Dose Finding for Drug Combination in Early Cancer Phase I Trials using Conditional Continual Reassessment Method

Mourad Tighiouart and Quanlin Li

Biostatistics and Bioinformatics Research Center, Cedars-Sinai Medical Center
8700 Beverly Blvd, Los Angeles, CA 90048

Abstract

We present a Bayesian adaptive design 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 for continuous dose levels of the two agents. Parametric models are used to describe the relationship between the doses and the probability of dose limiting toxicity (DLT). Trial design proceeds using the continual reassessment method, where at each stage of the trial, we seek the dose of one agent with estimated probability of DLT closest to a target probability of DLT given the current dose of the other agent. At the end of the trial, we estimate the MTD curve as a function of Bayes estimates of the model parameters. We evaluate design operating characteristics in terms of safety of the trial and percent of dose recommendation at dose combination neighbourhoods around the true MTD. We also examine the performance of the approach under model misspecifications for the true dose-toxicity relationship.

KeyWords: Cancer phase I trials; Maximum tolerated dose; Continual reassessment method; Drug combination; Dose limiting toxicity; Continuous dose.

1. Introduction

The primary objective of cancer phase I clinical trials 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. Single agent dose finding methods have been proposed and studied extensively in the past two decades, see [1-3] for a review. Combining several drugs can help reduce tumor resistance to chemotherapy by targeting different signaling pathways simultaneously [4]. 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. 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. The general problem can be stated as follows. Let A_i, $i = 1,\ldots,k$ be k drugs and S_i c 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:

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

For $k = 2$, a number of designs for estimating C have been proposed in the past decade, see [5-15]. By design, these methods do not apply to the case of continuous dose levels of the two agents. Except for [5, 6, 15], these 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. [16] 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 continual reassessment method (CRM) [17].

2. Model

2.1 Dose-Toxicity Model

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 necessary and 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

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

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.

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. Unlike the reparameterization used in [18] restricting the MTD curve to be in the lower triangle $\{(x,y): x+y \leq 1\}$, this reparametrization allows the MTD curve to lie anywhere in the x-y Cartesian plane. It can be shown that

$$\begin{cases} \mu = F^{-1}(\rho_{00}) \\ \beta = (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00})). \\ \gamma = (F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})) \end{cases} \tag{2.4}$$

The MTD (2.3) becomes

$$C = \left\{ (x^*, y^*) : y^* = \frac{\left(F^{-1}(\theta) - F^{-1}(\rho_{00})\right) - \left(F^{-1}(\rho_{10}) - F^{-1}(\rho_{00})\right)x^*}{\left(F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})\right) + \eta x^*} \right\}. \tag{2.5}$$

Figure 1shows MTD curves when $\rho_{00} = 0.05$, $\rho_{01} = \rho_{10} = 3\times10^{-6}$, and three values for the interaction coefficient $\eta = 20, 40, 60$.

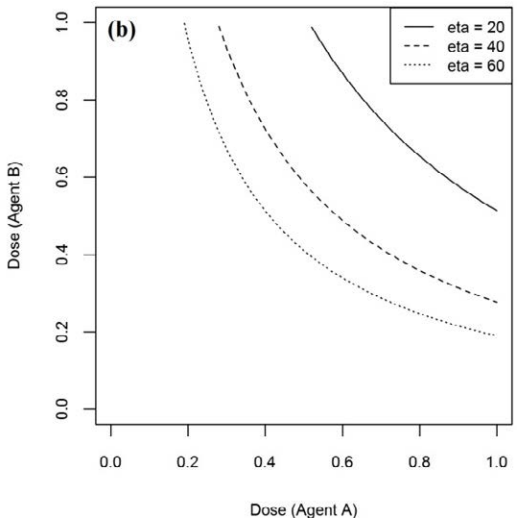

Figure 1.MTD curves for selected values of the interaction coefficient η. Target probability of DLT $\theta = 0.33$ and logistic link $F(u) = (1 + e^{-u})^{-1}$.

The target probability of DLT is $\theta = 0.33$ and the link function is the logistic $F(u) = (1 + e^{-u})^{-1}$. Figure 1 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 is

$$L(\rho_{00},\rho_{01},\rho_{10},\eta \mid D_n) = \prod_{i=1}^{n} \left(H(\rho_{00},\rho_{01},\rho_{10},\eta;x_i,y_i)\right)^{z_i} \times \left(1-H(\rho_{00},\rho_{01},\rho_{10},\eta;x_i,y_i)\right)^{1-z_i}, \tag{2.6}$$

where

$$\begin{aligned} &H(\rho_{00},\rho_{01},\rho_{10},\eta;x_i,y_i) \\ &= F\left(F^{-1}(\rho_{00}) + \left(F^{-1}(\rho_{10}) - F^{-1}(\rho_{00})\right)x_i + \left(F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})\right)y_i + \eta x_i y_i\right). \end{aligned} \tag{2.7}$$

2.2 Prior and Posterior Distributions

We note that (2.4) 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 $\rho_{01} \sim \text{beta}(a_1, b_1)$, $\rho_{10} \sim \text{beta}(a_2, b_2)$, and conditional on (ρ_{01}, ρ_{10}), $\rho_{00} / \min(\rho_{01}, \rho_{10}) \sim \text{beta}(a_3, b_3)$. Vague priors for these parameters are achieved by taking $a_j = b_j = 1$, $j = 1, 2, 3$. The prior for η is $\eta \sim \text{gamma}(a, b)$ with mean $E(\eta) = a / b$ and variance $Var(\eta) = a / b^2$. Using Bayes rule, the posterior distribution of the model parameters is proportional to the product of the likelihood and prior distribution

$$\begin{aligned} \pi(\rho_{00},\rho_{01},\rho_{10},\eta \mid D_n) \propto &\prod_{i=1}^{n} \left(H(\rho_{00},\rho_{01},\rho_{10},\eta;x_i,y_i)\right)^{\delta_i} \left(1-H(\rho_{00},\rho_{01},\rho_{10},\eta;x_i,y_i)\right)^{1-\delta_i} \\ &\times \pi(\rho_{01})\pi(\rho_{10})\pi(\rho_{00} \mid \rho_{01},\rho_{10})\pi(\eta), \end{aligned} \tag{2.8}$$

where

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

Features of this posterior distribution are estimated using WinBUGS[19] and JAGS.

2.3 Trial Design

Tighiouart et al. [18] used escalation with overdose control (EWOC) [20-22] principle to determine the dose to be allocated to each subsequent patient. Here, we extend that algorithm by treating cohorts of two patients simultaneously and determine the next dose using CRM. The adaptive design proceeds as follows:

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
   $$x_3 = \underset{u}{\text{argmin}} \left|\hat{P}\text{rob}(\delta = 1 \mid u, y_1) - \theta\right| \text{ and } y_4 = \underset{v}{\text{argmin}} \left|\hat{P}\text{rob}(\delta = 1 \mid x_2, v) - \theta\right|.$$
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}) and patient $2i$ receives dose (x_{2i-2}, y_{2i}), where
   $$x_{2i-1} = \underset{u}{\text{argmin}} \left|\hat{P}\text{rob}(\delta = 1 \mid u, y_{2i-3}) - \theta\right| \text{ and } y_{2i} = \underset{v}{\text{argmin}} \left|\hat{P}\text{rob}(\delta = 1 \mid x_{2i-2}, v) - \theta\right|.$$

A similar expression is derived when i is old. Here $\hat{P}\text{rob}(\delta = 1 \mid x, y)$ is the estimated probability of DLT obtained by replacing the parameters ρ_{00}, ρ_{01}, ρ_{10}, and η in $H(\rho_{00}, \rho_{01}, \rho_{10}, \eta; x, y)$ by their posterior medians given the data.

4. Repeat step 3 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.

At the end of the trial, we estimate the MTD curve using Bayes estimates of the parameters defining this curve as

$$\hat{C} = \left\{ (x^*, y^*) : y^* = \frac{\left(F^{-1}(\theta) - F^{-1}(\hat{\rho}_{00})\right) - \left(F^{-1}(\hat{\rho}_{10}) - F^{-1}(\hat{\rho}_{00})\right) x^*}{\left(F^{-1}(\hat{\rho}_{01}) - F^{-1}(\hat{\rho}_{00})\right) + \hat{\eta} x^*} \right\}, \qquad (2.9)$$

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 three scenarios for the true MTD curve. 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 $\rho_{00} = 10^{-7}$, $\rho_{01} = \rho_{10} = 3\times10^{-6}$, $\eta = 10$. The corresponding true MTD curve is shown by a solid line at the top left corner of Figure 2(a). This is a case where each agent is very safe within its range of doses. In the second scenario, $\rho_{00} = 0.01$, $\rho_{01} = 0.2$, $\rho_{10} = 0.9$, $\eta = 20$, see Figure 3(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. In the last scenario, we took $\rho_{00} = 0.05$, $\rho_{01} = 0.9$, $\rho_{10} = 0.9$, $\eta = 20$, see Figure 4(a). This is a case where the MTDs of both agents A and B are within the range of their respective dose levels when the other agent is at its minimum dose level. Vague priors for ρ_{00}, ρ_{01}, ρ_{10} were selected as described in Section 2.2. In all 3 scenarios, a vague prior for η was selected by taking $E(\eta) = 21$ and $Var(\eta) = 542$. Many other scenarios were studied but are not included here due to space limitation. 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 5(a-c) in the case $\rho_{00} = 0.05$, $\rho_{01} = \rho_{10} = 0.46$ 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. Under the reparameterization, the estimate is

$$\hat{C} = \left\{ (x^*, y^*) : y^* = \frac{\left(F^{-1}(\theta) - F^{-1}(\hat{\rho}_{00})\right) - \left(F^{-1}(\hat{\rho}_{10}) - F^{-1}(\hat{\rho}_{00})\right)x^*}{\left(F^{-1}(\hat{\rho}_{01}) - F^{-1}(\hat{\rho}_{00})\right) + \hat{\eta}x^*} \right\}, \tag{3.1}$$

where $F(\cdot)$ is the logistic function and $\hat{\rho}_{00}, \hat{\rho}_{01}, \hat{\rho}_{10}, \hat{\eta}$ are the average posterior medians of the parameters ρ_{00}, ρ_{01}, ρ_{10} and 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

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

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 $d^{(i)}_{(x,y)}$ is positive. Otherwise, it is negative. Let

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

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

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

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 $(100\times p)\%$ 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

recommendation within (100×p)% of the true MTD for a given tolerance p. Figure 5(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 $\rho_{00} = 0.05$, $\eta = 5$, $\rho_{01} = \rho_{10} = 0.46$, and $\rho_{01} = \rho_{10} = 1$. 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.

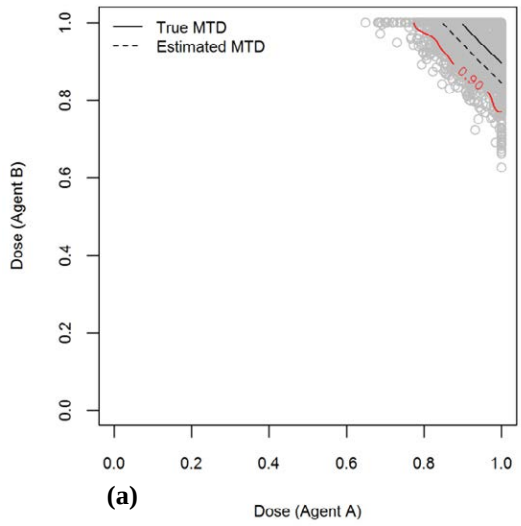

True model	Average % DLT	% w. DLT rate $>\theta + 0.05$	% w. DLT rate $>\theta + 0.10$
Logistic	16	0.0	0.0
Probit	16	0.0	0.0
Normal	16	0.0	0.0
LogLog	16	0.0	0.0

Table 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.

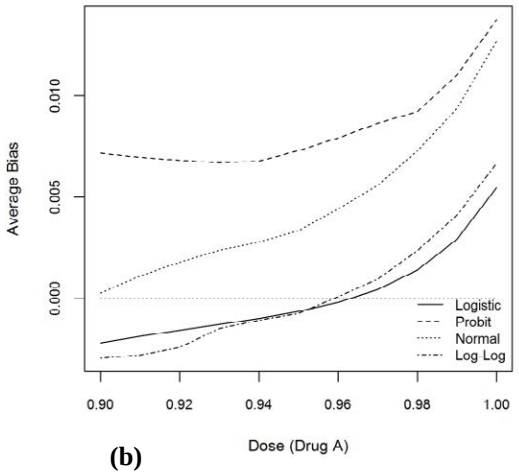

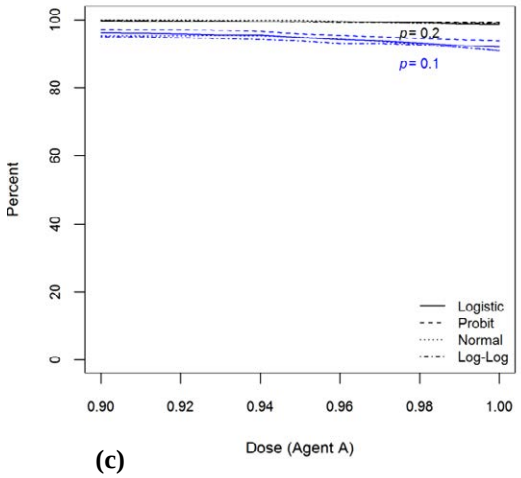

Figure 2. Summary statistics from m =1000 simulated trials under scenario 1. (a) shows the true and estimated MTD curve. The grey diamonds represent the last dose combination from each simulated trial along with a 90% confidence region, (b) pointwise average bias, (c) pointwise percent MTD recommendation for tolerances p = 0.1, 0.2.

3.3 Results

3.3.1 Trial Safety

Tables 1-3 shows that the average percent of DLTs varies between 16% and 38% across the three 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 $\theta + 0.1$ 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 this rate does not exceed 10%. Based on these findings, we conclude that the methodology is safe in general.

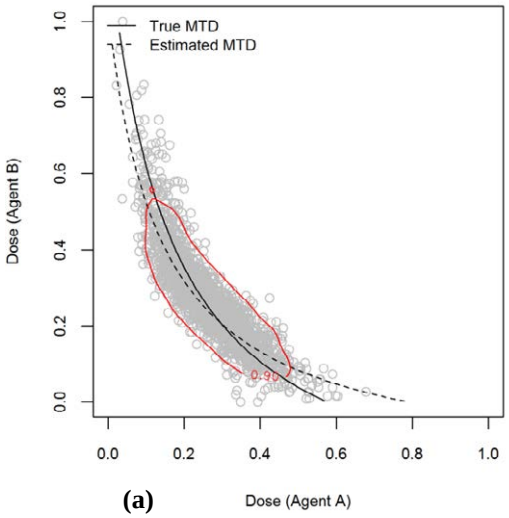

(a)

True model	Average % DLT	% w. DLT rate $>\theta + 0.05$	% w. DLT rate $>\theta + 0.10$
Logistic	34	16	1.1
Probit	33	8	0.3
Normal	34	18	2.3
LogLog	33	10	0.6

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.

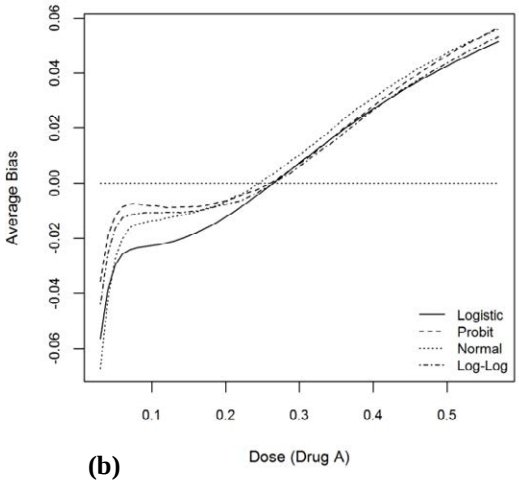

(b)

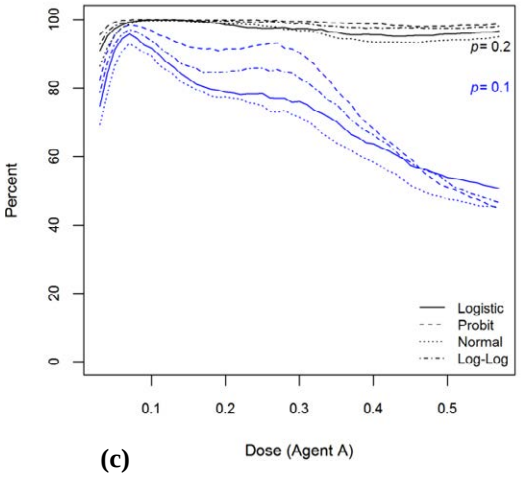

(c)

Figure 3. 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.

3.3.1 Trial Efficiency

Figures 2-4(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) 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 ρ_{01} and ρ_{10}. 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 [9].

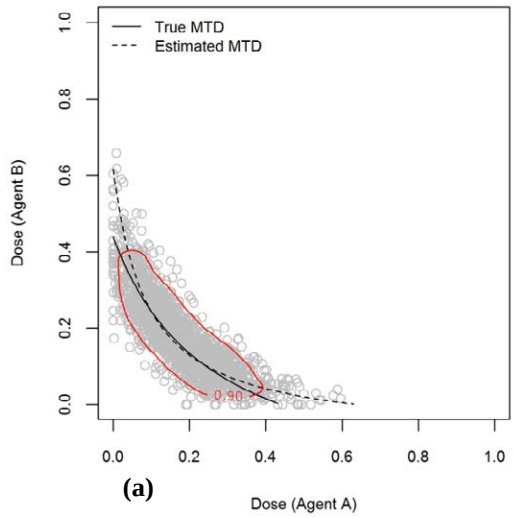

True model	Average % DLT	% w. DLT rate $>\theta+0.05$	% w. DLT rate $>\theta+0.10$
Logistic	38	43	10
Probit	37	35	5
Normal	38	42	10
LogLog	37	35	7

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.

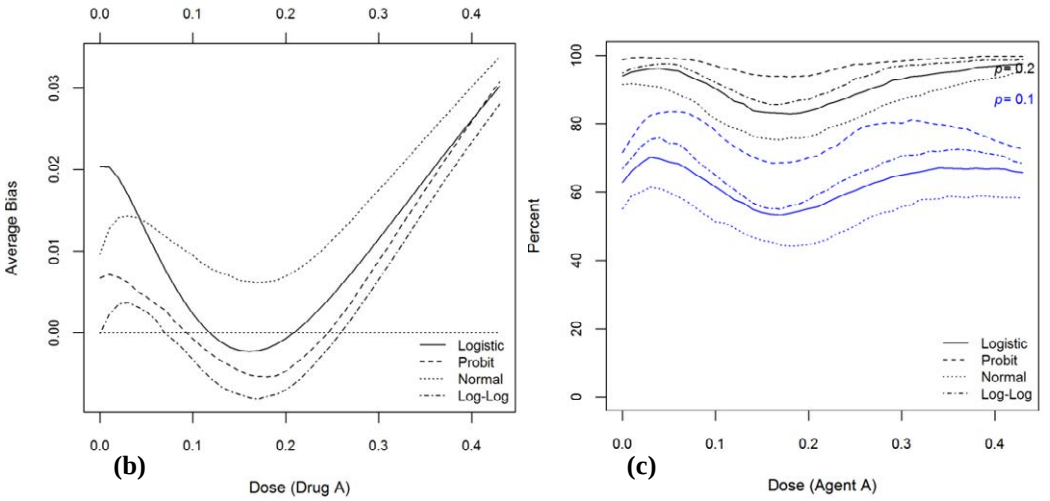

Figure 4. Summary statistics from m =1000 simulated trials under scenario 3. (a) shows the true and estimated MTD curve. The grey diamonds represent the last dose combination from each simulated trial along with a 90% confidence region, (b) pointwise average bias, (c) pointwise percent MTD recommendation for tolerances p = 0.1, 0.2.

Figures 2-4(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 3 (Figure 4(b)), the maximum average bias is about 0.03 when DLT responses are generated from the true logistic model. This corresponds to 6% of the distance from the minimum dose combination (0, 0) to the true MTD dose combination (0, 0.5). Scenarios 1 and 2 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. We conclude that the pointwise average bias is fairly small along the whole MTD curve.

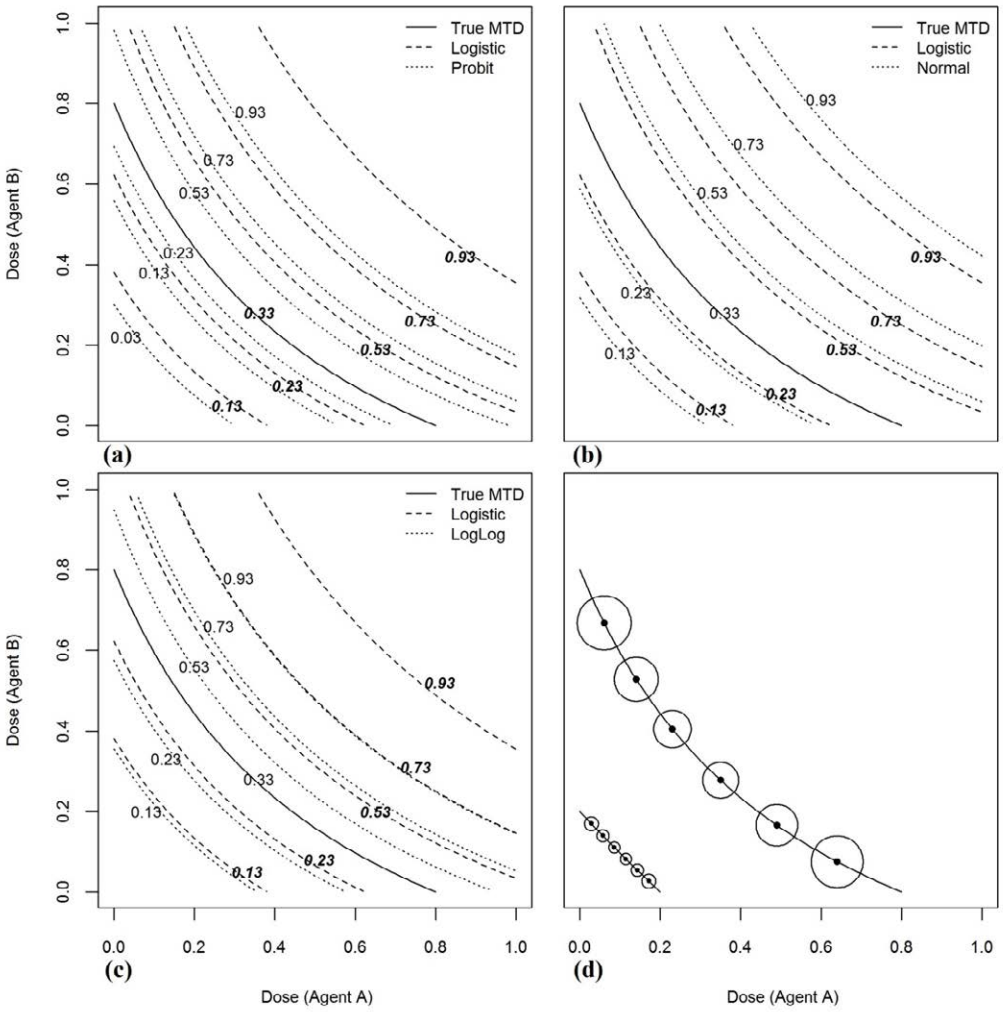

Figure 5. Contour plots from the logistic and probit (a), normal (b), and complementary log-log (c) dose-toxicity models. The probabilities of DLT corresponding to the logistic model are shown in bold italic. (d) shows the circles for calculating the percent of MTD recommendation for two scenarios when the tolerance $p = 0.1$.

Figures 2-4(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 50% and 95% under all three scenarios and is above 90% with a tolerance $p = 0.2$. 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 1-3 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 2-4(b) are negligible relative to the dose range of both agents. Similar conclusion holds for the pointwise percent of MTD recommendation shown in Figures 2-4(c). The largest difference of about 15% is observed under scenario 3 (Figure 4(c)) when the true model is the normal or probit 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 Prob(DLT | 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 introduce a new reparameterization that relaxes conditions that the MTD of both drugs are within the range of doses available in the trial. 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 using univariate CRM algorithm. We studied design operating characteristics of the method under a large number of practical scenarios (only three 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, 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 [23]. Finally, we plan to assess the performance of the method when the doses of the two agents are discretized using the method discussed in [20, 22] and compare the performance of the resulting design with the methods described in [7, 9, 10].

Acknowledgment

This work is supported in part by the National Institute of Health Grant Number 1 R01CA188480-01A1 (M.T, Q.L), the National Center for Research Resources, Grant

UL1RR033176, and is now at the National Center for Advancing Translational Sciences, Grant UL1TR000124 (M.T), and 2 P01 CA098912 (M.T).

References

1. Ting N. *Dose Finding in Drug Development.* (First edn). Springer: (New York, 2006.
2. Chevret S. *Statistical Methods for Dose-finding Experiments.* Wiley: Chichester, 2006.
3. 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.
4. 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.
5. Thall PF, Millikan RE, Mueller P, Lee SJ. Dose-finding with two agents in phase I oncology trials. *Biometrics* 2003; **59**: 487-496.
6. Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics* 2005; **61**: 217-222.
7. Yin GS, Yuan Y. A Latent Contingency Table Approach to Dose Finding for Combinations of Two Agents. *Biometrics* 2009; **65**: 866-875.
8. 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.
9. Braun TM, Wang SF. A Hierarchical Bayesian Design for Phase I Trials of Novel Combinations of Cancer Therapeutic Agents. *Biometrics* 2010; **66**: 805-812.
10. Wages NA, Conaway MR, O'Quigley J. Continual Reassessment Method for Partial Ordering. *Biometrics* 2011; **67**: 1555-1563.
11. Wages NA, Conaway MR, O'Quigley J. Dose-finding design for multi-drug combinations. *Clinical Trials* 2011; **8**: 380-389.
12. Sweeting MJ, Mander AP. Escalation strategies for combination therapy Phase I trials. *Pharm Stat* 2012; **11**: 258-266.
13. Shi Y, Yin G. Escalation with overdose control for phase I drug-combination trials. *Stat Med* 2013: in press.
14. Riviere MK, Yuan Y, Dubois F, Zohar S. A Bayesian dose-finding design for drug combination clinical trials based on the logistic model. *Pharm Stat* 2014; **13**: 247-257.
15. Mander AP, Sweeting MJ. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Stat Med* 2015; **34**: 1261-1276.
16. 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; **33**: 3815-3829.
17. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: A practical design for phase I clinical trials in cancer. *Biometrics* 1990; **46**: 33-48.
18. Tighiouart M, Piantadosi S, Rogatko A. Dose finding for drug combination in early cancer phase I trials using conditional escalation with overdose control. In JSM Proceedings, Biopharmaceutical section, Alexandria, VA: American Statistical Association. 2014.

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. Babb J, Rogatko A, Zacks S. Cancer Phase I clinical Trials: efficient dose escalation with overdose control. *Stat Med* 1998; **17**: 1103-1120.
21. Tighiouart M, Rogatko A, Babb JS. Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Stat Med* 2005; **24**: 2183-2196.
22. Tighiouart M, Rogatko A. Dose Finding with Escalation with Overdose Control (EWOC) in Cancer Clinical Trials. *Statistical Science* 2010; **25**: 217-226.
23. 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.