = b_j = 1$, $j = 1, 2, 3$. The hyperparameters for η are taken as in reparameterization 1.

2.3 Trial Design

We describe two algorithms for dose allocation. Both of these use the EWOC principle where at each stage of the trial, the posterior probability of overdosing a future patient is bounded by a feasibility bound α. The first one is a review of the dose allocation scheme described in [15] and the second algorithm enrolls cohorts of two patients simultaneously receiving different dose combinations.

2.3.1 Algorithm 1

1. The first patient receives dose (x_1, y_1) = (0, 0) and suppose the patient has no DLT, $z_1 = 0$.
2. Fix $x_1 = 0$, and calculate the posterior distribution of the MTD of agent *B*, given that the level of agent *A* is $x_1 = 0$, $\pi(\Gamma_{B|A=0} \mid D_1)$. The dose for patient 2 is (x_2, y_2) where $x_2 = x_1$ and y_2 is the α-th percentile of this posterior distribution.

3. Fix y_2, and calculate the posterior distribution of the MTD of agent *A*, given that the level of agent *B* is y_2, $\pi(\Gamma_{A|B=y2} \mid D_1)$. The dose for patient 3 is (x_3, y_3) where x_3 is the α-th percentile of this posterior distribution and $y_3 = y_2$.

In general, when we move from dose (x_i, y_i) to (x_{i+1}, y_{i+1}), either $x_i = x_{i+1}$ or $y_i = y_{i+1}$. Specifically, if i is even, then $x_{i+1} = \Pi^{-1}_{\Gamma_{A|B=y_i}}(\alpha \mid D_i)$ and $y_{i+1} = y_i$. If i is odd, then $x_{i+1} = x_i$ and $y_{i+1} = \Pi^{-1}_{\Gamma_{B|A=x_i}}(\alpha \mid D_i)$. Here, $\Pi^{-1}_{\Gamma_{A|B=y_i}}(\alpha \mid D_i)$ is the inverse cdf of the posterior distribution $\pi(\Gamma_{A|B=y_i} \mid D_i)$.

4. Repeat step 3 by fixing either dose x_i or y_i, depending on whether i is even or odd, until n patients are enrolled to the trial subject to the following stopping rule.

Stopping rule: We stop enrollment to the trial if $P(P(\text{DLT}|(x,y) = (0,0)) \geq \theta+\delta_1 \mid \text{data}) > \delta_2$, i.e. if the posterior probability that the probability of DLT at the minimum available dose combination in the trial exceeds the target probability of DLT is high. δ_1 and δ_2 are design parameters chosen to achieve desirable model operating characteristics.

2.3.2 Algorithm 2

1. Each patient in the first cohort of two patients receives the same dose combination $(x_1, y_1) = (x_2, y_2) = (0, 0)$. Let $D_2 = \{(x_1, y_1, \delta_1), (x_2, y_2, \delta_2)\}$.
2. In the second cohort of two patients, patient 3 receives dose (x_3, y_1) and patient 4 receives dose (x_2, y_4), where x_3 is the α-th percentile of $\pi(\Gamma_{A|B=y1} \mid D_2)$ and y_4 is the α-th percentile of $\pi(\Gamma_{B|A=x2} \mid D_2)$.
3. In the i-th cohort of two patients, if i is even, then patient $2i-1$ receives dose (x_{2i-1}, y_{2i-3}), patient $2i$ receives dose (x_{2i-2}, y_{2i}), where $x_{2i-1} = \Pi^{-1}_{\Gamma_{A|B=y_{2i-3}}}(\alpha \mid D_{2i-2}), y_{2i} = \Pi^{-1}_{\Gamma_{B|A=x_{2i-2}}}(\alpha \mid D_{2i-2})$. If i is odd, then patient $2i-1$ receives dose (x_{2i-3}, y_{2i-1}), patient $2i$ receives dose (x_{2i}, y_{2i-2}), where $y_{2i-1} = \Pi^{-1}_{\Gamma_{B|A=x_{2i-3}}}(\alpha \mid D_{2i-2}), x_{2i} = \Pi^{-1}_{\Gamma_{A|B=y_{2i-2}}}(\alpha \mid D_{2i-2})$.
4. Repeat step 3 until n patients are enrolled to the trial subject to the stopping rule described in Algorithm 1.

At the end of the trial, we estimate the MTD curve using Bayes estimates of the parameters defining this curve. For example, using reparameterization 2, an estimate of the MTD curve is obtained using (2.7) as

$$\hat{C} = \left\{ (x^*, y^*) : y^* = \frac{\left(F^{-1}(\theta) - F^{-1}(\hat{\rho}_{00})\right) - \left(F^{-1}(\hat{\rho}_{10}) - F^{-1}(\hat{\rho}_{00})\right)x^*}{\left(F^{-1}(\hat{\rho}_{01}) - F^{-1}(\hat{\rho}_{00})\right) + \hat{\eta}x^*} \right\}, \qquad (2.13)$$

where $\hat{\rho}_{00}, \hat{\rho}_{01}, \hat{\rho}_{10}, \hat{\eta}$ are the posterior medians given the data D_n.

3. Simulation Studies

3.1 Simulation Set-up and Scenarios

We evaluate design operating characteristics by assuming a logistic link function $F(u) = (1 + exp(-u))^{-1}$ for the working model. DLT responses are generated assuming both a logistic link function and three other link functions to assess the performance of the method under model misspecification. These are (1) the probit link $F(u) = \Phi(u)$, where

$\Phi(\cdot)$ is the cdf of the standard normal distribution, (2) the normal link $F(u) = \Phi(u/\sigma)$, and (3) the complementary log-log link $F(u) = 1 - exp(-exp(u))$. We present four scenarios for the true MTD curve, two for each model reparameterization. In all cases, the target probability of DLT is fixed at $\theta = 0.33$ and the trial sample size is $n = 40$ patients. The first scenario corresponds to $\rho_{00} = 0.05$, $\Gamma_{A|0} = \Gamma_{B|0} = 0.5$, $\eta = 20$ and the corresponding true MTD curve is shown by the solid line in Figure 3(a). In the second scenario, $\rho_{00} = 0.05$, $\Gamma_{A|0} = 0.5$, $\Gamma_{B|0} = 0.8$, $\eta = 10$ and the true MTD curve is displayed by a solid line in Figure 4(a). Uniform priors were used for ρ_{00}, $\Gamma_{A|0}$, $\Gamma_{B|0}$. In scenario 3, we took $\rho_{00} = 10^{-7}$, $\rho_{01} = \rho_{10} = 3\times10^{-6}$, $\eta = 10$. The corresponding true MTD curve is shown by a solid line at the top right corner of Figure 5(a). This is a case where each agent is very safe within its range of doses. In the last scenario, $\rho_{00} = 0.01$, $\rho_{01} = 0.2$, $\rho_{10} = 0.9$, $\eta = 100$, see Figure 6(a). This is a case where the MTD of agent A when agent B is at its minimum dose level is within the range of doses of drug A but the MTD of agent B when drug A is at its minimum dose level is above the maximum dose level of agent B. Vague priors for ρ_{00}, ρ_{01}, ρ_{10} were selected as described in Section 2.2. In all 4 scenarios, a vague prior for η was selected by taking $E(\eta) = 30$ and $Var(\eta) = 900$. Many other scenarios were studied but are not included here due to space limitation. Algorithm 1 was used for scenarios 1 and 2 and algorithm 2 was used for scenarios 2 and 3. For each scenario, $m = 1000$ trials were simulated using the logistic link function for the working model and logistic, probit, normal with $\sigma = 2$, and complementary log-log link functions for the true model. The parameter values μ, β, γ, η of these models were selected in such a way that they all have the same true MTD curve. The extent of departure of the true model from the working model is illustrated in Figure 2(a-c) in the case $\rho_{00} = 0.05$, $\Gamma_{A|0} = \Gamma_{B|0} = 0.8$, and $\eta = 5$.

3.2 Design Operating Characteristics

We evaluate the performance of the methods by assessing the safety of the trial designs and the efficiency of the estimate of the MTD curve.

3.2.1 Safety

We assess trial safety by reporting the average percent of DLTs across all m = 1000 trials and the percent of trials that have a DLT rate exceeding $\theta + \delta$, for $\delta = 0.05, 0.1$. The threshold $\theta + 0.1$ is usually considered to be an indication of an excessive DLT rate.

3.2.2 Efficiency

We present an estimate of the MTD curve using the average posterior medians of the model parameters. For example under reparameterization 1, the estimate is

$$\bar{C} = \left\{ (x^*, y^*) : 0 \le x^*, y^* \le 1, y^* = \frac{\left(F^{-1}(\theta) - F^{-1}(\bar{\rho}_{00})\right)\left(1 - x^* / \bar{\Gamma}_{A|0}\right)}{\left(\left(F^{-1}(\theta) - F^{-1}(\bar{\rho}_{00})\right) / \bar{\Gamma}_{B|0}\right) + \bar{\eta} x^*} \right\}, \qquad (3.1)$$

where $F(\cdot)$ is the logistic function and $\bar{\rho}_{00}, \bar{\Gamma}_{A|0}, \bar{\Gamma}_{B|0}, \bar{\eta}$ are the average posterior medians of the parameters ρ_{00}, $\Gamma_{A|0}$, $\Gamma_{B|0}$, η from all $m = 1000$ trials. The next measure of efficiency is the pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve. For $i = 1,\ldots,m$, let C_i be the estimated MTD curve and C_{true} be the true MTD curve. For every point $(x,y) \in C_{true}$, let

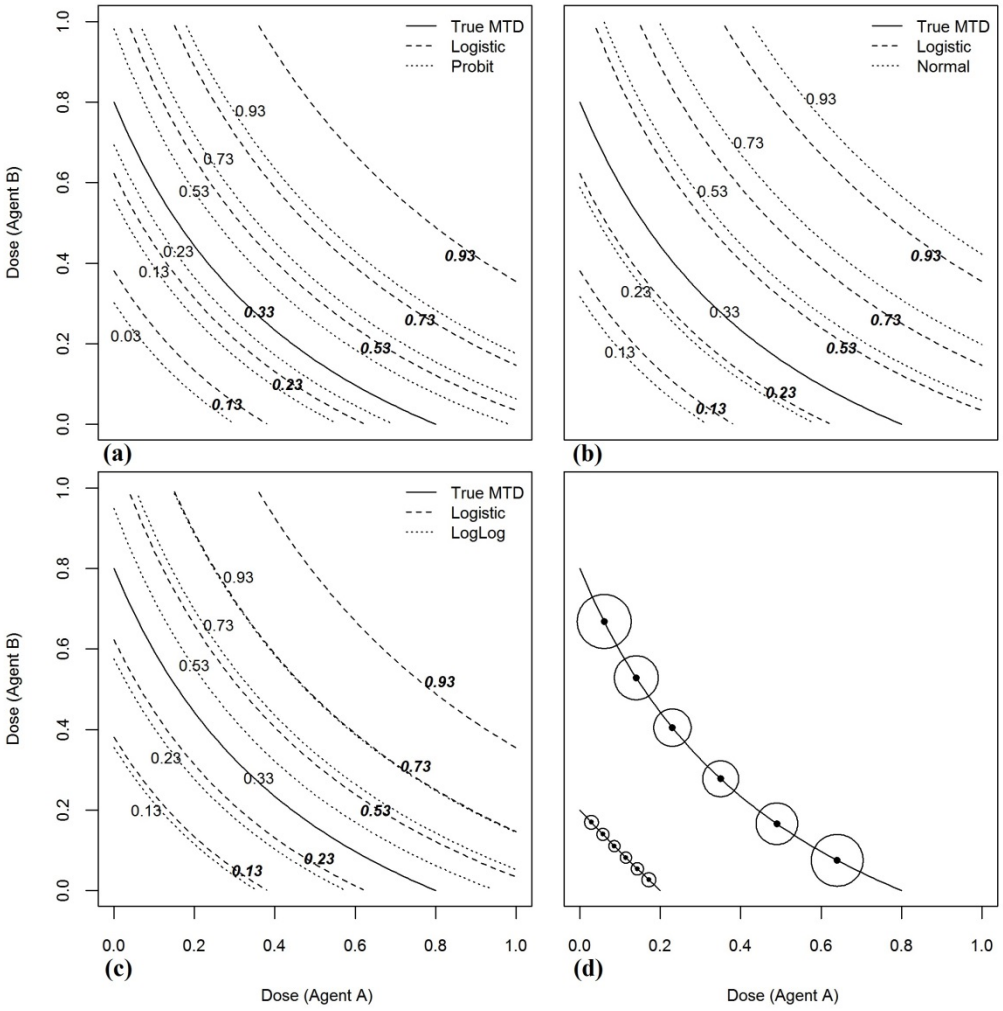

Figure 2. Contour plots from the logistic and probit (a), normal (b), and complementary log-log (c) dose-toxicity models. The probabilities of DLT corresponding to the logistic model are shown in bold italic. (d) shows the circles for calculating the percent of MTD recommendation for two scenarios when the tolerance $p = 0.1$.

$$d_{(x,y)}^{(i)} = sign(y' - y) \times \min_{\{(x^*,y^*):(x^*,y^*)\in C_i\}} \left((x - x^*)^2 + (y - y^*)^2\right)^{1/2} \tag{3.2}$$

where y' is such that $(x, y') \in C_i$. This is the minimum relative distance of the point (x,y) on the true MTD curve to the estimated MTD curve C_i. If the point (x,y) is below C_i, then $d_{(x,y)}^{(i)}$ is positive. Otherwise, it is negative. Let

$$d_{(x,y)} = m^{-1} \sum_{i=1}^{m} d_{(x,y)}^{(i)}. \tag{3.3}$$

This is the pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve and can be interpreted as the pointwise average bias in estimating the MTD. Let $\Delta(x,y)$ be the Euclidian distance between the minimum dose combination

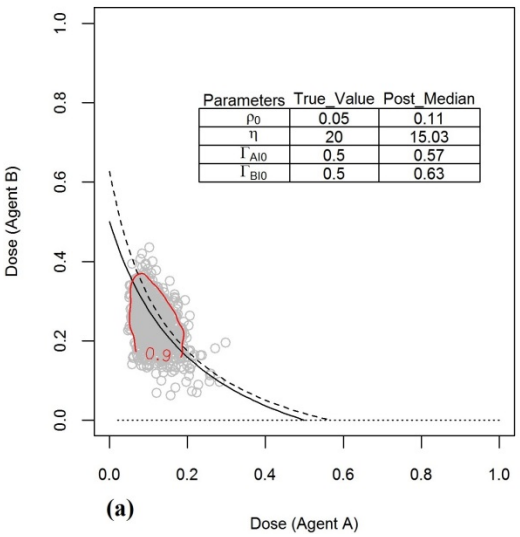

True model	Average % DLT	% w. DLT rate > $\theta + 0.05$	% w. DLT rate > $\theta + 0.10$
Logistic	28	0.5	0.0
Probit	28	0.3	0.0
Normal	28	0.3	0.0
LogLog	28	0.2	0.0

Table 1. Average percent of DLTs and percent of trials with DLT rate exceeding $\theta + \delta$ under scenario 1 for various true models. Working model is logistic.

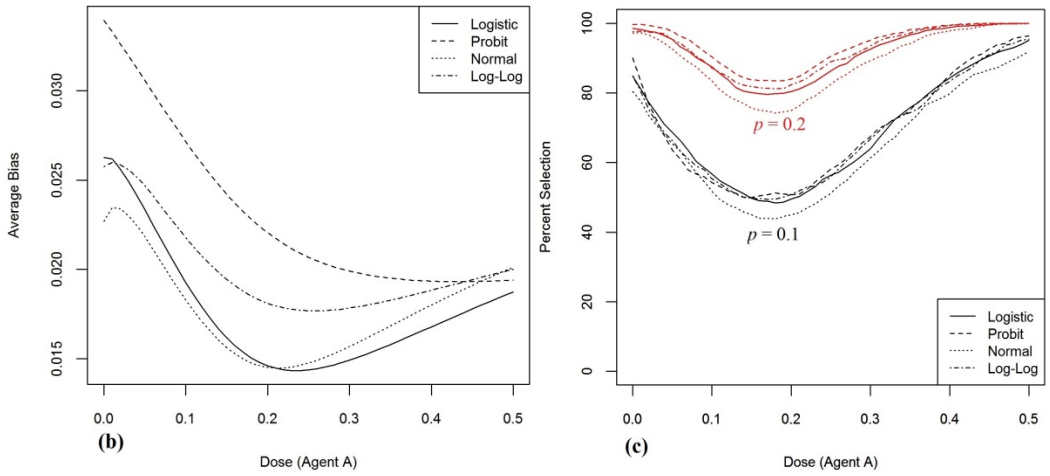

Figure 3. Summary statistics from m =1000 simulated trials under scenario 1. (a) shows the true and estimated MTD curve. The grey diamonds represent the last dose combination from each simulated trial along with a 90% confidence region, (b) pointwise average bias, (c) pointwise percent MTD recommendation for tolerances p = 0.1, 0.2.

(0,0) and the point (x,y) on the true MTD curve and $0 < p < 1$. The last measure of efficiency we consider is

$$p_{(x,y)} = m^{-1}\sum_{i=1}^{m} I\left(|\, d_{(x,y)}^{(i)} \,| \leq p\Delta(x,y)\right). \tag{3.4}$$

This is the pointwise percent of trials for which the minimum distance of the point (x,y) on the true MTD curve to the estimated MTD curve C_i is no more than (100×p)% of the true MTD. This statistic is equivalent to drawing a circle with center (x,y) on the true MTD curve and radius $p\Delta(x,y)$ and calculating the percent of trials with MTD curve estimate C_i falling inside the circle. This will give us the percent of trials with MTD

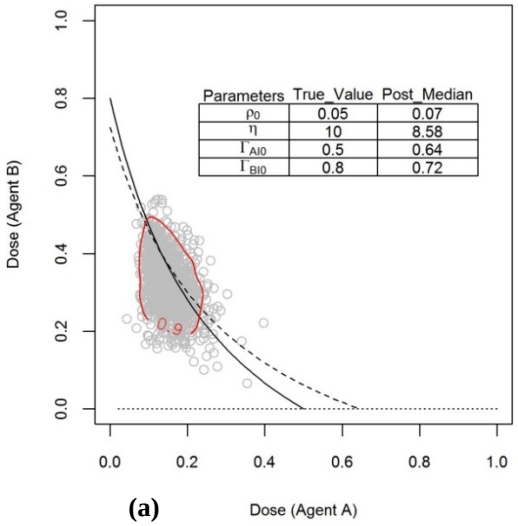

True model	Average % DLT	% w. DLT rate > θ + 0.05	% w. DLT rate > θ + 0.10
Logistic	23	0.0	0.0
Probit	24	0.0	0.0
Normal	23	0.0	0.0
LogLog	23	0.1	0.0

Table 2. Average percent of DLTs and percent of trials with DLT rate exceeding $\theta + \delta$ under scenario 2 for various true models. Working model is logistic.

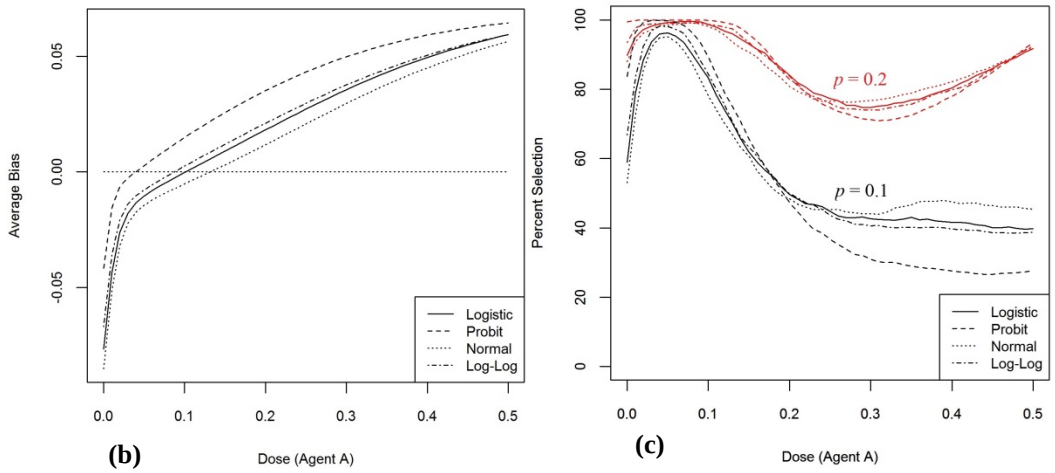

Figure 4. Summary statistics from m =1000 simulated trials under scenario 2. (a) shows the true and estimated MTD curve. The grey diamonds represent the last dose combination from each simulated trial along with a 90% confidence region, (b) pointwise average bias, (c) pointwise percent MTD recommendation for tolerances p = 0.1, 0.2.

recommendation within (100×p)% of the true MTD for a given tolerance p. Figure 2(d) illustrates the variability of the radius of the tolerance circles for various locations of dose combinations on the true MTD curve under two different scenarios ρ_{00} = 0.05, η = 5, $\Gamma_{A|0}$ = $\Gamma_{B|0}$ = 0.2, and $\Gamma_{A|0}$ = $\Gamma_{B|0}$ = 0.8. Here, p = 0.1. We can see that the farther the true MTD is from the minimum dose combination, the larger is the tolerance for estimating the percent of MTD recommendation. The MTD curve common to all models is represented by a solid line. The probabilities of DLT of the toxicity profiles of the logistic model are shown in bold italic below the diagonal and the corresponding probabilities for the other models are included above the diagonal. Figure 2(a) is a situation where DLT responses are generated from a model with tighter toxicity profiles relative to the working model, Figure 2(b) is a case where the working model is tighter relative to the true model, and

Figure 2(c) is a situation where the working model is wider relative to the true model above the MTD curve but is tighter below the MTD curve.

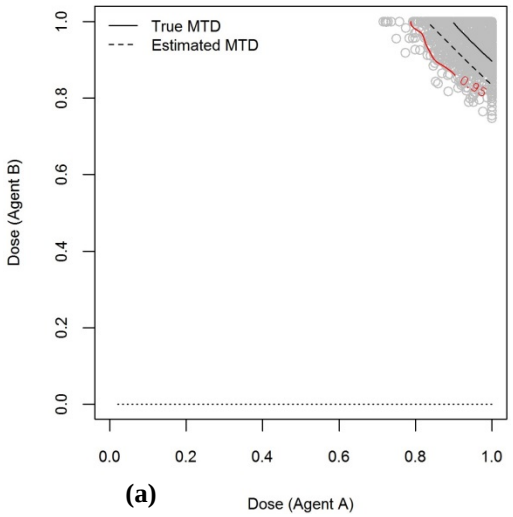

True model	Average % DLT	% w. DLT rate > $\theta + 0.05$	% w. DLT rate > $\theta + 0.10$
Logistic	16	0.0	0.0
Probit	16	0.0	0.0
Normal	16	0.0	0.0
LogLog	16	0.0	0.0

Table 3. Average percent of DLTs and percent of trials with DLT rate exceeding $\theta + \delta$ under scenario 3 for various true models. Working model is logistic.

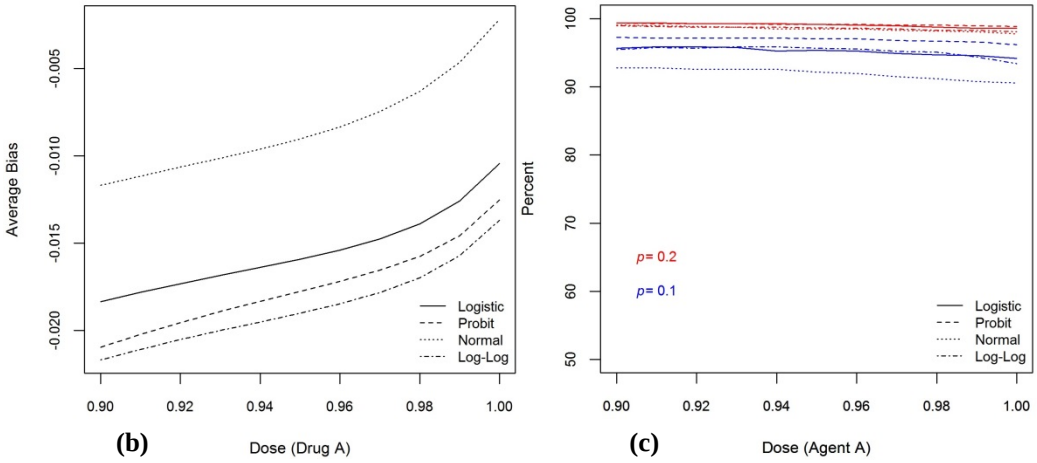

Figure 5. Summary statistics from m =1000 simulated trials under scenario 3. (a) shows the true and estimated MTD curve. The grey diamonds represent the last dose combination from each simulated trial along with a 90% confidence region, (b) pointwise average bias, (c) pointwise percent MTD recommendation for tolerances p = 0.1, 0.2.

3.3 Results

3.3.1 Trial Safety

Tables 1-4 shows that the average percent of DLTs varies between 16% and 34% across the four scenarios. In general, the average DLT rate tends to be lower when the true MTD curve is farther away from the minimum dose combination. These tables also show that the percent of trials with an excessive number of DLTs as defined by a DLT rate exceeding θ + 0.1 is very small. Further simulations (results not shown) under scenarios where the true MTD curve is close to the minimum dose combination available in the trial show that his rate does not exceed 5%. Based on these findings, we conclude that the methodology is safe in general.

3.3.1 Trial Efficiency

Figures 3-6(a) show the plots of the true and estimated MTD curves. The estimated MTD curve shown by a dashed line was obtained using (3.1) or the corresponding equation under model reparameterization 2 and DLT responses were simulated using the true logistic model. In all cases, the estimated MTD curve is very close to the true MTD except at the edges of the curve. This is probably due to the fact that we are using uniform priors for $\Gamma_{A|0}$ and $\Gamma_{B|0}$ with prior means equal to 0.5 for scenarios 1 and 2 and uniform priors for ρ_{01} and ρ_{10} for scenarios 3 and 4. For each scenario, scatter plot of the last dose combinations from each of the m = 1000 simulated trials along with 90% confidence region is also included. These statistics are useful for clinicians who plan to use the last dose combination in a phase I trial as the recommended phase II dose combinations as in [10].

True model	Average % DLT	% w. DLT rate > θ + 0.05	% w. DLT rate > θ + 0.10
Logistic	34	8.1	0.0
Probit	33	1.9	0.0
Normal	34	11.8	0.4
LogLog	33	3.0	0.1

Table 4. Average percent of DLTs and percent of trials with DLT rate exceeding $\theta + \delta$ under scenario 4 for various true models. Working model is logistic.

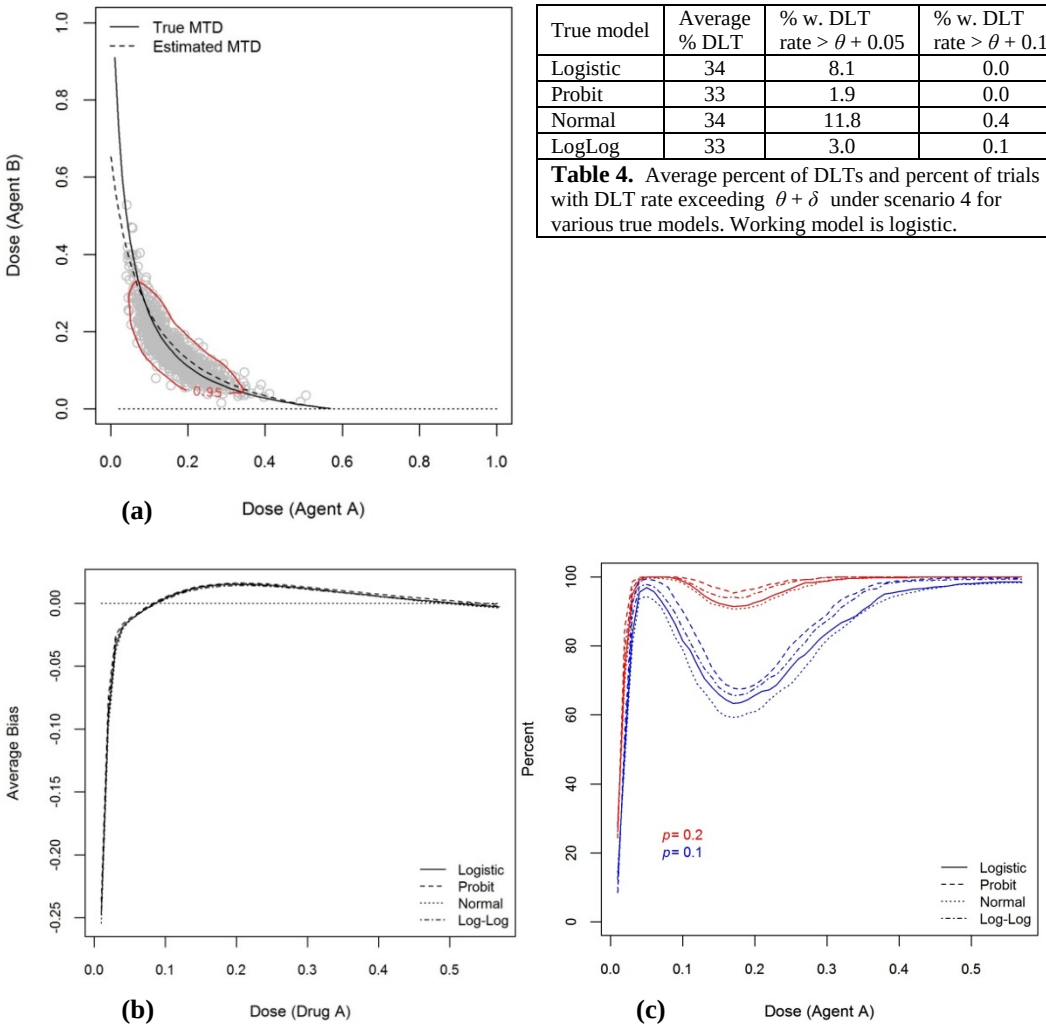

Figure 6. Summary statistics from m =1000 simulated trials under scenario 1. (a) shows the true and estimated MTD curve. The grey diamonds represent the last dose combination from each

simulated trial along with a 90% confidence region, (b) pointwise average bias, (c) pointwise percent MTD recommendation for tolerances $p = 0.1, 0.2$.

Figures 3-6(b) display the pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve as defined by (3.3). This is a measure of pointwise average bias of the estimate of the MTD. In scenario 1 (Figure 3(b)), the maximum average bias is about 0.025 when DLT responses are generated from the true logistic model. This corresponds to 5% of the distance from the minimum dose combination (0, 0) to the true MTD dose combination (0, 0.5). Scenarios 2 and 3 also show that the maximum average bias is no more than 10% of the distance from (0,0) to the corresponding true MTD dose combination. Scenario 4 (Figure 6(b)) show a much larger bias when dose level of Drug *A* is 0. This is due to the fact that the MTD of agent *B* is greater than the maximum dose level of agent *B* available in the trial. However, the pointwise average bias is negligible elsewhere. We conclude that the pointwise average bias is fairly small along the whole MTD curve except at the edges of the curve.

Figures 3-6(c) show the pointwise percent of trials for which the minimum distance from the true MTD curve to the estimated MTD curve is no more than $(100\times p)\%$ of the true MTD for $p = 0.1$ and $p = 0.2$. This can be interpreted as the percent of MTD recommendation for a given tolerance p. With a tolerance of $p = 0.1$, the percent of trials with correct MTD recommendation varies between 40% and 95% under scenarios 1, 2, and 3. Under scenario 4, the percent recommendation is low (Figure 6(c)) at the minimum dose level of agent *A*, consistent with the bias found at this dose level. The percent increases as we move away from this dose level. Based on these results and others from scenarios not shown here, we conclude that the design is practically efficient in general in recommending the MTD curve estimate.

3.3.1 Model Robustness

The average DLT rates and percent of trials with DLT rate $> \theta + \delta$ shown in Tables 1-4 when the true model is misspecified are very close. Furthermore, differences in the pointwise average bias between the different models and the logistic model shown in Figures 3-6(b) are negligible relative to the dose range of both agents. Similar conclusion holds for the pointwise percent of MTD recommendation shown in Figures 3-6(c). The largest difference of about 10% is observed under scenario 2 (Figure 4(c)) when the true model is the complementary log-log and the dose of agent *A* is between 0.3 and 0.5 and the tolerance is $p = 0.1$. We conclude that our method is fairly robust to model misspecifications under the dose-toxicity family of models of the form $\text{Prob}(\text{DLT} \mid x,y) = F(\mu + \beta\, x + \gamma\, y + \eta\, x\, y)$ for some selected link functions $F(\cdot)$.

4. Discussion

We described Bayesian adaptive designs for cancer phase I clinical trials using two drugs with continuous dose levels. The goal is to estimate the MTD curve in the two-dimensional Cartesian plane. We reviewed the reparameterization proposed in [15] which assumes that the MTDs $\Gamma_{A|0}$ and $\Gamma_{B|0}$ are within the range of doses available in the trial and introduces a new reparameterization that relaxes these conditions. In each case, vague priors were used for quantifying the toxicity profile of each agent *a priori*. We also introduced another algorithm for dose escalation where cohorts of two patients are enrolled simultaneously and the patients receive different dose combinations. We studied design operating characteristics of the method under a large number of practical scenarios (only four of them are included due to space limitation) and under several model

misspecifications. In all simulations, we used a sample size of $n = 40$ patients. We found that in general, the methodology is safe in terms of the probability that a prospective trial will results in an excessively high number of DLTs. Under scenarios where the true MTD curve is near the minimum dose combination or below it using reparameterization 2, there is a high probability of stopping the trial. We used several measures to assess the efficiency of the estimate of the MTD and in the majority of scenarios, the percent of MTD recommendation is good and increases as the true MTD curve drifts away from the minimum dose combination. We recommend that clinicians select dose combinations around the middle of the MTD curve for efficacy evaluation since the percent of recommendation is high and dose combinations where the level of one of the agents is very low may not be of interest for efficacy studies. We also showed that the method is practically robust with respect to trial safety and efficiency under a reasonable class of model misspecification. We also plan to study the performance of the proposed design under a class of models which allow synergistic and antagonistic relation between the drugs as described in [20]. Finally, we plan to assess the performance of the method when the doses of the two agents are discretized using the method discussed in [16, 18] and compare the performance of the resulting design with the methods described in [8, 10, 11].

Acknowledgment

This work is supported in part by the National Center for Research Resources, Grant UL1RR033176, and is now at the National Center for Advancing Translational Sciences, Grant UL1TR000124 (M.T and A.R), Grant 5P01CA098912-05 (M.T and A.R.), P01DK046763 (A.R.), and 2R01CA108646-07A1 (A.R.). The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

References

1. Roberts TG, Goulart BH, Squitieri L, Stallings SC, Halpern EF, Chabner BA, Gazelle GS, Finkelstein SN, Clark JW. Trends in the risks and benefits to patients with cancer participating in phase 1 clinical trials. *Jama-Journal of the American Medical Association* 2004; **292**: 2130-2140.
2. Ting N. *Dose Finding in Drug Development.* (First edn). Springer: (New York, 2006.
3. Chevret S. *Statistical Methods for Dose-finding Experiments.* Wiley: Chichester, 2006.
4. Le Tourneau C, Lee JJ, Siu LL. Dose Escalation Methods in Phase I Cancer Clinical Trials. *Journal of the National Cancer Institute* 2009; **101**: 708-720.
5. Frey E, III., Karon M, Levin RH, Freireich EJ, Taylor RJ, Hananian J, Selawry O, Holland JF, Hoogstraten B, Wolman IJ, Abir E, Sawitsky A, Lee S, Mills SD, Burgert EO, Spurr CL, Patterson RB, Ebaugh FG, James GWI, Moon JH. The Effectiveness of Combinations of Antileukemic Agents in Inducing and Maintaining Remission in Children with Acute Leukemia. *Blood* 1965; **26**: 642-656.
6. Thall PF, Millikan RE, Mueller P, Lee SJ. Dose-finding with two agents in phase I oncology trials. *Biometrics* 2003; **59**: 487-496.

7. Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics* 2005; **61**: 217-222.
8. Yin GS, Yuan Y. A Latent Contingency Table Approach to Dose Finding for Combinations of Two Agents. *Biometrics* 2009; **65**: 866-875.
9. Yin GS, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society Series C-Applied Statistics* 2009; **58**: 211-224.
10. Braun TM, Wang SF. A Hierarchical Bayesian Design for Phase I Trials of Novel Combinations of Cancer Therapeutic Agents. *Biometrics* 2010; **66**: 805-812.
11. Wages NA, Conaway MR, O'Quigley J. Continual Reassessment Method for Partial Ordering. *Biometrics* 2011; **67**: 1555-1563.
12. Wages NA, Conaway MR, O'Quigley J. Dose-finding design for multi-drug combinations. *Clinical Trials* 2011; **8**: 380-389.
13. Sweeting MJ, Mander AP. Escalation strategies for combination therapy Phase I trials. *Pharmaceutical Statistics* 2012; **11**: 258-266.
14. Shi Y, Yin G. Escalation with overdose control for phase I drug-combination trials. *Statistics in Medicine* 2013: in press.
15. Tighiouart M, Piantadosi S, Rogatko A. Dose finding with drug combinations in cancer phase I clinical trials using conditional escalation with overdose control. *Stat Med* 2014; **In press**.
16. Babb J, Rogatko A, Zacks S. Cancer Phase I clinical Trials: efficient dose escalation with overdose control. *Stat Med* 1998; **17**: 1103-1120.
17. Tighiouart M, Rogatko A, Babb JS. Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Statistics in Medicine* 2005; **24**: 2183-2196.
18. Tighiouart M, Rogatko A. Dose Finding with Escalation with Overdose Control (EWOC) in Cancer Clinical Trials. *Statistical Science* 2010; **25**: 217-226.
19. Lunn DJ, Thomas A, Best N, Spiegelhalter D. WinBUGS - A Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing* 2000; **10**: 325-337.
20. Gasparini M. General classes of multiple binary regression models in dose finding problems for combination therapies. *Journal of the Royal Statistical Society Series C-Applied Statistics* 2013; **62**: 115-133.