



Published in final edited form as:

*Biometrics*. 2010 September ; 66(3): 805–812. doi:10.1111/j.1541-0420.2009.01363.x.

## A Hierarchical Bayesian Design for Phase I Trials of Novel Combinations of Cancer Therapeutic Agents

Thomas M. Braun and Shufang Wang

University of Michigan, Department of Biostatistics, Ann Arbor, MI 48109

### Summary

We propose a hierarchical model for the probability of dose-limiting toxicity (DLT) for combinations of doses of two therapeutic agents. We apply this model to an adaptive Bayesian trial algorithm whose goal is to identify combinations with DLT rates close to a pre-specified target rate. We describe methods for generating prior distributions for the parameters in our model from a basic set of information elicited from clinical investigators. We survey the performance of our algorithm in a series of simulations of a hypothetical trial that examines combinations of four doses of two agents. We also compare the performance of our approach to two existing methods and assess the sensitivity of our approach to the chosen prior distribution.

### Keywords

dose-finding study; dose-escalation study; two-dimensional; adaptive design; Bayesian statistics

## 1. Background and Significance

Phase I trials of combination cancer therapies have been published for a variety of cancer types, including small-cell lung cancer (Rudin et al., 2004), gastric cancer (Inokuchi et al., 2006), melanoma (Azzabi et al., 2005), ovarian cancer (Benepal et al., 2005), and renal cell carcinoma (Amato et al., 2006). Unfortunately, all of these trials, and many others similar to them, suffer from poor study designs that have two distinct limitations. The first limitation is that four out of five of the cited trials escalated doses of only one of the agents, while fixing the dose of the other agent at some pre-determined dose. However, it is very possible that the safest dose of each agent will depend upon which dose of the other agent is used. In an ideal design, simultaneous modification of doses for both agents will be possible. The second limitation is that the cited trials used a variant of the so-called ‘3+3 design’ (Storer, 1989), which has been shown to have poor operating characteristics, including a strong propensity to select doses below the actual MTD (Ahn, 1998; Lin and Shih, 2001). A preferred design would incorporate a parametric model describing how doses of both agents contribute to the probability of dose-limiting toxicity (DLT).

Research into parametric models for two-agent combinations has been ongoing for over forty years, starting with the work of Plackett and Hewlett (1967). However, this research was primarily theoretic and was not specifically motivated by dose-finding studies. In the past decade, a handful of dual-agent Phase I study designs have been published. Kramar et al. (1999) first noted the limitations of an algorithmic approach and instead employed the maximum-likelihood version (O’Quigley and Shen, 1996) of the continual reassessment method (CRM) of O’Quigley et al. (1990). However, as the CRM is an adaptive design for single-agent Phase I trials, Kramar and colleagues then developed an ad-hoc formula allowing them to combine a dose of each agent into a single imputed dose in an effort to provide an ordering for the various dose combinations under study. A related approach was

proposed by Conaway et al. (2004), in which the ordering restrictions were more formally incorporated into parameter estimation and selection of the optimal combination. A very recent CRM-based design was proposed by Yuan and Yin (2008).

Two model-based, adaptive approaches that reflect the individual contributions of both agents also exist. Thall et al. (2003) proposed a design that identifies an entire “contour” of combinations by modeling the probability of DLT as a function of both doses using a six-parameter logistic regression model. One unique aspect of Thall et al. (2003) is that the design first studies specific combinations of two agents, and at the occurrence of the first DLT, the study is widened to examine a continuum of doses in a neighborhood of the combination in which the first DLT occurred. An alternate two-stage design was developed by Wang and Ivanova (2005), in which the first stage is viewed as a “start-up” for the study when little data is available for parameter estimation. Once the first stage has developed a sufficient toxicity profile of the combinations using algorithmic approaches, those combinations deemed “acceptable” in the first stage are more fully examined with adaptive, Bayesian approaches applied to their proposed parametric model. The design of Wang and Ivanova designates one agent as primary (referred to as “dose 1”) and seeks one suitable dose of the secondary agent for every dose of the primary agent that leads to acceptable DLT rates.

In this manuscript, we propose another approach that has distinct differences from the designs of Thall et al. and Wang and Ivanova. First, patients will be enrolled continuously in a single stage, with the parametric model and corresponding Bayesian methods applied throughout the entire study. To limit the selected prior distribution’s influence early in the study, we will implement a stopping rule based solely upon the cumulative number of DLTs observed in the study. Second, in an approach novel to dose-finding study designs, our design does not assume that the probability of DLT is a fixed quantity for every subject receiving the same combination and we choose to model the effects of each agent on the parameters describing the distribution of DLT probabilities for each combination. As a result, our design will accommodate subject heterogeneity better than the competing designs that assume the probability of DLT is constant for each combination. In Section 2, we describe our hierarchical model and the resulting likelihood, and in Section 3, we describe how to use elicited clinical information to develop appropriate prior distributions for each model parameter. Section 4 outlines the actual design and conduct of an actual trial, and Section 5 numerically examines the performance of our design in a variety of settings, and includes a direct comparison to the CRM and method of Wang and Ivanova. We conclude with summarizing remarks in Section 6.

## 2. Proposed Hierarchical Model

We have a Phase I study designed to examine combinations of  $m$  doses of Agent A, denoted  $a_1 < a_2 < \dots < a_m$ , and  $n$  doses of Agent B, denoted  $b_1 < b_2 < \dots < b_n$ . Let  $(j, k)$  represent the combination of dose  $a_j$ ,  $j = 1, 2, \dots, m$ , and dose  $b_k$ ,  $k = 1, 2, \dots, n$ . Note that the value ascribed to each  $a_j$  and  $b_k$  will not be the actual clinical values of the doses, but will be “effective” dose values that will lend stability to our dose-toxicity model, an approach used in numerous Phase I trial designs. We will describe how to determine reasonable effective dose values in Section 3.

We let  $p_{jk}$  denote the probability of a DLT for a subject receiving combination  $(j, k)$ , and we let  $p^*$  denote the desired DLT rate of the optimal combination. Thus, the combination  $(j, k)$  whose corresponding probability of DLT is closest to  $p^*$  should be selected as the maximum tolerated combination (MTC); we denote this combination as  $(j^*, k^*)$ . We will assume that each  $p_{jk}$  has a beta distribution with parameters  $\alpha_{jk}$  and  $\beta_{jk}$ , which is a useful probability

model for our setting because  $\alpha_{jk}$  ( $\beta_{jk}$ ) can be interpreted as the prior number of subjects assigned to combination  $(j, k)$  expected to have (not have) a DLT. It is also natural to incorporate variability into the DLT rates of each combination as dose-finding studies tend to enroll very heterogeneous samples of subjects. As we would expect  $\alpha_{jk}$  ( $\beta_{jk}$ ) to increase (decrease) with both  $a_j$  and  $b_k$ , we explicitly model  $\alpha_{jk}$  and  $\beta_{jk}$  with the following parametric functions of  $a_j$  and  $b_k$ :

$$\log\{\alpha_{jk}(\theta)\}=\theta_0+\theta_1a_j+\theta_2b_k \quad (1)$$

$$\log\{\beta_{jk}(\varphi)\}=\varphi_0-\varphi_1a_j-\varphi_2b_k, \quad (2)$$

in which  $\theta = \{\theta_0, \theta_1, \theta_2\}$  has a multivariate normal distribution with mean  $\boldsymbol{\mu} = \{\mu_0, \mu_1, \mu_2\}$ ,  $\varphi = \{\varphi_0, \varphi_1, \varphi_2\}$  has a multivariate normal distribution with mean  $\boldsymbol{\omega} = \{\omega_0, \omega_1, \omega_2\}$ , and both  $\theta$  and  $\varphi$  have variance  $\sigma^2\mathcal{Q}_3$ , in which  $\mathcal{Q}_3$  is a  $3 \times 3$  identity matrix.

We have chosen to omit the interactive effects of doses. This decision is based, in part, from the work of Wang and Ivanova (2005), who found that the no interaction version of their model performed better than the interaction version, at least in the settings examined. Furthermore, we do not seek to correctly model the entire dose-response curve for every combination, but only those in a neighborhood of the actual MTC, and it has been documented that underparameterized models can provide adequate local fit sufficient for dose-finding studies (O'Quigley and Paoletti, 2003). Nonetheless, our model could be easily generalized to include an interaction term in each of Equations (1) and (2), although additional parameters will necessarily increase the sample size required for the study. We have also assumed that all six regression parameters are independent and have the same prior variance  $\sigma^2$ ; either assumption could be relaxed by incorporating additional variance and/or covariance parameters. However, we have found in simulations (results not shown) that additional parameters add needless complexity to the model with little gain in algorithm performance.

Let  $N_{jk}$  denote the number of subjects assigned to combination  $(j, k)$ , of whom  $Y_{jk}$  subjects have experienced a DLT. If we then define  $\mathbf{Y} = \{Y_{jk} : j = 1, 2, \dots, m; k = 1, 2, \dots, n\}$  and  $\mathbf{N} = \{N_{jk} : j = 1, 2, \dots, m; k = 1, 2, \dots, n\}$ , the posterior distribution for  $(\theta, \varphi)$  is

$$f(\theta, \varphi | \mathbf{Y}, \mathbf{N}) = \frac{\prod_j \prod_k f_{jk}(Y_{jk} | \theta, \varphi, N_{jk}) g(\theta) h(\varphi)}{\int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} \prod_j \prod_k f_{jk}(Y_{jk} | \theta, \varphi, N_{jk}) g(\theta) h(\varphi) d\theta d\varphi}$$

where  $g(\cdot)$  and  $h(\cdot)$  are the respective multivariate normal priors for  $\theta$  and  $\varphi$  described earlier. Although a closed-form expression for the posterior distribution is not available, samples from the posterior distribution are easily obtained using Markov Chain Monte Carlo (MCMC) methods (Robert and Casella, 1999). These samples lead to posterior distributions for each element of  $\theta$  and  $\varphi$ , which in turn lead to a posterior distribution for each  $p_{jk}$ . The corresponding posterior means,  $\tilde{p}_{jk}$ , will then be used to determine which combination is deemed the current estimate of the MTC as more formally explained in Section 4.

### 3. Developing Priors and Effective Dose Values

In order to identify appropriate hyperparameter and effective dose values, we need the investigator to supply  $\tilde{p}_{j1}$  and  $\tilde{p}_{1k}$ , the respective *a priori* values for  $E\{p_{j1}\}$  and  $E\{p_{1k}\}$ , for  $j$

$= 1, 2, \dots, m$  and  $k = 1, 2, \dots, n$ . Thus, at a minimum, we expect the investigator to have sufficient historical information regarding the toxicity profiles of the doses of both agents when combined with the lowest dose of the other agent. To simplify computation, we set the lowest dose of each agent to zero ( $a_1 = b_1 = 0$ ). As a result,  $\log(\alpha_{11}) = \theta_0$  and  $\log(\beta_{11}) = \varphi_0$  so that  $\theta_0$  and  $\varphi_0$  describe the expected numbers of DLTs for the combination using the lowest dose of each agent, and the remaining parameters in Equations (1) and (2) will describe how the expected numbers of DLTs for the other combinations differ from combination (1, 1).

We then use the fact that

$$\frac{K\tilde{p}_{11}}{K(1 - \tilde{p}_{11})} = \frac{\alpha_{11}}{\beta_{11}} = \frac{\exp\{\theta_0\}}{\exp\{\varphi_0\}} = \frac{\exp\{\mu_0\}}{\exp\{\omega_0\}},$$

leading to the solutions  $\mu_0 = \log(K\tilde{p}_{11})$  and  $\omega_0 = \log(K[1 - \tilde{p}_{11}])$ , where  $K = 1000$  was chosen as a scaling factor to keep both hyperparameters sufficiently above zero. We have selected normal priors for our regression parameters, and have therefore allowed for the possibility of a decreased probability of DLT with increasing doses of either agent, although we do expect there to be low probability of such an occurrence. As a result, we have selected the values  $\mu_1 = \mu_2 = \omega_1 = \omega_2 = 2\sqrt{\sigma^2}$  so that 97.5% of the prior distributions for  $\theta_1, \theta_2, \varphi_1$ , and  $\varphi_2$  will lie above zero, depending upon the value of  $\sigma^2$ .

Given these values, we define the elicited odds ratios

$$\begin{aligned} \widetilde{OR}_j &= \exp\left\{\frac{\tilde{p}_{j1}/[1-\tilde{p}_{j1}]}{\tilde{p}_{11}/[1-\tilde{p}_{11}]}\right\} = \exp\{(\theta_1 + \varphi_1)a_j\} \\ &\approx \exp\{(\mu_1 + \omega_1)a_j\} \end{aligned}$$

and

$$\begin{aligned} \widetilde{OR}_{,k} &= \exp\left\{\frac{\tilde{p}_{1k}/[1-\tilde{p}_{1k}]}{\tilde{p}_{11}/[1-\tilde{p}_{11}]}\right\} = \exp\{(\theta_2 + \varphi_2)b_k\} \\ &\approx \exp\{(\mu_2 + \omega_2)b_k\}. \end{aligned}$$

As a result, our effective dose values are  $a_j = (\mu_1 + \omega_1)^{-1} \log(\widetilde{OR}_j)$  and  $b_k = (\mu_2 + \omega_2)^{-1} \log(\widetilde{OR}_{,k})$ , meaning all doses are rescaled to be proportional to log-odds ratios relative to combination (1, 1). Because the elicited probabilities of toxicity will increase with an increase in the dose of either Agent A or Agent B, we will have  $a_1 < a_2 < \dots < a_m$  and  $b_1 < b_2 < \dots < b_n$  as desired.

All of the above computations require a value for  $\sigma^2$ , which we identify through a grid search of candidate values. For each candidate value, a series of small simulation studies should be performed to assess the performance of the algorithm when the prior means are correctly specified as well as situations when the prior means are too high or too low. An appropriate value of  $\sigma^2$  is one that is small enough so that the prior is sufficiently informative when there is limited data at the beginning of a trial, but large enough so that the prior becomes sufficiently non-informative when there is enough data later in the trial.

Although each trial setting will require fine-tuning of  $\sigma^2$ , we have found in our settings that values of  $\sigma^2$  in the interval [5, 10] are often sufficient. Note that because Equations (1) and (2) are additive in both doses, the variance of  $p_{jk}$  necessarily grows with  $a_j$  and  $b_k$ . We examined situations in which we scaled  $\sigma^2$  by  $a_j$  or  $b_k$  as an attempt to keep the variance stable among all combinations. However, we found that the algorithm did no better with this variance stabilization than it did without it. Furthermore, the non-constant variance model has practical suitability, as it is plausible that there is more certainty about prior values of  $p_{jk}$  for combinations of lower dose values than for combinations of higher dose values.

#### 4. Trial Design and Conduct

Before the trial begins enrolling subjects, the investigators should first specify: (a) the  $(m+n-1)$  prior probabilities of DLT for combinations containing the lowest dose of Agent A and/or Agent B, and (b) the targeted probability of DLT,  $p^*$ . From this information, the study statistician can determine values for all hyperparameters and effective doses as described in Section 3. At this point, subject  $i = 1$  can be enrolled and assigned to combination  $(j_1, k_1) = (1, 1)$ . In order to determine the dose assignment  $(j_i, k_i)$  for each subject  $i = 2, 3, \dots, N$ , the conduct of the trial proceeds as follows:

1. Compute a 95% confidence interval for the overall DLT rate among all combinations using the cumulative number of observed DLTs for subjects 1, 2, ...,  $(i-1)$ .
2. If the lower bound of the confidence interval from step (1) is greater than  $p^*$ , terminate the trial.
3. If the lower bound of the confidence interval from step (1) is no more than  $p^*$ :
  - a. Use the outcomes and assignments of subjects 1, 2, ...,  $(i-1)$  to determine the posterior distribution of each  $p_{jk}$ , with posterior mean  $\bar{p}_{jk}$ , as described in Section 2;
  - b. Define the set  $S = \{(j, k) : j_{i-1} - 1 \leq j \leq j_{i-1} + 1, k_{i-1} - 1 \leq k \leq k_{i-1} + 1\}$  that contains combinations that are within one dose level of the corresponding doses in the combination assigned to the most recently enrolled subject;
  - c. Identify the combination  $(j^*, k^*)$  in  $S$  as the one with smallest  $d_{jk} = |\bar{p}_{jk} - p^*|$ ;
  - d. Assign subject  $i$  to  $(j_i, k_i) = (j^*, k^*)$ .
4. If all  $N$  subjects have been enrolled and followed, repeat steps (3a)-(3c) with the data of all  $N$  subjects to identify the MTC.

Recall that our design enrolls all subjects in a single stage, rather than in two-stages. Although some authors feel that the first stage is needed to limit the influence of the prior when little data has been collected, we feel that a well-designed sensitivity analysis of  $\sigma^2$  prior to inception of the study will serve to limit the influence of the prior in the early portion of the study. Also, step (1) in our trial is completely data-driven and allows for early termination of the study without input from the prior distribution. We did examine a stopping rule based upon the percentage of the posterior distributions of each  $p_{jk}$  above  $p^*$  but found this stopping rule did not lead to study termination often enough if all of the combinations were overly toxic. The width of the confidence interval used in step (1) can also be modified to increase or decrease the probability of early termination.

Like most dose-finding studies, we enforce a “do not skip” rule in step (3b) of our conduct; however, we apply this rule to escalation as well as to de-escalation. Specifically, subject  $i$  must be assigned to a combination that is in a “close neighborhood”, as defined in step (3b), of the combination assigned to subject  $(i - 1)$ . By doing so, we hope to limit the number of patients exposed to overly toxic combinations while still allowing the algorithm to fully search the entire grid of possible combinations. We examined putting no limit on de-escalation (as is done in most single-agent studies) but found that our algorithm then had difficulty identifying the optimal combination if it existed at high-dose combinations of both agents, as a single DLT had large influence on the posterior distributions of the  $p_{jk}$ .

Our approach allows for simultaneous dose escalations of both agents, which may appear to be overly aggressive and increase the likelihood of exposing too many subjects to overly toxic combinations. However, we cite the work of Wang and Ivanova (2005), as well as that of Braun et al. (2007) in the setting of simultaneous dose/schedule finding, who found escalation of both dimensions simultaneously does not lead to an increased observed toxicity rate. Furthermore, allowing simultaneous escalations in both doses increases the ability of the algorithm to fully explore all combinations in a neighborhood of the actual MTC when the MTC is a combination of high doses for both agents. Note that escalation could be slowed by requiring that a cohort of  $M$  patients be assigned to the same combination before escalation can be considered. However, we have found that using  $M = 1$  is sufficient for limiting escalation and using values of  $M > 1$  will treat too many subjects at sub-optimal combinations and unnecessarily increase the overall sample size of the study when the MTC is a combination of high doses of both agents.

## 5. Numerical Studies

### 5.1 Operating Characteristics

We examine the performance of our algorithm in six scenarios (A through F) for a hypothetical clinical trial of  $N = 35$  subjects, where the sample size was selected for feasibility yet also satisfactory operating characteristics across all six scenarios. We also examined sample sizes of  $N = 40$  and  $N = 50$  in small simulation studies and found little improvement in the operating characteristics seen with  $N = 35$ . The trial is designed to determine which of four doses of Agent A ( $m = 4$ ) and which of four doses of Agent B ( $n = 4$ ), when given in combination, lead to a DLT rate close to  $p^* = 0.20$ . The actual DLT rates for each combination under each scenario are displayed in the first four columns of values in Table 1. Scenario A has an abundance of tolerable combinations at higher dose levels of one or both agents, while Scenario B has all combinations with DLT probabilities under the target  $p^*$ . Scenario C has very few tolerable combinations and Scenario D has no tolerable combinations. Scenarios E and F are settings in which the probability of DLT increases little with dose increases of one agent yet increases steeply with dose increases of the other agent, increasing the difficulty of identifying a single best choice for the MTC.

The *a priori* rates of DLT elicited from investigators are equal to the actual DLT rates in Scenario A. Based upon these values, the investigators believe combination (2, 3) is optimal as it has a DLT rate equal to that desired. Using the methods described in Section 3, we produce the effective dose values  $a_1 = 0.000$ ,  $a_2 = 0.058$ ,  $a_3 = 0.094$ , and  $a_4 = 0.120$  for Agent A and  $b_1 = 0.000$ ,  $b_2 = 0.078$ ,  $b_3 = 0.120$ , and  $b_4 = 0.151$  for Agent B and hyperparameter values  $\sigma^2 = 10$ ,  $\mu = \{3.69, 6.32, 6.32\}$ , and  $\omega = \{6.87, 6.32, 6.32\}$ . The value  $\sigma^2 = 10$  was selected from a grid search of values 1, 2, ... 12. Specifically, we performed small simulation studies of 100 replications with each possible value of  $\sigma^2$  in each of the six scenarios and found that  $\sigma^2 = 10$  lead to a prior distribution that was sufficiently informative during the early portion of a study and also allowed the data to dominate during the latter portion of a study across all scenarios.



We ran 1000 simulations of our algorithm under each scenario; the performance of our algorithm is summarized in the final eight columns of Table 1. The first four of the eight columns display the percentage of simulations in which each combination was identified as the MTC at the end of the study, and the last four of the eight columns display the average percentage of 35 subjects among all simulations that were assigned to each combination. All simulations were done in the statistical package *R*. 2000 draws from the posterior distributions of each  $\alpha_{jk}$  and  $\beta_{jk}$  were generated with the MCMC machinery supplied in the JAGS (Just Another Gibbs Sampler) library. User-friendly *R* code is available for applying our methods to both simulation studies and actual trials; the code can be downloaded at [www.sph.umich.edu/~tombraun/software.html](http://www.sph.umich.edu/~tombraun/software.html).

As we discuss the results, we note that although we wish to target combinations whose DLT probability is 0.20, selecting combinations whose DLT probabilities are within 10 points of 0.20 (which we call the “10-point window”) still indicates satisfactory performance of our algorithm, as: (1) a sample size of 35 subjects is insufficient for discrimination between DLT rates in the interval [0.10, 0.30], and (2) in most clinical applications, the desired DLT rate is often a rough guess and finding combinations with DLT rates close to the target rate selected by the investigator will still prove of interest to the investigator. A narrower window could certainly be used. However, a level of performance with this narrower window equivalent to that we observed with our 10-point window would directly increase the sample size chosen for the study.

In Scenario A our algorithm selects combinations within the 10-point window as the MTC in 89% of simulations and assigns an average of 76% of subjects to those combinations. Furthermore, within the 10-point window, our algorithm is more likely to select combinations and assign patients to those combinations with DLT rates within four points of 0.20. In Scenario B, we see combinations including doses 3 and 4 of either agent are selected as the MTC and assigned to subjects more often than they were in Scenario A, reflecting the increased safety of those combinations that would be demonstrated in the data. Conversely in Scenario C, we see combinations including doses 1 and 2 of either agent are selected as the MTC and assigned to subjects more often than they were in Scenario A, reflecting the increased toxicity of those combinations that would be demonstrated in the data. Scenarios A, B, and C demonstrate that the prior we selected for  $\theta$  and  $\phi$  is sufficiently non-informative so as to allow the algorithm to “move with the data.” Early termination occurred in 1%, 0%, and 16% of simulations in Scenarios A-C, respectively; the increased early termination rate in Scenario C is due to a majority of combinations being unacceptably toxic.

In Scenario D, where no combinations were tolerable, we see that an MTC was very rarely identified, with only 5% of simulations fully enrolling 35 patients and identifying the MTC at combination (1, 1), which was also the combination most frequently assigned to subjects. From the results presented in Table 1 for scenarios E and F, we see that our algorithm is able to adjust for distinctly differential dose-toxicity patterns of the two agents and still identify the MTC at combinations in a neighborhood of the true MTC and assign a preponderance of subjects to those combinations. However, we do see in Scenario F that the algorithm selects combinations with dose 3 of Agent B more often than desired. This result is due to the dramatic increase in DLT rates between doses 2 and 3 of Agent B that does not fit our assumed model in Equations (1) and (2). Our model underestimates the toxicity probabilities of combinations including doses 3 and 4 of Agent B, making them more likely to be identified as the MTC and assigned to subjects more than desired. Nonetheless, the dose-toxicity pattern of Agent B in Scenario F would challenge any parametric model applied to this setting. Early termination occurred in 11% and 12% of simulations in Scenarios E and F, respectively. One may argue that since there was at least one optimal combination in

Scenarios C, E, and F, it is not desirable to observe early termination in those scenarios. However, in order to have a design that stops early with high probability when all combinations are overly toxic, we have to allow for some early termination when very few optimal combinations exist. As stated earlier, we examined stopping rules based upon the posterior distributions of the parameters rather than the actual number of observed DLTs, but found those stopping rules were unable to stop the study soon enough when all combinations were overly toxic unless the prior was skewed toward having all combinations being overly toxic.

As a final summary of the performance of our algorithm, Figure 1 displays the average assignment of each subject across all 1000 simulations in terms of the actual DLT rates of the combinations. For example, in Scenario A, if two subjects were assigned to combinations (1, 2) and (2, 1), their respective assignments were given values of 0.10 and 0.08. All subjects not receiving an assignment due to early termination were given a value of 0.00. The hope is that as more and more subjects are enrolled, the average combination assigned to the final subject has a DLT rate close to the target of  $p^* = 0.20$  (the horizontal line in Figure 1). In all scenarios except Scenario D, we see that the final subject is assigned on average to a combination whose DLT rate is quite close to the targeted DLT rate. The pattern for Scenario D reflects the fact that latter subjects tended to receive no assignment and the average DLT rate tends toward 0.00.

## 5.2 Comparison to Alternate Methods

We also compared the performance of our approach to the performance of the CRM and the approach of Wang and Ivanova (2005) in all six scenarios. Table 2 displays results for Scenarios A, B and E; Scenarios C, D, and F are omitted, as the results for each were similar to the results for one of the presented scenarios. For the approach of Wang and Ivanova, we selected Agent A as “dose 1” and applied the model, prior distributions, and study design exactly as described in their manuscript, except that in the second stage, we enrolled patients in cohorts of size one rather than three. For the CRM, we ordered the 16 combinations by their prior DLT rates (the actual DLT rates in Scenario A). For these 16 “doses”,  $d_\ell$ ,  $\ell = 1, 2, \dots, 16$ , we used the model  $p_\ell = \tilde{d}_\ell^\zeta$ , where  $p_\ell$  is the probability of DLT for “dose”  $\ell$  and  $\zeta$  had an exponential distribution with mean one. Note that combinations (1, 3) and (4, 1) had equal prior DLT rates and thus equal ordering for the CRM. If those combinations were selected by the CRM as the MTC, the algorithm was programmed to choose one of the combinations with equal probability as the assignment for the next subject. The same approach was used with combinations (4, 2) and (1, 4).

In Scenarios A and B, we see fairly comparable performance between our proposed method and the CRM. However, in scenario A, the CRM assigns 80% of subjects to combinations in the 10-point window, compared to 76% for our method, and in Scenario B, our method selects and assigns combination (4, 4) much more often than the CRM does. The method of Wang and Ivanova has noticeably poorer performance in Scenario A, tending to select and assign combinations that include dose 3 or 4 of both agents. This result is likely due to the prior distribution for each parameter suggested by the authors (exponential with mean one). It is possible that changing the exponential distribution to a gamma distribution with mean one but flexible variance would allow for a less informative prior distribution and move the algorithm away from the higher dose combinations. This preference for higher dose combinations also leads to improved performance in Scenario B over our method and the CRM.

Scenario E demonstrates the primary limitation of using the CRM, which was designed for studies of a single agent, rather than a design specifically created to examine combinations



of two agents. The forced ordering of the combinations reflected in the prior DLT rates is drastically different from the true ordering of the combinations. Thus, the CRM tends to focus upon combinations that include the lowest two doses of either agent, as opposed to our method, which is able to more often select and assign combinations that include the middle two doses of either agent.

### 5.3 Assessing Influence of Prior

In Table 3, we present results for Scenario A using six prior distributions with means and/or variance different from that used for the simulations presented in Table 1. The prior DLT rates used in Scenarios A1, A3, and A5 were the actual DLT rates of Scenario F shown in Table 1 and the prior DLT rates used in Scenarios A2, A4, and A6 were the actual DLT rates of Scenario B. The prior variance in Scenarios A1 and A2 equals that for the simulations of Section 5.1 ( $\sigma^2 = 10$ ), while Scenarios A3 and A4 have a smaller prior variance ( $\sigma^2 = 5$ ) and Scenarios A5 and A6 have a larger prior variance ( $\sigma^2 = 15$ ). By comparing the results in Table 3 to the results of Scenario A in Table 1, we see there are slight variations in which combinations are selected as the MTC. For example, in Scenarios A1, A3, and A5, combination (4,2) is selected as the MTC more often than it was before. However, patient assignments and the algorithm's ability to identify combinations within the 10-point window as the MTC are generally unaffected by the choice of prior.


## 6. Conclusion

Our manuscript has proposed a Bayesian hierarchical design for identifying doses of two agents with optimal DLT rates when given in combination. The important aspects of our design include the sequential enrollment of subjects in a single stage, an explicit procedure for developing model parameter prior distributions from a set of DLT rates elicited from investigators, and a design that formally accommodates patient heterogeneity. Future planned extensions of our model include the incorporation of weights into our likelihood to allow for decision-making with partial follow-up of currently enrolled subjects, thus allowing subjects to be enrolled as soon as they are eligible and shortening the duration of the trial. We also want to generalize our model to accommodate the actual administration times of the two agents, as many dual-agent therapies do not administer both agents simultaneously and the length of delay of the second agent may be important when assessing the DLT rate of the combination.

## References

- Ahn C. An evaluation of Phase I cancer clinical trials. *Statistics in Medicine*. 1998; 17:1537–1549. [PubMed: 9699228]
- Amato RJ, Morgan M, Rawat A. Phase I/II study of thalidomide in combination with interleukin-2 in patients with metastatic renal cell carcinoma. *Cancer*. 2006; 106:1498–1506. [PubMed: 16475152]
- Azzabi A, Hughes AN, Calvert PM, Plummer ER, Todd R, Griffin MJ, Lind MJ, Maraveyas A, Kelly C, Fishwick K, Calvert AH, Boddy AV. Phase I study of temozolomide plus paclitaxel in patients with advanced malignant melanoma and associated in vitro investigations. *British Journal of Cancer*. 2005; 92:1006–1012. [PubMed: 15756276]
- Benepal T, Jackman A, Pyle L, Bate S, Hardcastle A, Aherne W, Mitchell F, Simmons L, Ruddle R, Raynaud F, Gore M. A Phase I pharmacokinetic and pharmacodynamic study of BGC9331 and carboplatin in relapsed gynaecological malignancies. *British Journal of Cancer*. 2005; 93:868–875. [PubMed: 16222320]
- Braun TM, Thall PF, Nguyen H, de Lima M. Simultaneously optimizing dose and schedule of a new cytotoxic agent. *Clinical Trials*. 2007; 4:113–124. [PubMed: 17456511]
- Conaway MR, Dunbar S, Peddada SD. Designs for single- or multiple-agent Phase I trials. *Biometrics*. 2004; 60:661–669. [PubMed: 15339288]

- Inokuchi M, Yamashita T, Yamada H, Kojima K, Ichikawa W, Nihei Z, Kawano T, Sugihara K. Phase I/II study of S-1 combined with irinotecan for metastatic advanced gastric cancer. *British Journal of Cancer*. 2006; 94:1130–1135. [PubMed: 16570038]
- Kramar A, Lebecq A, Candalh E. Continual reassessment methods in Phase I trials of the combination of two drugs in oncology. *Statistics in Medicine*. 1999; 18:1849–1864. [PubMed: 10407256]
- Lin Y, Shih WJ. Statistical properties of the traditional algorithm-based designs for Phase I cancer clinical trials. *Biostatistics*. 2001; 2:203–215. [PubMed: 12933550]
- O'Quigley J, Paoletti X. Continual reassessment method for ordered groups. *Biometrics*. 2003; 59:430–440. [PubMed: 12926728]
- 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. [PubMed: 2350571]
- O'Quigley J, Shen L. Continual reassessment method: A likelihood approach. *Biometrics*. 1996; 52:163–172.
- Plackett RL, Hewlett PS. A comparison of two approaches to the construction of models for quantal responses to mixtures of drugs. *Biometrics*. 1967; 23:27–44. [PubMed: 6050470]
- Robert, CP.; Casella, G. *Monte Carlo Statistical Methods*. New York: Springer; 1999.
- Rudin CM, Kozloff M, Hoffman PC, Edelman M, Karnauskas R, Tomek R, Szeto L, Vokes EE. Phase I study of G3139, a bcl-2 antisense oligonucleotide, combined with carboplatin and etoposide in patients with small-cell lung cancer. *Journal of Clinical Oncology*. 2004; 22:1110–1117. [PubMed: 15020613]
- Storer BE. Design and analysis of Phase I clinical trials. *Biometrics*. 1989; 45:925–937. [PubMed: 2790129]
- Thall PF, Millikan RE, Mueller P, Lee S-J. Dose-finding with two agents in Phase I oncology trials. *Biometrics*. 2003; 59:487–496. [PubMed: 14601749]
- Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics*. 2005; 61:217–222. [PubMed: 15737096]
- Yuan Y, Yin G. Sequential continual reassessment method for two-dimensional dose finding. *Statistics in Medicine*. 2008; 27:5664–5678. [PubMed: 18618901]



**Figure 1.**  
Pattern of dose assignments.

**Table 1**

Summary of true DLT probabilities and simulation results for Scenarios A through F. All probabilities and percentages are multiplied by 100. Values in boldface correspond to combinations with a true DLT probability within 10 points of the target probability  $p^* = 0.20$ . In all six settings, the prior probabilities of DLT for each combination equal the actual probabilities of DLT in Scenario A, leading to hyperparameter values  $\mu = \{3.69, 6.32, 6.32\}$ ,  $\omega = \{6.87, 6.32, 6.32\}$ , and  $\sigma^2 = 10$ .

Scenario	Dose of Agent B	True DLT Probability															
		% of Simulations Selected as MTC															
		Mean % of Subjects Assigned															
		Dose of Agent A															
		1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
A	1	4	8	12	16	0	2	2	4	4	6	3	3	2			
	2	10	14	18	22	3	9	11	7	4	10	8	4				
	3	16	20	24	28	6	12	13	7	5	8	13	6				
	4	22	26	30	34	5	6	6	6	3	4	6	15				
B	1	2	4	6	8	0	0	0	0	4	1	1	1				
	2	5	7	9	11	0	1	2	2	1	5	3	2				
	3	8	10	12	14	1	2	7	8	1	3	10	6				
	4	11	13	15	17	1	2	11	61	1	2	7	52				
C	1	10	20	30	40	15	27	9	3	23	21	8	3				
	2	25	35	45	55	20	3	0	0	14	9	2	1				
	3	40	50	60	70	6	0	0	0	6	2	3	0				
	4	55	65	75	85	0	0	0	0	1	0	0	1				
D	1	44	48	52	56	5	0	0	0	27	4	1	0				
	2	50	54	58	62	0	0	0	0	2	3	0	0				
	3	56	60	64	68	0	0	0	0	1	0	1	0				
	4	62	66	70	74	0	0	0	0	0	0	0	0				
E	1	8	18	28	38	2	9	6	5	10	8	5	3				
	2	9	19	29	39	8	11	11	3	6	12	7	3				
	3	10	20	30	40	6	12	6	2	5	8	9	3				
	4	11	21	31	41	5	4	3	2	4	3	3	8				
F	1	12	13	14	15	6	8	9	7	14	9	7	5				
	2	16	18	20	22	16	14	6	1	11	14	7	2				
	3	44	45	46	47	11	4	1	0	8	5	5	1				

Scenario	True DLT Probability				% of Simulations Selected as MTC				Mean % of Subjects Assigned			
	1	2	3	4	1	2	3	4	1	2	3	4
4	50	52	54	55	2	0	0	0	2	1	1	3

**Table 2**

Comparison of proposed design to CRM and design of Wang & Ivanova. All percentages are multiplied by 100. Values in boldface correspond to combinations with a true DLT probability within 10 points of the target probability  $p^* = 0.20$ .

Scenario	Dose of Agent B	Dose of Agent A												
		Braun				CRM				Wang & Ivanova				
<i>% of Simulations Selected as MTC</i>														
A	1	0	2	2	4	0	2	7	6	0	0	0	3	
	2	3	9	11	7	5	11	13	4	0	0	2	8	
	3	6	12	13	7	5	11	10	6	7	7	21	15	
	4	5	6	6	6	5	7	4	4	2	5	11	16	
B	1	0	0	0	0	0	0	1	1	0	0	0	0	
	2	0	1	2	2	0	1	3	2	0	0	0	0	
	3	1	2	7	8	1	4	7	8	0	0	1	4	
E	4	1	2	11	61	2	8	12	50	0	2	10	82	
	1	2	9	6	5	8	15	12	4	0	0	8	6	
	2	8	11	11	3	17	10	6	3	0	1	14	8	
	3	6	12	6	2	4	6	4	2	6	11	19	7	
A	4	5	4	3	2	2	4	1	1	2	6	6	4	
	<i>Mean % of Subjects Assigned</i>													
	1	6	3	3	2	7	9	9	6	6	5	5	9	
	2	4	10	8	4	8	10	9	4	2	4	3	5	
B	3	5	8	13	6	6	8	5	3	4	4	10	8	
	4	3	4	6	15	5	4	3	3	1	5	9	20	
	1	4	1	1	1	4	4	5	4	6	5	5	11	
	2	1	5	3	2	4	5	5	4	1	3	2	2	
E	3	1	3	10	6	4	6	6	6	1	1	2	5	
	4	1	2	7	52	4	6	9	24	0	2	5	50	
	1	10	8	5	3	18	17	11	4	6	5	5	5	
	2	6	12	7	3	12	8	5	2	4	5	8	6	
A	3	5	8	9	3	4	4	3	2	5	7	11	6	



	Braun				CRM				Wang & Ivanova				
	Dose of Agent A												
Scenario	Dose of Agent B	1	2	3	4	1	2	3	4	1	2	3	4
	4	4	3	3	8	3	3	3	2	1	2	7	8

**Table 3**

Assessing the impact of the chosen prior distribution on the operating characteristics presented for Scenario A in Table 1. All percentages are multiplied by 100. All combinations have a true DLT probability within 10 points of the target probability  $p^* = 0.20$ , except combinations (1,1), (2,1), and (4,4).

Scenario	Prior Parameters		% of Simulations Selected as MTC										Mean % of Subjects Assigned						
	$\mu$	$\omega$	$\sigma^2$	Dose of Agent B		Dose of Agent A						Dose of Agent A							
				1	2	1	2	3	4	5	6	7	8	9	10	1	2	3	4
A1	4.79	6.78	10	1	2	2	2	4	3	8	3	3	3	3	2				
	6.32	6.32		2	1	5	10	25	2	8	7	18							
	6.32	6.32		3	16	8	7	2	9	7	10	3							
A2				4	2	4	3	5	1	3	3	12							
	3.00	6.98	10	1	0	1	2	3	6	3	3	2							
	6.32	6.32		2	2	6	8	6	3	10	7	4							
	6.32	6.32		3	5	9	10	5	5	8	12	6							
A3				4	4	5	6	5	3	4	6	18							
	4.79	6.78	5	1	1	3	4	3	8	3	3	2							
	4.47	4.47		2	3	7	9	23	2	9	8	17							
	4.47	4.47		3	16	8	6	1	10	7	10	2							
A4				4	2	4	4	4	1	3	4	11							
	3.00	6.98	5	1	0	2	4	2	5	3	2	2							
	4.47	4.47		2	3	9	11	7	4	10	8	4							
	4.47	4.47		3	5	13	12	6	5	9	13	5							
A5				4	6	6	5	7	3	4	6	16							
	4.79	6.78	15	1	2	2	3	4	7	3	2	2							
	7.75	7.75		2	2	5	9	24	2	8	7	17							
	7.75	7.75		3	17	8	8	2	10	7	11	3							
A6				4	2	4	4	5	1	3	4	13							
	3.00	6.98	15	1	0	2	2	3	6	3	2	2							
	7.75	7.75		2	3	11	8	7	3	10	8	4							
	7.75	7.75		3	5	12	14	8	4	8	13	6							
			4	6	6	7	6	3	4	6	18								