Published in final edited form as:

Stat Biopharm Res. 2013 January 1; 5(2): 105–115. doi:10.1080/19466315.2013.767213.

A Generalized Continual Reassessment Method for Two-Agent Phase I Trials

Thomas M. Braun and Nan Jia

University of Michigan School of Public Health 1415 Washington Heights, Ann Arbor, MI 48109, U.S.A.

1 Background

The primary goal of traditional phase I oncology trials is identifying the maximum tolerated dose (MTD), defined as the largest possible dose that maintains the probability of a dose-limiting toxicity (DLT) that is closest to a pre-specified target rate, $\Gamma \in (0, 1)$. Identification of the MTD is usually done in an adaptive fashion, with the outcomes of some or all of the previously enrolled subjects used to determine the dose assignment of the next subject or cohort of subjects. The first model-based adaptive design was the Continual Reassessment Method (CRM) of O'Quigley *et al.* (1990), which was later modified by Faries (1994) and Goodman *et al.* (1995). The CRM assumes that the probability of DLT is related to dose d through a simple parametric function $F(\beta; d)$ that is monotonic in dose, with monotonicity enforced by requiring $\beta > 0$. The two most commonly-used forms for $F(\beta; d)$ are the logistic model $\text{logit}[F(\beta; d)] = a + \beta d$ with known value a , where $\text{logit}(p) = \log(p) - \log(1 - p)$, and the power model $F(\beta; d) = d^\beta$.

Although β can conceptually be estimated using maximum-likelihood theory (O'Quigley and Shen, 1996), the need for heterogeneous outcomes, i.e. at least one patient with a DLT and one patient without a DLT, requires the use of a first stage to determine the dose-assignment of each patient until the first DLT is observed. Alternatively, β can be estimated in the CRM using Bayesian methods without the need for heterogeneous outcomes. A prior distribution is placed on β and the posterior mean of β is sequentially updated as each enrolled subject completes their follow-up. The first patient is usually assigned to the lowest dose, and

patient i , $i = 2, 3, \dots, N$ is assigned to the dose in which $F(\tilde{\beta}_{i-1}; d)$ is closest to Γ , where $\tilde{\beta}_{i-1}$ is the posterior mean of β based upon the data from the first $(i - 1)$ patients. At the end of the study, the MTD is the dose, d^* , for which $F(\tilde{\beta}_N; d^*)$ is closest to Γ .

Many recent Phase I trials have begun to study the combined toxicity of two agents, with multiple doses of both agents being studied. For example, Rowinsky *et al.* (1996) conducted a trial where cisplatin and topotecan with five and two dose levels, respectively, were simultaneously administered to study the combined toxicity of the two drugs. Yuan and Yin (2008) also presented a study for the combination of bortezomib with chemotherapy that examined eight dose levels for each agent. In such studies, identifying a maximum tolerated combination (MTC) cannot be done using single-agent designs. This is because even with an assumption of monotonicity of DLT rates between dose levels in one agent when the other agent is fixed, we still do not have a complete ordering of the probabilities of DLT for all possible dose combinations with the two agents. For instance, in a simple case of two doses each of Agents A and B, we have no knowledge regarding the order of the DLT rates of the combinations of the high dose of Agent A and the low dose of Agent B versus the low dose of Agent A and the high dose of Agent B.

Numerous statistical approaches exist for the design of two-agent Phase I trials. Thall *et al.* (2003) propose a design in which the DLT rate of a combination is modeled as a function of the doses of the two agents via a logistic model. Identification of the MTC is based on a contour map, in which combinations on the same contour have the same DLT rate. Wang and Ivanova (2005) use an alternate parametric model and define a “start-up” stage when little data are available for parameter estimation. After collection of enough data for a sufficient toxicity profile, the dose combinations deemed “acceptable” in the “start-up” stage are more fully examined using the parametric model and Bayesian methods for estimation. Yuan and Yin (2008) define a search range of combinations by first running a one-dimensional trial and later updating it, then eliminating the doses that lay outside of the search range based on the partial ordering of the DLT rates, thereby largely reducing the two-dimensional dose-finding space. Additional adaptive, parametric approaches have been proposed by Yuan and Yin (2009); Yin and Yuan (2009) and Braun and Wang (2010). Most recently, using the methods of partial ordering first suggested in the single-agent setting by Conaway *et al.* (2004), Wages *et al.* (2011a,b) develop designs for two-agent Phase I trials based upon an examination of all possible orders of combinations, when the number of combinations can be feasibly enumerated, or a subset of all possible orders when enumeration of all possible orders is infeasible.

Given the multitude of methods available for the design of two-agent Phase I trials, it is frustrating that none of these methods have been adopted as a standard in actual clinical trials. We posit that a major limitation to the use of most of the designs is the complexity of the models that are used. Although these models have appealing statistical traits, like the use of a hierarchical model in Braun and Wang (2010) and a copula in Yuan and Yin (2009); Yin and Yuan (2009), there is no obvious benefit from the incorporation of complex statistical ideas given that all the methods have similar operating characteristics. Given that the CRM is gradually being adopted as a standard for single-agent Phase I trials, it seems plausible that an approach that mimics the CRM would appeal to investigators and may provide a standard for the design of two-agent trials. Nonetheless, it should be emphasized that all existing methods, including ours, require careful selection of the prior distributions of the model parameters, which is another issue that may explain the limited adoption of existing methods.

To that end, we propose an alternative design of two-agent Phase I trials that is a straightforward generalization of the CRM, which we denote gCRM. We visualize the combinations of m doses of Agent A and n doses of Agent B as a grid of m columns and n rows. We then adopt the CRM separately for each row in the grid, with each CRM specific to a single dose of Agent B in combination with doses of Agent A. We include a parameter β that is shared among all the row-wise models, while each model contains a second parameter (usually fixed in the one-agent CRM) that is specific to the dose of Agent B being studied. These rowspecific parameters are given a prior distribution that reflects their correlation and allows for “borrowing of strength” among all combinations being studied. In Section 2, we explicitly develop our model and describe how to determine hyperparameter values for the prior distributions of the model parameters, as well as rescaled dose values for one of the agents, using information elicited from investigators. We also explicitly outline the conduct of an actual trial using our methods. We examine the operating characteristics of our approach via simulations and compare our approach to other existing approaches in Section 3. We present concluding remarks in Section 4.

2 Methods

2.1 Description of gCRM Model

As previously stated, the logistic model and the power model are commonly used as the dose-toxicity function $F(\beta; d)$ for the traditional CRM. We now generalize the CRM to a Phase I trial of two agents specific to the logistic model; a short description of generalization for the power model appears at the end of Section 2.4. We have a trial of two agents, in which there are m doses, $a_1 < a_2 < \dots < a_m$, of Agent A, and n doses, $b_1 < b_2 < \dots < b_n$, of Agent B. We let (j, k) represent the combination of Agent A at dose a_j , $j = 1, 2, \dots, m$ and Agent B at dose b_k , $k = 1, 2, \dots, n$. We let p_{jk} denote the true DLT probability for (j, k) and define the MTC, (j^*, k^*) , as the combination with a DLT rate closest to the target rate Γ . For each value of k (each dose of Agent B), we adopt a traditional logistic regression model $\text{logit}(p_{jk}) = \alpha_k + \beta a_j$ describing how the probability of DLT varies with each dose of Agent A combined with the given dose of Agent B. Note that the intercept is now a parameter specific to the dose of Agent B under study, rather than a fixed value traditionally used in the CRM; see Figure 1 for a visual representation of this idea. In more general applications, this model is known as a proportional odds logistic regression model. We attempt to “join” the row-wise CRM models through a joint prior distribution that forces correlation among their intercept parameters. We note that because the row-wise CRM models share the same parameter β , the gCRM assumes that no interaction exists between doses of the two agents, although we demonstrate in Section 3 that the operating characteristics of the gCRM are quite insensitive to this assumption. We also note that like the traditional CRM, the dose values $a_1, a_2, \dots, a_m$ of Agent A are not actual clinical dose values, but are rescaled dose values selected to improve model fit. However, there is no need to compute rescaled dose values for Agent B because the doses of Agent B do not contribute directly to our model.

Given that we would like to enforce non-decreasing DLT rates among increasing doses of Agent A, we assume β has a Gamma distribution with mean μ_β and variance σ_β^2 , denoted $\beta \sim Ga(\mu_\beta, \sigma_\beta^2)$ which can be alternately viewed as a Gamma distribution with shape $\mu_\beta^2/\sigma_\beta^2$ and scale σ_β^2/μ_β . We also assume that the intercept for the model at the lowest dose of Agent B has a Normal distribution with mean μ_α and variance σ_α^2 , denoted $\alpha_1 \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2)$. Instead of directly modeling the joint distribution of the intercepts, we focus on the difference between intercepts corresponding to adjacent doses of Agent B.

Specifically, we assume $\Delta_k = \alpha_k - \alpha_{k-1} \sim \mathcal{N}(\delta_k, 2\sigma_\alpha^2)$, $k=2, 3, \dots, n$. We chose a normal distribution for the differences so that the joint distribution of the intercepts would be a multivariate normal distribution. Although such an approach does allow for the possibility of non-increasing DLT rates among doses of Agent B for a fixed dose of Agent A, we also tried using a prior gamma distribution for the differences (with the same mean and variance as the normal distributions) to maintain positive values for the differences. We found in simulations that the operating characteristics using the gamma distribution (results not shown) were similar to those presented in Table 2.

2.2 Alternative Models

As we have stated, we have chosen our model for p_{jk} because it is a direct generalization of that used with the CRM. Certainly, there are a myriad of other models that could be used and we focus upon three that have been proposed in recent publications. Two publications suggested that p_{jk} could be viewed as a joint probability of DLT that is a function of the marginal DLT rates of each agent and any possible interaction of the two agents. Let $0 < a_j < 1$ denote a rescaled value for dose j of Agent A, usually a marginal rate of DLT for dose j ,

and $0 < b_k < 1$ denote a rescaled value for dose k of Agent B, usually a marginal rate of DLT for dose j and $0 < b_k < 1$. Yuan and Yin (2009) used a copula credited to Gumbel to model p_{jk} as

$$p_{jk} = 1 - \left(1 - a_j^\alpha\right) \left(1 - b_k^\beta\right) \left[1 + a_j^\alpha b_k^\beta (e^\gamma - 1) / (e^\gamma + 1)\right],$$

in which α and β each have uniform priors over the interval $[0.2, 2]$ and γ has a prior Gamma distribution with shape and scale each set at 0.1, so that γ has a mean of 1 and variance of 10. Yin and Yuan (2009) used a copula credited to Clayton to model p_{jk} as

$$p_{jk} = 1 - \left[\left(1 - a_j^\alpha\right)^{-\gamma} + \left(1 - b_k^\beta\right)^{-\gamma} - 1 \right]^{-1/\gamma},$$

in which α and β each have a prior Gamma distribution with shape and scale each set at 2.0 and γ has a prior Gamma distribution with shape and scale each set at 0.1.

In an alternate approach, Braun and Wang (2010) assumed that each p_{jk} had a marginal Beta distribution with parameters a_{jk} and β_{jk} , both of which varied with doses of each agent through the parametric equations $\log(a_{jk}) = \gamma_0 + \gamma_1 a_j + \gamma_2 b_k$ and $\log(\beta_{jk}) = \varphi_0 - \varphi_1 a_j - \varphi_2 b_k$, in which each of the six parameters γ_0 , φ_0 , γ_1 , φ_1 , γ_2 , and φ_2 were independent Normal random variables each with their own mean and a common variance.

In Section 3, we will present a direct comparison of the proportional odds logistic regression model in the gCRM with the models of Yuan and Yin (2009) and Yin and Yuan (2009). The simulations will describe settings that have already been studied by Braun and Wang (2010), so that one could also be able to indirectly compare the performance of the gCRM with their method, although they used a smaller sample size and a different stopping rule. We again emphasize that we are not seeking to show that the gCRM is better than the other methods described. We simply wish to show that our simpler modeling approach leads to similar operating characteristics as the other modeling approaches.

2.3 Likelihood and Posterior Distribution

Let N_{jk} be the number of patients assigned to (j, k) , and Y_{jk} be the number of the N_{jk} patients with DLT. Based upon the accumulated data $\mathcal{Y} = \{Y_{jk}; j=1, 2, \dots, m; k=1, 2, \dots, n\}$ and $\mathcal{N} = \{N_{jk}; j=1, 2, \dots, m; k=1, 2, \dots, n\}$, the gCRM uses the likelihood

$$L_M(\beta, \alpha \mid \mathcal{Y}, \mathcal{N}) = \prod_{j=1}^m \prod_{k=1}^n p_{jk}(\beta, \alpha)^{Y_{jk}} [1 - p_{jk}(\beta, \alpha)]^{N_{jk} - Y_{jk}} \quad (1)$$

in which the subscript $M = \sum_{j=1}^m \sum_{k=1}^n N_{jk}$ refers to the total number of enrolled subjects. From this likelihood, we can compute the posterior distribution

$$f_M(\beta, \alpha \mid \mathcal{Y}, \mathcal{N}) = \frac{L_M(\beta, \alpha \mid \mathcal{Y}, \mathcal{N}) g(\beta) f(\alpha)}{\int \dots \int L_M(\beta, \alpha \mid \mathcal{Y}, \mathcal{N}) g(\beta) f(\alpha) d\beta d\alpha_1 \dots d\alpha_n},$$

in which $g(\beta)$ and $f(\alpha)$ are the respective Gamma and multivariate Normal distributions defined in Section 2.1. Given that this posterior distribution does not have a closed form expression, we can use Gibbs sampling to generate a sample of observations from the joint

posterior distribution of (β, α) . From this posterior sample, we can compute a posterior estimate for the probability of DLT for combination (j, k) as $\text{logit}(\tilde{p}_{jk}) = \tilde{\alpha}_k + \tilde{\beta}a_j$, in which $\tilde{\beta}$ and $\tilde{\alpha}_k$ are the respective posterior means of β and α_k . A natural estimate for the MTC would be the combination with $\tilde{p}_{jk}$ closest to the targeted DLT rate Γ .

2.4 Deriving Hyperparameter Values and Rescaled Doses

In the gCRM, there are $n + 3$ hyperparameter values $\mu_\beta, \sigma_\beta^2, \mu_\alpha, \sigma_\alpha^2$, and $\delta = \{\delta_2, \delta_3, \dots, \delta_n\}$ and m rescaled dose values for Agent A that need values assigned to them. Our approach is to have investigators supply the values p_{j1}^* and p_{1k}^* , which are the *a priori* estimates for p_{j1} and p_{1k} , the DLT rates for combinations that use the lowest dose of either or both of the agents, i.e. the first column and row of Figure 1(a). By doing so, we do not require *a priori* estimates for the DLT rates of every combination, as many of the combinations have not been studied and investigators would not have reliable *a priori* estimates for their DLT rates anyway. This leads to $m + n - 1$ values from which we seek to determine $m + n + 3$ values in our model. Due to this overparameterization, we choose to treat the prior means and variances for β and α_0 as four “tuning parameters” whose values will be determined via simulation, with the remaining values based directly upon the elicited information.

Specifically, if we assume from our logistic model that $\text{logit}(p_{j1}^*) = E(\alpha_1) + a_j E[\beta]$, where $E(\cdot)$ denotes prior expectation, we have the rescaled dose value

$a_j \approx [\text{logit}(p_{j1}^*) - \mu_\alpha] / \mu_\beta$. Since δ_k is approximately the average difference between the logits of the DLT rates of doses b_k and b_{k-1} , $k = 2, 3, \dots, n$ when each is combined with the same dose of Agent A, we set $\delta_k = \text{logit}(p_{1k}^*) - \text{logit}(p_{1,k-1}^*)$.

Although theoretically one might consider doing a grid search over numerous possible combinations of $\mu_\alpha, \mu_\beta, \sigma_\alpha^2$ and σ_β^2 , such an approach is computationally burdensome and likely not worth the amount of effort it would require. Instead, we first set $\mu_\beta = 1$ as has been adopted in most single-agent designs and also set $\mu_\alpha = -8$. We note that most applications of a logistic model with single-agent trials set the intercept to be 3. However, we have centered the distribution of α_0 at -8 instead of 3 because we want the *a priori* value of p_{11} to be relatively low, i.e. < 0.01 and a value of $\mu_\alpha = -8$ met this criterion, while a value of $\mu_\alpha = 3$ created a much larger value for the *a priori* value of p_{11} and then led to unrealistically large *a priori* DLT rates for all other combinations as well. For example, in our simulations, the lowest dose was assigned a value of $a_1 = 0.0325$, and the prior mean for $p_{11} < 0.01$ when $\mu_\alpha = -8$, but was instead $p_{11} = 0.95$ when $\mu_\alpha = 3$.

Identifying appropriate prior variances is less straightforward than identifying prior mean values. Our approach is to select different values for each variance within the range of [0.5, 2]. We then run a small, i.e. 100 replications, simulation study of the gCRM in three different settings: where the MTC is in close proximity to combination (m, n) , where the MTC is in close proximity to combination $(1, 1)$, and where the MTD lies at a combination of the median doses of both agents. We also assume the *a priori* DLT rates elicited from investigators match the actual DLT rates for those combinations. We then look at the percentage of the 100 simulations in which the MTC is identified correctly, or at least close to the actual MTC, and try to find a combination of values for σ_α^2 and σ_β^2 that lead to the best identification of the MTC. Through our own simulations, we found that $\sigma_\alpha^2 = 1$ and $\sigma_\beta^2 = 1$ were suitable choices. Parameter calibration is an inexact science and we remind readers that attempting to find the “best” hyperparameter values is a futile goal. One only need find values that allow the priors to be informative against the relative lack of data when only a

few patients have been enrolled, yet also allows the data to dominate the priors when more data has been gathered, especially near the end of the study.

We note that although the methods described in Section 2.1 are specific to the use of a logistic model, our idea is easily generalized to any one-parameter model used with the CRM. For example, a simple transformation of the power model gives us $\log[-\log F(\beta; d)] = \alpha + \log[-\log(d)]$, where $\alpha = \log(\beta)$, which is a standard logistic model with complementary log-log link and slope parameter fixed at 1.0. Application of our methods to this complementary log-log model is straightforward. Adopting the previous parameterization used for the logistic model, for each dose k of Agent B, we have $\log[-\log(p_{jk})] = \alpha_k + \beta \log[-\log(a_j)]$, in which the slope parameter is now allowed to vary from the fixed value of 1.0, if necessary. We once again define $\Delta_k = \alpha_k - \alpha_{k-1}$, $k = 2, 3, \dots, n$ and adopt the same form for the prior distributions of α_1 , β , and Δ . We use the elicited information to set

$\delta_k = \log[-\log(p_{1k}^*)] - \log[-\log(p_{1,k-1}^*)]$ and the rescaled dose values a_j are determined from the equation $\log[-\log(p_{j1}^*)] \approx \mu + \exp(\theta) \log[-\log(a_j)]$. Values for the other hyperparameters would be found through simulations as described in Section 2.4.

2.5 Trial Conduct

Two outstanding issues remain for our design. First, like most Phase I trial designs, we restrict the set of possible assignments to each patient in the trial so as to limit exposure of patients to toxic combinations but also promote sufficient exploration of the combinations and increase the likelihood of correctly identifying the MTC. The idea of sufficient exploration is of greater import when studying combinations of two agents instead of a single agent because every additional dose level in an agent leads to a significant increase in the total number of dose combinations. To this end, the first patient is assigned to combination (1, 1) and each future patient $i = 2, 3, \dots, N$ must be assigned to a combination that has doses of Agent A and Agent B that are within one dose (higher or lower) than the respective doses of Agent A and B assigned to patient $i - 1$. Note that this rule allows for the possibility of escalation of both doses at the same time, which may seem aggressive, but has been shown to not lead to increased exposure to unsafe combinations (Wang and Ivanova, 2005; Braun and Wang, 2010). We limit the amount of dose de-escalation to within one dose lower than the respective doses of Agent A and B because in two-agent trials, especially when few patients are entered into a study, a single DLT can be quite influential and result in dramatic dose de-escalation, which slows later dose escalation when the true MTC is among combinations including the highest dose of either agent. However, we did run simulations in which no limit on de-escalation was used (results not shown) and saw little change to the results presented in Section 3.

Second, the study should be terminated if there is sufficient evidence that all of the combinations are overly toxic. We have chosen to use the same stopping rule as that proposed in Yuan and Yin (2009) and Yin and Yuan (2009) based upon the posterior distribution for the probability of DLT for the combination of the lowest dose of both agents. If more than 95% of the distribution lies above the targeted DLT rate Γ , the trial is terminated, with no further enrollment and no combination selected as the MTC since all combinations are overly toxic.

Thus, the conduct of a Phase I trial for two agents using the gCRM is as follows:

1. Elicit *a priori* estimates for the probabilities of DLT for combinations including the lowest dose of either agent;

2. Determine appropriate values for the hyperparameters and rescaled dose values as described in Section 2.4;
3. Assign the first patient to combination (1, 1);
4. For patient $i = 2, 3, \dots, N$, compute the posterior means $p_{jk}^{\sim}$ for each combination (j, k) . If we let (j_{i-1}, k_{i-1}) denote the combination assigned to patient $i - 1$, we define the set of acceptable combinations as $S = \{(j, k) : \max(1, j_{i-1} - 1) \leq j \leq \min(j_{i-1} + 1, m), \max(1, k_{i-1} - 1) \leq k \leq \min(k_{i-1} + 1, n)\}$. The current estimate of the MTC, (j^*, k^*) , is the combination in S whose value of $p_{jk}^{\sim}$ is closest to Γ ;
5. Determine if the stopping rule described earlier has been reached. If so, terminate the trial; otherwise, assign patient i to the current MTC;
6. If data have been collected on all N patients, use step (4) to compute the final estimate of the MTC.

As noted by one referee, step (5) above uses a “greedy” algorithm that assigns each subject to the most recent estimate of the MTC. Although the greedy algorithm is used in a majority of adaptive phase I trial designs, we could also consider assigning each patient to a combination that satisfies some measure of informativeness, i.e. the combination with greatest posterior variance among those within an acceptable range of toxicity. By doing so, we hope to gain more information about the entire grid of combinations, rather than focusing on a narrow area of the grid, in hopes of better identifying the MTC at the end of the study.

3 Simulation Studies

3.1 Description of Simulation Settings

We examine the operating characteristics of our design in seven scenarios, scenarios A-G, with $m = 4$ doses of Agent A and $n = 4$ doses of Agent B. Scenarios A-F were previously examined in Braun and Wang (2010). We have a target DLT rate of $\Gamma = 0.20$. The true DLT rates of the 16 combinations in each scenario are displayed in Table 1. In scenario A, combinations of lower doses of both agents have DLT rates less than Γ while the combinations of higher doses of both agents have DLT rates greater than Γ . Scenarios B and D reflect settings when all the combinations are safe and too toxic, respectively. Scenario C is a situation when only a few dose combinations are tolerable while all others are overly toxic. Scenario E is an example when the DLT rates grow slowly with doses of one agent but rapidly with doses of the other agent. Scenario F represents a situation when there is a jump in the DLT rates of one agent, thus only the dose combinations associated with certain levels of that agent are safe. Scenario G reflects a situation when there are large deviations from the assumed proportional odds logistic model.

In Table 2, we have also included for each scenario a value quantifying the goodness-of-fit (GOF) of the assumed proportional odds model to the given DLT rates. The GOF value was found by simulating 10,000 observations for each dose combination from a Bernoulli distribution with the indicated DLT rate. We then fit the assumed proportional odds model to the 160,000 simulated observations to estimate the DLT rate for each dose combination. The GOF value is the sum of the sixteen squared differences of the fitted DLT rates and the actual DLT rates, divided by 1,000, so that a value of GOF=0 indicates perfect fit. These GOF values highlight the fact that Scenarios C and G deviate greatly from a proportional odds model. We also performed simulations in which there were a different number of doses for the two agents (results not shown) and found little change to the conclusions reached from these seven scenarios.

Since we have 16 dose combinations for each scenario, we chose to use $N = 50$ as the maximum sample size of each trial in our simulations, which is approximately three patients per dose combination on average, and we consider that to be the minimal possible sample size for exploration of all dose combinations. We used the first row and first column of true DLT rates for scenario A as the elicited values supplied by the investigators. Thus, scenario A represents the situation when the elicited values equal the truth, while the remaining scenarios represent the more realistic setting when the elicited values do not equal the truth. These elicited values generated rescaled dose values of $a_1 = 0.0325$, $a_2 = 0.0374$, $a_3 = 0.0405$, and $a_4 = 0.0427$ and hyperparameters $\delta_2 = 0.98$, $\delta_3 = 0.54$, and $\delta_4 = 0.39$. As we stated in Section 2.4, we use hyperparameter values $\mu_\alpha = -8$, $\mu_\beta = 1$, $\sigma_\alpha^2 = 1$, and $\sigma_\beta^2 = 1$. We ran 2,000 simulations for each of the seven scenarios; 2,000 parallel simulations were also performed by replacing the model used in the gCRM with each of the copula models of Yuan and Yin (2009) and Yin and Yuan (2009) described in Section 2.2.

In each simulation, we drew 5,000 samples from the posterior distributions with a burn-in of 1,000 samples. The samples were obtained via the “rjags” library in R, which also requires use of the free software JAGS (Just Another Gibbs Sampler) available on the internet. Among the 2,000 simulations in each scenario, we calculated the mean percentage that each dose combination was selected as the MTC at the end of the trial, the mean percentage of patients assigned to each dose combination, and the mean probability of DLT that each patient is exposed to based upon their assignments in the 2,000 simulations. All simulations were performed with the statistical language R and source code is available upon request from the authors.

3.2 Summary of Simulation Results

The operating characteristics of the three models in each scenario are summarized in Table 2. Columns 3-5 in Table 2 present the percentage of simulations in which each of the three models identifies the MTC at a combination with a DLT rate exactly equal to Γ , within 1-10 points of Γ , and beyond 10 points from Γ , respectively. Column 6 is the percentage of simulations that terminate early with no combination selected as the MTC. Likewise, columns 7-9 in Table 2 present the average proportion of patients assigned to combinations with a DLT rate exactly equal to Γ , within 1-10 points of Γ , and beyond 10 points from Γ , respectively. Column 10 is the average proportion of patients not assigned a combination due to early termination.

We reiterate that the goal of the simulations is to show that the gCRM and the models of Yuan and Yin (2009) and Yin and Yuan (2009) lead to generally similar operating characteristics rather than to show that one approach is better than the others. To that end, a general conclusion from the results in Table 2 is that all three models do lead to similar operating characteristics, although there are some subtle differences that we now highlight. The first difference, most apparent in scenarios C, E, and F, is that the adopted stopping rule tends to lead to a greater likelihood of early termination with the gCRM than with the other two models. Nonetheless, when the gCRM terminates a trial early, it tends to do so quite late in the trial, with only a handful of the maximum sample size not receiving an assignment.

The second difference is highlighted by the results of scenarios E and G, where it is apparent that the models of Yuan and Yin (2009) and Yin and Yuan (2009) are more likely to identify the MTC at combinations with DLT rates more than 10 points from the desired DLT rate. In fact, the model of Yin and Yuan (2009) appears to be more conservative than that of Yuan and Yin (2009). Specifically, in the 25% of simulations in scenario E and 13% of simulations in scenario G where the model of Yin and Yuan (2009) identified the MTC at combinations with a DLT rate 10 points beyond the target, the DLT rate of the chosen MTC

was more often below the target (14% in scenario E and 9% in scenario G) than above the target (11% in scenario E and 4% in scenario G). In contrast, the DLT rate of the MTC selected with the model of Yuan and Yin (2009) was more often above the target in scenario E, but was more often below the target in scenario G. These results, of course, bring up the typical quandry in Phase I trials: is it better to assign patients to less toxic, but possibly less effective, combinations (as with Yuan and Yin, 2009) or to terminate the trial (as with the gCRM) or assign patients to more toxic, but possibly more effective, combinations (as with Yin and Yuan, 2009)? The answer to that question will vary from trial to trial and is dependent upon the severity and potential reversibility of the DLTs and the disease severity of the patients being studied.

Although there is no apparent trend in Table 2 with regard to patient assignments, a definite trend is obvious in Figure 2, which presents the patterns of assignment in six of the seven scenarios for all three methods. We do not present a pattern for Scenario D, since all combinations were overly toxic and nearly all trials were terminated relatively early. The y-axis is the average DLT rate of the 2,000 dose assignments of each patient and the x-axis is the number of each patient (1 denotes the first patient, 2 denotes the second patient, etc). Note that the average DLT rate for each patient is only based upon trials in which that patient was enrolled, i.e. the mean dose for patient 50 would be based upon less than 2,000 simulations if some of the simulations had trials that terminated early. In all six scenarios, we see that although the assignments from all three models tend to converge to a combination with a DLT rate close to the target, escalation is fastest with the gCRM and slowest with the model of Yin and Yuan (2009). These plots also suggest that in terms of a sample size necessary for identifying the MTC, 30 patients tends to be sufficient, although the added accrual allows for a larger sample of patients who actually receive the recommended MTC.

Table 3 summarizes the sensitivity of the gCRM in scenarios C and G to the prior variances and the skeleton used to calibrate the dose values and some of the parameters. The first row of results for each scenario contains the values from Table 3 when $\sigma_\alpha^2 = \sigma_\beta^2 = 1$, the next four rows summarize the results when the values σ_α^2 and σ_β^2 are increased or decreased, and the last row summarizes the results when the skeleton values come from the first row and column of DLT rates for scenario E listed in Table 1. For scenario C, we see that all combinations of values for the prior variances lead to similar results when using a skeleton based upon the true DLT rates in scenario A. However, we see that the ability of the gCRM to identify the MTD is hindered when the skeleton instead based upon the true DLT rates in scenario E. This is because the gCRM is more likely to select combinations with the lowest dose of Agent A, which the skeleton believes are less toxic than they truly are. This indicates that the skeleton leads to an overly informative prior and a larger value of σ_α^2 and/or σ_β^2 is warranted. The inter-relationship of prior variance and skeleton has been well-documented in single-agent adaptive designs Lee and Cheung (2011), and our results emphasize the same is true for two-agent adaptive designs.

Scenario G is a setting that proved challenging for all three models and we see that although using prior variance values $\sigma_\alpha^2 = \sigma_\beta^2 = 1$ leads to relatively good operating characteristics, the gCRM identifies the MTC slightly better if at least σ_α^2 is doubled. We also see that the operating characteristics in scenario G differ less between the two skeletons than they did with scenario C, mostly likely because the first column and row of DLT for scenario E (the skeleton) are much closer to the first column and row of DLT rates for scenario G than for scenario C.

4 Discussion

The gCRM is a logical extension of the original CRM for single-agent Phase I trials to the study of two agents by replacing the logistic model used in the traditional CRM with a proportional odds logistic model. However, the gCRM only assumes proportional odds of DLT for adjacent doses of one of the agents (Agent A) when combined with the same dose of the other agent (Agent B). The odds of DLT among adjacent doses of Agent B, when combined with the same dose of Agent A, are not necessarily required to be proportional nor monotonic, although monotonicity is expected and can be promoted through appropriate selection of the prior variance parameter. In essence, we have reduced the problem of estimating $m \times n$ probabilities of DLT to the estimation of $n + 1$ model parameters. This fact motivates our suggestion that the label “Agent B” be given to the agent with fewer dose levels being studied, as this will minimize the number of parameters that need to be estimated with the traditionally small sample sizes used in Phase I trials.

Our approach of adding a series of correlated dose-specific parameters to the model used in the CRM is in the same spirit of O’Quigley and Paoletti (2003), who sought to fit a CRM to groups of patients whose DLT rates were expected to follow a specific order. Thus, an individual model was applied to each group, with one parameter in each of those models having a joint distribution that enforced the ordering of the groups, yet allowed for sharing of data among all groups during estimation.

The results presented in Section 3 demonstrate that no one design for identifying the MTC will be best in all scenarios. In fact, the simulation results suggest that the choice of model is likely less important than careful selection of the skeleton values and hyperparameter values in the prior distributions for the parameters in that model. Careful simulation across a multitude of expected true DLT rates and different choices of prior are crucial before any design can be adopted in practice. Furthermore, it remains undetermined what an appropriate sample size might be for two agent designs, which is obviously a function of the number of combinations under study, but is more subtly a function of how quickly a design can converge to a neighborhood of the MTC. The plots we have presented in Figure 2 are an excellent tool for assessing this convergence and we recommend their use for determining an appropriate sample size in actual trials.

Acknowledgments

This research was supported by National Institute of Health grant 5R01CA148713. We thank the editor and two referees for their constructive and thoughtful reviews that helped the manuscript significantly.

References

- Braun T, Wang S. A hierarchical Bayesian design for phase I trials of novel combinations of cancer therapeutic agents. *Biometrics*. 2010; 66:805–812. [PubMed: 19995354]
- Conaway M, Dunbar S, Peddada S. Designs for single or multiple-agent phase I trials. *Biometrics*. 2004; 60:661–669. [PubMed: 15339288]
- Faries D. Practical modifications of the continual reassessment method for phase I cancer clinical trials. *Journal of Biopharmaceutical Statistics*. 1994; 4:147–164. [PubMed: 7951271]
- Goodman S, Zahurak M, Piantadosi S. Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine*. 1995; 14:1149–1161. [PubMed: 7667557]
- Lee SM, Cheung YK. Calibration of prior variance in the Bayesian continual reassessment method. *Statistics in Medicine*. 2011; 30:2081–2089. [PubMed: 21413054]
- O’Quigley J, Paoletti X. Continual reassessment method for ordered groups. *Biometrics*. 2003; 59:430–440. [PubMed: 12926728]

- O'Quigley J, Shen L. Continual reassessment method: A likelihood approach. *Biometrics*. 1996; 52:673–684. [PubMed: 8672707]
- 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]
- Rowinsky E, Kaufmann S, Baker S, Grochow L, Chen T, Peereboom D, Bowling M, Sartorius S, Ettinger D, Forastiere A, Donehower R. Sequences of topotecan and cisplatin: Phase I, pharmacologic, and in vitro studies to examine sequence dependence. *Journal of Clinical Oncology*. 1996; 14:3074–3084. [PubMed: 8955652]
- Thall P, Millikan RE, Mueller P, Lee S-J. Dose-finding with two agents in phase I oncology trials. *Biometrics*. 2003; 59:487–496. [PubMed: 14601749]
- Wages N, Conaway M, O'Quigley J. Continual reassessment method for partial ordering. *Biometrics*. 2011a; 67:1555–1563. [PubMed: 21361888]
- Wages N, Conaway M, O'Quigley J. Dose-finding design for multi-drug combinations. *Clinical Trials*. 2011b; 8:380–389. [PubMed: 21652689]
- Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics*. 2005; 61:217–222. [PubMed: 15737096]
- Yin G, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society*. 2009; 58:211–224.
- Yuan Y, Yin G. Sequential continual reassessment method for two-dimensional dose finding. *Statistics in Medicine*. 2008; 27:5664–5678. [PubMed: 18618901]
- Yuan Y, Yin G. A latent contingency table approach to dose finding for combinations of two agents. *Biometrics*. 2009; 65:866–875. [PubMed: 18759848]

		Dose of Agent A			
		j=1	j=2	j=3	j=4
Dose of Agent B	k=1	$\text{logit}(p_{j1}) = \alpha_1 + a_j\beta$			
	k=2	$\text{logit}(p_{j2}) = \alpha_2 + a_j\beta$			
	k=3	$\text{logit}(p_{j3}) = \alpha_3 + a_j\beta$			
	k=4	$\text{logit}(p_{j4}) = \alpha_4 + a_j\beta$			

Figure 1.
Schematic representation of gCRM model for phase I trial of two agents, with four doses of each agent.

Figure 2.

Average dose combination (in terms of DLT rate) assigned to each patient in scenarios A-C and E-G. YY09a and YY09b refer to the methods of Yuan and Yin (2009) and Yin and Yuan (2009), respectively. The horizontal line represents the target DLT rate $\Gamma = 0.20$.

True DLT rates ($\times 100$) for the seven scenarios examined; GOF = metric quantifying goodness-of-fit of proportional odds model to DLT rates, with higher values indicating poorer fit.

Table 1

Agent A						Agent A					
Scenario	Agent B	1	2	3	4	Scenario	Agent B	1	2	3	4
A	1	4	8	12	16	E	1	8	18	28	29
	2	10	14	18	22		2	9	19	29	30
	3	16	20	24	28		3	10	20	30	31
	4	22	26	30	34		4	11	21	31	41
GOF=3.1						GOF=0.8					
Agent A						Agent A					
Scenario	Agent B	1	2	3	4	Scenario	Agent B	1	2	3	4
B	1	2	4	6	8	F	1	12	13	14	15
	2	5	7	9	11		2	16	18	20	22
	3	8	10	12	14		3	44	45	46	47
	4	11	13	15	17		4	50	52	54	55
GOF=1.0						GOF=1.0					
Agent A						Agent A					
Scenario	Agent B	1	2	3	4	Scenario	Agent B	1	2	3	4
C	1	10	20	30	40	G	1	1	2	3	4
	2	25	35	45	55		2	4	10	15	20
	3	40	50	60	70		3	6	15	30	45
	4	55	65	75	85		4	10	30	50	80
GOF=10.0						GOF=20.7					
Agent A						Agent A					
Scenario	Agent B	1	2	3	4	Scenario	Agent B	1	2	3	4
D	1	44	48	52	56						
	2	50	54	58	62						
	3	56	60	64	68						
	4	62	66	70	74						

Scenario	Agent A				Agent B				Agent A			
	1	2	3	4	1	2	3	4	1	2	3	4
GOF=1.6												

Table 2

Operating characteristics of the gCRM and the models proposed by Yuan and Yin (2009) (denoted YY09a) and Yin and Yuan (2009) (denoted YY09b). All values are multiplied by 100. All results are in reference to a desired DLT rate $\Gamma = 0.20$.

Scenario	Model	At Γ	Prob of Selection as MTC			Mean Prop of Patients Assigned		
			1-10		None	1-10		None
			Pts of Γ	Pts of Γ		At Γ	Pts of Γ	
A	gCRM	10	82	3	5	6	72	17
	YY09a	13	82	5	0	13	72	15
	YY09b	11	81	6	2	10	70	20
B	gCRM	0	94	3	3	0	87	13
	YY09a	0	99	1	0	0	86	14
	YY09b	0	96	4	0	0	71	29
C	gCRM	45	39	5	11	30	41	18
	YY09a	41	50	5	4	27	54	16
	YY09b	42	47	5	6	29	55	11
D	gCRM	0	0	4	96	0	0	22
	YY09a	0	0	1	99	0	0	20
	YY09b	0	0	1	99	0	0	16
E	gCRM	9	70	14	7	5	56	32
	YY09a	6	65	27	2	9	55	34
	YY09b	7	67	25	1	6	54	38
F	gCRM	13	70	6	11	10	64	16
	YY09a	14	76	6	4	7	75	14
	YY09b	12	74	7	7	7	77	9
G	gCRM	25	68	5	2	18	57	24
	YY09a	12	76	12	0	3	71	26
	YY09b	15	72	13	0	7	61	32

Table 3

Sensitivity of the operating characteristics of the gCRM in scenarios C and G to different prior distributions or skeleton. All values are multiplied by 100. All results are in reference to a desired DLT rate $\Gamma = 0.20$.

Scenario C:		Prob of Selection as MTC				Mean Prop of Patients Assigned			
σ_a^2	σ_β^2	Skeleton	At Γ	1-10		> 10		1-10	
				Pts of Γ	None	Pts of Γ	None	Pts of Γ	None
1.0	1.0	A	45	39	5	11	30	41	18
0.5	0.5	A	49	44	4	3	32	50	15
2.0	0.5	A	49	44	5	2	33	49	15
0.5	2.0	A	43	42	3	12	28	45	16
2.0	2.0	A	44	41	5	10	28	41	19
1.0	1.0	E	32	51	5	12	18	51	17

Scenario G:		Prob of Selection as MTC				Mean Prop of Patients Assigned			
σ_a^2	σ_β^2	Skeleton	At Γ	1-10		> 10		1-10	
				Pts of Γ	None	Pts of Γ	None	Pts of Γ	None
1.0	1.0	A	25	68	5	2	18	57	24
0.5	0.5	A	19	76	4	1	14	64	22
2.0	0.5	A	28	62	8	2	21	51	26
0.5	2.0	A	22	73	5	0	15	62	22
2.0	2.0	A	31	58	8	3	24	49	25
1.0	1.0	E	21	71	8	0	15	60	25