Published in final edited form as:

Stat Methods Med Res. 2016 April ; 25(2): 659–673. doi:10.1177/0962280212464541.

Phase I/II trial design when response is unobserved in subjects with dose-limiting toxicity

Thomas M. Braun, Shan Kang, and Jeremy M. G. Taylor

Department of Biostatistics, University of Michigan School of Public Health, Ann Arbor, MI, USA

Abstract

We propose a Phase I/II trial design in which subjects with dose-limiting toxicity are not followed for response, leading to three possible outcomes for each subject: dose-limiting toxicity, absence of therapeutic response without dose-limiting toxicity, and presence of therapeutic response without dose-limiting toxicity. We define the latter outcome as a ‘success,’ and the goal of the trial is to identify the dose with the largest probability of success. This dose is commonly referred to as the most successful dose. We propose a design that accumulates information on subjects with regard to both dose-limiting toxicity and response conditional on no dose-limiting toxicity. Bayesian methods are used to update the estimates of dose-limiting toxicity and response probabilities when each subject is enrolled, and we use these methods to determine the dose level assigned to each subject. Due to the need to explore doses more fully, each subject is not necessarily assigned the current estimate of the most successful dose; our algorithm may instead assign a dose that is in a neighborhood of the current most successful dose. We examine the ability of our design to correctly identify the most successful dose in a variety of settings via simulation and compare the performance of our design to that of competing approaches.

Keywords

Dose-finding studies; early-phase clinical trials; most successful dose; adaptive design

1 Background

Traditionally, the safety and efficacy of a new oncologic agent is studied sequentially in two separate phases. Safety is first examined in a Phase I clinical trial, in which one seeks to find the maximum tolerated dose (MTD), the dose that has the probability of dose-limiting toxicities (DLT) closest to a pre-defined target probability, usually in the range of 0.20–0.35. Efficacy, or the probability of clinical response, of the MTD is then typically examined in a separate Phase II trial. Due to statistical and resource insufficiency of this two-phase approach, several designs have been proposed for studies that examine toxicity and efficacy simultaneously in the same group of subjects.^{1–5} We categorize these previous methods into two design types, Type A and Type B. Type A designs, like that of Braun² and Yin et al.,⁵

assume that toxicity and response are each binary outcomes whose marginal probability of occurring is a parametric function of dose. A joint distribution is then developed to incorporate any possible correlation between the two outcomes, most often through the conditional distribution of one outcome given the other occurs or through a copula. Type B designs, like that of Thall and Russell,¹ create an ordinal variable whose possible values are based upon the possible combinations of toxicity and response outcomes that will be observed. It is then the marginal distribution of this ordinal variable that is modeled as a function of dose. Implicit in these designs is modeling the probability of success, defined as joint probability of response and no DLT. The dose with the largest value of this probability has been termed the most successful dose (MSD). We note that although the the Type A and Type B designs differ in their modeling approach, one can envision specific parameterizations of each model that would lead to the same probability of success.

The concept of an MSD was first proposed by O'Quigley et al.⁶ in the design of trials examining the safety and efficacy of anti-retroviral agents in children with HIV. In their proposed design, any dose d is associated with a probability of DLT, $P(d, \gamma)$, and a probability of response given no DLT, $Q(d, \theta)$, in which both $P(\cdot)$ and $Q(\cdot)$ are monotonic in dose, and γ and θ are unknown parameters that are estimated using methods adopted directly from the continual reassessment method (CRM).⁷ Most notably, the likelihood used for the estimation of γ and θ assumes that all four combinations of DLT (yes/no) and response (yes/no) will be observed in the study. The product $R(d, \gamma, \theta) = Q(d, \theta)[1 - P(d, \gamma)]$ leads to the probability of success for any dose, and each subject is assigned to a dose with the dose with the largest value of $R(\cdot)$ based upon the outcomes of previously enrolled subjects. The application of these concepts to oncology were presented by Zohar and O'Quigley,⁸ and discussion of optimal designs for finding the MSD were presented in Zohar and O'Quigley.⁹ Earlier theoretic work related to these concepts was presented by Li et al.¹⁰

Ivanova³ presented a design in which dose assignments were based on an up-and-down algorithm rather than a parametric model. Specifically, let $D_{(i)}$ denote the dose assigned to subject i . The dose assigned to subject $(i + 1)$ will be dose $D_{(i)} - 1$ if subject i had a DLT (regardless of response), dose $D_{(i)} + 1$ if subject i had no DLT and no response, and dose $D_{(i)}$ if subject i had no DLT and response. The marginal probability of DLT was also modeled as a monotonic one-parameter function of dose as a way to monitor the toxicity of each dose. If the algorithm would assign to subject i a dose with an estimated DLT rate above the desired threshold, then the dose assigned to subject $(i + 1)$ would instead be largest dose that had an estimated DLT rate below the threshold. One unique feature of this design is that it does not require that response be measured on subjects who experience DLT. Specifically, in most Phase I/II studies, investigators follow subjects for DLT for a shorter period of time, usually a few weeks, than they follow subjects for response, which can take months to occur. As a result, if a subject experiences a DLT, their follow-up may end and they cannot be observed for response; see Lissitchkov et al.¹¹ for an example of a trial in which this was the case. It is this setting that motivates our methods.

Given that response is not measured on subjects who experience a DLT, Type A Phase I/II designs are not suitable in our setting, as the marginal distribution for response cannot be determined. One might consider using a Type B Phase I/II design in which the ordinal

variable could take the value 0, 1, or 2 depending on whether the subject experienced DLT, experienced no response without DLT, or experienced response without DLT, respectively. However, given that response data is naturally collected conditional upon no DLT, we have instead chosen to model the marginal probability of DLT and the conditional probability of response given no DLT. Like O'Quigley et al.,⁶ we assume that the probability of DLT is monotonic in dose. However, given that it is questionable that the probability of response is also monotonic in dose for subjects not experiencing DLT, we make no such assumption in our methods. Furthermore, given the relatively small amount of response data that will be observed, we allow for some correlation among the probabilities of response for adjacent doses so that we can 'smooth' the data in order to estimate the conditional response probabilities for each dose. Our methods are described in Section 2 and the performance of our design is compared to the design of Ivanova via simulation in Section 3. Concluding remarks are in Section 4.

2 Proposed methods

2.1 Models and prior distributions

Consider a study that is designed to examine K dose levels in a sample of n subjects, with each subject $i = 1, 2, \dots, n$ receiving one of the K doses. Subject i will be observed for two outcomes: Y_i , which is a binary indicator of DLT, and Z_i , which is a binary indicator of therapeutic response. It is assumed that each subject will be followed for τ time units after starting treatment and DLT can occur any time during that period. In contrast, each subject will be assessed for response only after being observed for τ time units. Moreover, response can only be observed in subjects that did not experience DLT by time τ , so that the bivariate combination of DLT and response ($Y_i = 1$ and $Z_i = 1$) cannot be observed.

For dose $k = 1, 2, \dots, K$, we let p_k be the probability of DLT and q_k be the conditional probability of response given no DLT; we assume that p_k is monotonically increasing in dose, although no such assumption is made for q_k . We let $r_k = (1 - p_k)q_k$ be the probability of success, defined as the probability of response and no DLT. Our goal is to find the dose with the largest value of r_k at the end of the study among all doses in which $p_k \leq \pi^{\max}$, with the value $\pi^{\max}$ denoting an unacceptably high probability of DLT. Before the study begins, the investigator supplies two values for each dose k : (1), p_k^* , the a priori probability of DLT, and (2) q_k^* , the a priori probability of response in subjects without DLT. In traditional Phase I trial designs, these values are known as the 'skeleton.' The skeleton also leads to values $r_1^*, r_2^*, \dots, r_K^*$ for the probability of success for each dose, in which $r_k^* = (1 - p_k^*)q_k^*$.

We follow the convention of the CRM and model the probability of DLT as $\text{logit}(p_k) = 3 + \beta D_k$, where $\text{logit}(x) = \log(x) - \log(1 - x)$ and β is an unknown parameter assumed to have a prior $\mathcal{N}(1, \sigma^2)$ distribution. We note that although our regression model allows β to be negative and does not strictly enforce increasing monotonicity among our estimates of p_k , the prior mean of 1.0 for β does encourage monotonicity as a majority of the mass of the prior distribution will lie above zero. Each dose value D_k is not an actual clinical dose, but is a value assigned to dose k that promotes the fit of our model. Following the convention of the authors of the CRM, we take our model for the probability of DLT and substitute p_k by

its corresponding skeleton value p_k^* and β by its prior mean of 1 and then solve for D_k , resulting in $D_k = \text{logit}(p_k^*) - 3$.

We choose to use a value of $\sigma^2 = 0.5$ for the prior variance of β , which we determined through simulations over a range of possible values of σ^2 . Specifically, we simulated hypothetical trials in a variety of settings (Table 1) using values for σ^2 in the set $\{0.25, 0.50, 1.00, 1.50, 2.00\}$. We found that $\sigma^2 = 0.50$ led to the best operating characteristics, i.e. correctly identified the MSD and assigned many subjects to the MSD, on average among all the settings. Like any application of the CRM, a suitable value for σ^2 will vary in practice from study to study and needs to be determined heuristically, although a more systematic approach to determining a value for σ^2 was proposed by Lee and Cheung.¹²

However, the logistic model adopted for modeling the association of dose with probability of DLT is too restrictive for estimating q_k , as it is unclear that the probability of response, conditional on no DLT, must necessarily increase with dose. Instead, we assume each q_k has a marginal prior Beta distribution with fixed parameters a_k and $A - a_k$, where A is a value that represents the ‘prior effective sample size’ on which each of the prior distributions is based. An appropriate value for A reflects how informative the prior marginal Beta distributions should be relative to the data collected in the study, with larger values of A leading to a smaller prior variance. Via simulation, we have found that value of A in the interval $[1, 5]$ is sufficient in terms of correctly identifying the MSD. Once a value for A is determined, the value for a_k results by setting a_k/A , the prior mean of q_k , to its corresponding skeleton value q_k^* , so that $a_k = Aq_k^*$.

The joint prior distribution for $q_1, q_2, \dots, q_K$ is then fully specified with one additional parameter $\rho \in [0, 1]$ by defining $\text{Corr}(q_k, q_\ell) = \rho^{|k-\ell|}$, better known as an autoregressive correlation structure, as we expect adjacent doses to provide some information regarding each of their probabilities of response. Although this correlation structure may not completely reflect the underlying correlation of the conditional response rates, we selected this correlation structure because of its simple parametric structure and its having more biologic plausibility than other parametric correlation structures like an exchangeable correlation matrix. The parameter ρ is fixed during the trial because there will be insufficient data collected during the trial that would allow for estimation of ρ . However, the actual value of ρ used appears to be relatively irrelevant to the performance of our design, and this fact will be discussed further in Section 3.

For the sake of clarity, we re-iterate that we collect data to estimate the parameters β and $\mathbf{q} = \{q_1, q_2, \dots, q_K\}$ using Bayesian methods. The means of the prior distributions for β and $\mathbf{q}$ are chosen to match the values of the skeleton values for the DLT and conditional response rates. The prior distribution of β contains one parameter σ^2 that needs a value assigned to it, and the prior distribution of $\mathbf{q}$ contains two parameters, A and ρ , that need values assigned to them. In general application, appropriate values for these three parameters could be determined by selecting values for each and then applying our design to trials simulated from a specific setting to see how often the MSD is correctly determined and how many subjects are assigned to the MSD. New values would be assigned to one or all of σ^2 , A , and ρ and the simulations repeated until a specific combination of the three parameters produces

desirable operating characteristics. This process would be repeated over several different settings (like those in Table 1) until a combination of parameter values is found that works well in all settings. Given that this simulation process is quite time-consuming, we have stated earlier a range of values for each parameter that we feel should work well in practice for most trials.

2.2 Dose assignments

2.2.1 Posterior computations—Like most Phase I designs, we assign the first subject to the lowest dose to reduce the probability of DLTs occurring early in the trial. We further make the restriction that future subjects cannot be assigned to a dose until all lower doses have been assigned to at least one previous subject. Subject to this one constraint, the dose assignment for every future subject $j = 2, 3, \dots, n$, will be based upon the dose assignments and outcomes of subjects $1, 2, \dots, (j-1)$. For dose k , we let W_{1k} denote the number of subjects with DLT, W_{2k} denote the number of subjects without DLT who respond, and W_{3k} denote the number of subjects without DLT who fail to respond. This data from the first $(j-1)$ subjects leads to the likelihood

$$L(\beta, \mathbf{q}) = \prod_{k=1}^K \frac{[\exp(3+\beta D_k)]^{W_{1k}} q_k^{W_{2k}} (1-q_k)^{W_{3k}}}{[1+\exp(3+\beta D_k)]^{j-1}} \quad (1)$$

in which $\sum_{k=1}^K \sum_{\ell=1}^3 W_{\ell k} = (j-1)$. The posterior distribution for β and $\mathbf{q}$ is equal to

$$\frac{L(\beta, \mathbf{q}) f(\beta) g(\mathbf{q})}{\int \dots \int L(\beta, \mathbf{q}) f(\beta) g(\mathbf{q}) d\beta dq_1 \dots dq_K},$$

in which $f(\beta)$ and $g(\mathbf{q})$ are the respective prior distributions for β and $\mathbf{q}$ described earlier. We then compute the posterior mean of β and plug this value into our model for probability of DLT, leading to posterior estimates $p_1, p_2, \dots, p_K$. Combining these values with the posterior means $q_1, q_2, \dots, q_K$, we compute $r_1, r_2, \dots$, and r_K , where $r_k = q_k(1 - p_k)$.

We could draw samples from the posterior distribution of β and $\mathbf{q}$ using Markov Chain Monte Carlo (MCMC) methods. Specifically, although a closed-form expression for the joint prior distribution of $\mathbf{q}$ is not available, the assumed autoregressive correlation structure allows us to express each q_k , $k = 2$, as a linear function of q_{k-1} . As a result, we could first sample q_1 and then sample each remaining q_k conditional upon the value of q_{k-1} sampled. Nonetheless, this process has a large computational burden (that increases with k), which is trivial in a single application but increases dramatically for a simulation study. Therefore, we use the following approach, which although statistically less efficient than MCMC, is computationally faster than MCMC and produces valid posterior estimates.

For the DLT rates, we directly draw 100,000 samples from the prior distribution of the model parameter β , which in turn, provides 100,000 prior samples for each of $p_1, p_2, \dots, p_K$. We then sample the quantiles $Q_1, Q_2, \dots, Q_K$ from a multivariate normal distribution in

which each Q_k had mean zero and variance 1 and $Q_1, Q_2, \dots, Q_K$ have an autoregressive correlation structure with $\text{Corr}(Q_k, Q_\ell) = \rho^{|k-\ell|}$. If we let $W_k = \Phi(Q_k)$, where $\Phi(x)$ is the cumulative distribution function of a standard normal distribution, then $q_k = \mathcal{B}(W_k; a_k, A - a_k)$, where $\mathcal{B}(w; a_k, A - a_k)$ is the inverse CDF of a Beta distribution with parameters a_k and $A - a_k$.

This leads to a vector of correlated Beta random variables, although their correlation is not equal to ρ . In fact, the magnitude of the correlation is restricted by the values of the parameters in the marginal Beta distributions. Our approach is in the spirit of a copula, where the marginal distributions of each variate are known, but no form exists for their joint distribution. We repeated this process 100,000 times to produce a representative sample from the prior joint distribution of $q_1, q_2, \dots, q_K$. After data from each subject is collected, we then take the resulting likelihood and average it over the 100,000 draws to produce posterior means for $p_1, p_2, \dots, p_K$ and $q_1, q_2, \dots, q_K$.

2.2.2 Greedy dose assignment algorithm—If we were to adopt a ‘greedy’ dose-assignment approach, then we would assign the next subject to the dose with the largest value of r_k^* among doses with $p_k^* > \pi^{\max}$. However, any adaptive design for early-phase clinical trials has two competing goals. The first goal is local and is focused on assigning each subject to the dose that is deemed the MSD at the time the subject is enrolled. This goal is addressed through the ‘greedy’ approach just described. The second goal is global and is focused on finding the correct MSD at the end of the study. The dose assignment for each subject is then based on reaching this goal, which requires more exploration of all dose levels and may not coincide with the ‘greedy’ dose assignment. As we will show in Section 3, a ‘greedy’ approach leads to a poor ability of finding the MSD at the end of the study.

For example, suppose we have four dose levels and, for lack of prior information on the pattern of response with dose in subjects without DLT, an investigator gives us a ‘flat’ skeleton $(q_1^*, q_2^*, q_3^*, q_4^*) = (0.2, 0.2, 0.2, 0.2)$ that are constant among the doses. Suppose we set $A = 1$ and $\rho = 0$ so that all $q_k, k = 1, 2, 3, 4$ are mutually independent with the same $\text{Beta}(0.2, 0.8)$ prior distribution. If the first subject (who is assigned to the lowest dose) does not have a DLT and also has a response, the posterior distribution for q_1 will be $\text{Beta}(1.2, 0.8)$, while the posterior distribution for each of q_2, q_3 , and q_4 will be $\text{Beta}(0.2, 0.8)$ because $\rho = 0$ so that the first subject supplies no information for other than the lowest dose. This leads to a posterior value $q_1^* = 0.6$ that is larger than $q_2^* = q_3^* = q_4^* = 0.20$. Because of our assumption that the probability of DLT tends to increase with dose, we have $(1 - p_1^*) > (1 - p_2^*) > (1 - p_3^*) > (1 - p_4^*)$, and we naturally have the lowest dose with the largest value of r_k^* , which is then deemed the MSD and assigned to the next subject.

If that second subject also had no DLT, but failed to respond, q_1 would have a posterior $\text{Beta}(1.2, 1.8)$ distribution with mean $q_1^* = 0.4$, which is still larger than the posterior means of the other doses and the third subject would also be assigned to the lowest dose. In fact, all future dose assignments would remain at the lowest dose until q_1^* is sufficiently smaller than q_2^* to make $r_1^* < r_2^*$. Although values of $\rho > 0$ would allow for q_2^*, q_3^* , and q_4^* to change from their prior means, the amount of change would not be drastic and lead to the same pattern of dose assignments just presented. As a result, we need to assign doses in a way that allows

for greater exploration among the possible doses and increases the accuracy of the MSD determined at the end of the study.

2.2.3 Alternate dose assignment methods—The most common alternative to a ‘greedy’ dose assignment is to fully randomize the assignment of each subject to all doses based upon their estimated proximities to the true MSD; see Thall et al.¹³ for an example of this approach. As applied to our setting, if r_k is the current estimate of success for dose k , that dose would be assigned to the next subject with probability $\hat{r}_k^\lambda / \sum_{k=1}^K \hat{r}_k^\lambda$, where $\lambda > 0$ is a fixed value. Values of $\lambda \in \{1, 1.5, 2\}$ are commonly used values with adaptive randomization, although there is no methodology for selecting a ‘best’ value to use for a specific application. This dose assignment approach will be examined in the simulations presented in Section 3. As a compromise between the ‘greedy’ approach, which can be viewed as randomization with probability equal to one for the current MSD and equal to zero for all other doses, and randomization among all doses, we consider an approach in which each subject’s dose assignment is either the current MSD or one of the other doses, which is determined as follows.

At any point in the study, we define the standardized distance between two doses d_{k_1} and d_{k_2} as

$$\mathcal{D}(d_{k_1}, d_{k_2}) = \frac{(r_{k_1} - r_{k_2})}{\sqrt{r_{k_1}(1-r_{k_1})/n_{k_1} + r_{k_2}(1-r_{k_2})/n_{k_2}}} \quad (2)$$

where r_{k_1} and r_{k_2} are the respective current estimated success rates of d_{k_1} and d_{k_2} and n_{k_1} and n_{k_2} are the respective number of patients that have been assigned to d_{k_1} and d_{k_2} . Thus, any dose that has not yet been assigned to any subjects will have a distance of zero from every other dose. Because this distance function will be used twice to determine which dose to assign to each subject, we refer to our method as a ‘two-stage’ dose-assignment approach.

Suppose we have collected the DLT and conditional response outcomes for the first $(j-1)$ subjects and have computed the posterior means p_k , q_k , and r_k for each dose. Let $\Gamma_0 = \{d_k : p_k \sim \pi^{\max}\}$ be the set of doses with acceptable posterior DLT estimates, r^* be the largest value of r_k among the doses in Γ_0 , and d^* be the dose with the posterior probability of success equal to r^* e.g. the current estimate of the MSD, which has been assigned to n^* patients. We next define $\Gamma_1 = \{d_k : d_k \in \Gamma_0 \text{ \& } d_k \neq d^*\}$, which is the subset of doses in Γ_0 other than d^* . For all doses $d_k \in \Gamma_1$, we compute $\mathcal{D}(d^*, d_k)$ and let d^{**} denote the dose in Γ_1 with the smallest value of $\mathcal{D}(d^*, d_k)$, which has been assigned to n^{**} subjects and whose estimated probability of success is denoted r^{**} . We see that a dose will be selected as d^{**} if the posterior distribution for its probability of success is either (a) close to the current MSD on average, or (b) has a relatively large variance, which is approximated by a standard binomial variance equation and would occur by its having been assigned to fewer subjects than the other members of Γ_1 . Note that this interpretation implicitly assumes that r^* is independent of every other element of r_k , so that the denominator of equation (2) is the standard deviation of the numerator. We also remind readers that our restriction on dose

escalation limits d^{**} to be no more than one dose larger than the largest dose assigned to subjects 1, 2, ..., $(j-1)$.

We now need to decide if subject j (the next subject) should be assigned to dose d^* , which would be the same as the ‘greedy’ dose assignment, or to dose d^{**} , which would increase the exploration of doses around the current MSD. Our metric for this decision is the value of $\mathcal{D}(d^*, d^{**})$ that would result *after* subject j has been assigned and followed for DLT and response. Specifically, when a dose assignment is later needed for subject $(j+1)$, we should have assigned subject j to the dose that increased our confidence that doses d^* and d^{**} truly are different from each other, i.e. the dose that led to a *larger* value of $\mathcal{D}(d^*, d^{**})$. If we assume that the values of r^* and r^{**} do not change with the addition of data from subject j , then we can focus on the denominator of equation (2), which is equal to

$$\frac{\tilde{r}^*(1-\tilde{r}^*)}{n^*+1} + \frac{\tilde{r}^{**}(1-\tilde{r}^{**})}{n^{**}} \quad (3)$$

if subject j is assigned to dose d^* , and

$$\frac{\tilde{r}^*(1-\tilde{r}^*)}{n^*} + \frac{\tilde{r}^{**}(1-\tilde{r}^{**})}{n^{**}+1} \quad (4)$$

if subject j is assigned to dose d^{**} . The actual assignment for the next subject would be the dose (d^* or d^{**}) that leads to the smaller of equations (3) and (4). Put another way, the denominator of equation (2) is an approximation to the variance of $(r_{k1} - r_{k2})$ and making the variance of this difference small would provide more confidence that we can correctly identify which success rate is larger.

Recall that because we limit each the dose assignment of each subject to be no more than one dose above those already assigned, it is not possible for both n^* and n^{**} be equal to zero. However, it is possible for $n^*=0$, in which case d^* would be assigned to patient j because instead assigning d^{**} would result in an infinite value for equation (4). Similar reasoning results in d^{**} being assigned when $n^{**}=0$. Thus, our approach will always choose a dose that has not yet been assigned to any subjects.

After the last subject has been followed for DLT and response, all of the data is used to compute the values $r_1, r_2, \dots, r_K$, in which r_k is the product of the posterior means of $(1-p_k)$ and q_k . The dose corresponding to the largest of $r_1, r_2, \dots, r_K$ and accompanied by $p \sim \pi^{\max}$ is chosen as the MSD. We also could include a stopping rule if the data indicate that all doses have unacceptable DLT rates; see either Thall and Russell¹ or O’Quigley and Reiner¹⁴ for common choices of stopping rules. We also note that Rosenberger et al.¹⁵ presents alternate distance measures to equation (2) that might prove useful with our design.

3 Numerical studies

3.1 Simulation settings

We have a study designed to determine which of five candidate doses is the MSD in a sample of 35 subjects. We note that we also performed our simulations with a larger sample size of 50 subjects (results not shown) and found no significant change to our results. We did not study smaller sample sizes, but we find a good rule of thumb is six subjects per dose, which is the maximum sample size of the 3+3 method for designing a traditional Phase I trial, so our sample size of 35 seemed reasonable. Nonetheless, an appropriate sample size should be selected after simulations are done with varying sample sizes to select one that has good operating characteristics.

Table 1 contains the skeleton values of the probability of DLT and conditional probability of response. Based on these values, we believe a priori that the lowest dose is the MSD, but also believe that all five doses have very similar success rates, with only a four point difference in the success rates among the five doses. From the skeleton of DLT rates, we are able to compute the dose values $D_1=1.60$, $D_2=0.48$, $D_3=-0.06$, $D_4=-0.80$, $D_5=-1.61$. We have set a value $\pi^{\max}=0.50$ for the maximum acceptable DLT rate, which may seem unusually high to those used to traditional Phase I trial designs. However, we found in simulations that if we used a value of $\pi^{\max}=0.30$, the MSD tended to be identified at doses with DLT rates far below 0.30 and led to sub-optimal performance of the algorithms. In other words, whatever the desired DLT rate is for the MSD, the upper bound $\pi^{\max}$ should be set higher than that value to compensate for the inherent variability that would occur in a sample of 35 subjects. With regard to the prior distribution for the conditional response rates, we set $A=1$, which combined with the skeleton values given in Table 1, leads to a Beta distribution with parameters 0.2 and 0.8 for each q_k . We also set $\rho=0$ so that all the q_k are independent. Once we present the results of the simulations using these parameter values, we will present results for simulations using different values of the skeletons as well as different values of A and ρ .

Table 1 also contains the true rates of DLT, conditional response, and success for each of six settings examined. Each setting has success rates that were selected so as to have the true MSD occur at differing doses among the five settings. Settings 1 and 2 have probabilities of conditional response that are non-monotonic with dose. However, setting 1 has DLT rates that are equal to the skeleton, while setting 2 has DLT rates higher than the skeleton. Setting 3 has DLT rates also equal to the skeleton, but the conditional response rates are actually monotonic in dose. Setting 4 has monotonic DLT and conditional response rates that are generally higher than the skeleton, which are even higher in setting 5. Setting 6 is similar to setting 3 except that in setting 6 the conditional response rates are now higher than those in setting 3, so that responses will be observed more often in setting 6 than in setting 3. Setting 6 is also quite different from the skeleton and allows us to compare the robustness of the different design approaches to a highly misspecified skeleton.

Via simulation over the settings defined in Table 1, we examine our design's operating characteristics, which are the percentage of simulations in which each dose is selected as the MSD and the percentage of subjects assigned to each dose. Within our design, we examine

‘greedy’ dose assignments, fully randomizing among all doses relative to their estimated success rates with $\lambda=2$, and our new two-stage dose assignment algorithm described earlier. We also examined full randomization with $\lambda=1$ and $\lambda=1.5$ (results not shown), but found that the operating characteristics were not materially different than those when using $\lambda=2$. We also applied the method of Ivanova described in the introduction to each setting, as well as a so-called optimal result formulated by Zohar and O’Quigley.⁹ The optimal result is calculated from a virtual design that would not occur in practice but could be used as one uses the Cramer-Rao bound for unbiased estimators. The virtual design supposes that we have the outcomes for each subject if they were to receive every one of the doses. For each dose k , the observed proportion of subjects with DLT is used to estimate p_k and the proportion of responses in subjects with no DLT is used to estimate q_k , which then lead to an estimate of r_k . The dose level with the largest estimate of r_k and estimate of p_k no larger than $\pi^{\max}$ is selected as the MSD. Since the optimal result is based upon all possible information, it can be used as a benchmark for the performance of all of the designs.

At the end of the study, our method selects the MSD as the dose with the largest probability of success among all doses with a probability of DLT no larger than π . Ivanova suggests the following approach for her method for selecting the MSD at the end of the study. We let R denote the ratio of the maximum sample size to the number of doses. Among the doses assigned to at least R subjects, the MSD is chosen as the dose with the highest observed proportion of successes among all doses with a probability of DLT no larger than π . All of our simulations were done in the statistical software package **R** and code is available from the first author upon request.

3.2 Simulation results

Table 2 presents the operating characteristics of the methods averaged over 1000 simulations in each setting. In Settings 1 and 2, where the MSD is located at the third and second dose, respectively, we see that our proposed method with two-stage dose assignments identifies the MSD slightly better than either using randomization of all doses or greedy dose assignments, although the differences rather quite small. The two-stage approach tends to assign slightly fewer subjects to the MSD than either randomization or a greedy approach, although again the differences are slight. What is most striking is that in both settings, our design identifies the MSD and assigns far more subjects to the MSD much more often than Ivanova’s method. In settings like Settings 1 and 2, where the probabilities of DLT and response tend to be low, most subjects will be observed without DLT and without response, which in Ivanova’s algorithm, encourages escalation to higher and higher doses. As a result, as seen in Table 2, Ivanova’s method tends to identify the MSD at doses above the true MSD and assigns a majority of subjects to those doses. Because of this property, we see in Setting 3 that Ivanova’s method performs better than the other designs, both in terms of identifying the MSD and assigning subjects to the MSD. And because the other designs escalate doses much slower than Ivanova’s method, they all assign far more subjects to doses below the MSD than Ivanova’s method.

In Setting 4, we see that use of greedy dose assignments leads to poorer identification of and assignment to the MSD than the other designs, which all tend to have similar operating

characteristics. However, we note that due to the propensity of Ivanova's design toward higher doses, that design assigns far more subjects to the last dose, which is more toxic and less effective than the MSD, than our design with two-stage or randomized dose assignments. In Setting 5, once again use of greedy dose assignments performs poorer than the other methods, which all have similar abilities of identifying the MSD and assigning patients to the MSD.

Up to this point, our simulation results showed little difference between the performance of our two-stage approach and randomization. It is Setting 6 in which a difference between these two approaches is seen. Because the gradient of success rates changes less steeply with dose in Setting 6 than in Setting 3, the probabilities used with randomization all tend to be quite similar, leading to slowed dose escalation and more subjects being assigned to doses other than the MSD. Because our two-stage approach limits the assignment of each subject to two possible doses, instead of all doses, it promotes faster dose escalation and thereby identifies the MSD more often and assigns more subjects to the MSD than when randomizing among all doses. Because the MSD is the last dose, the method of Ivanova performs as well as our two-stage approach.

Relative to the optimal design, we see that all methods tend to perform less well, as expected, although the performance of our design is quite close to the optimal design in Settings 5 and 6, and Ivanova's method actually identifies the MSD more often than the optimal design in Setting 3. However, this result is partially due to the propensity of Ivanova's method for higher doses. Most notably, we see much less variability in the ability of our method to identify the MSD, which ranges from 35% to 54% of simulations, than Ivanova's method, which ranges from 3% to 75% of simulations.

As a first attempt at a sensitivity analysis of our proposed two-stage dose assignment approach, Table 3 contains simulation results for Settings 4 and 5 when three different values for A and three different values of ρ are used. We also performed this sensitivity analysis for the other settings, but do not present those results as they were similar to those presented in Table 3. Surprisingly, the performance of our design appears to be rather insensitive to values of either parameter. Since A measures the prior effective sample size, larger values of A would indicate smaller prior variance, and suggest a more informative prior distribution for $q_1, q_2, \dots, q_K$. Since the skeleton assumed a flat pattern of response rates with dose, we had expected that larger values of A would degrade the performance of our design and lead to selection of lower doses more often than with smaller values of A . The values in Table 3 make only a very modest support of this belief. The value of ρ would indicate how similar the response rates are of adjacent doses and the amount of 'smoothing' that can be done with the data. The results in Table 3 give no indication of any added benefit from smoothing the data and suggest that assumption of independence of $q_1, q_2, \dots, q_K$ is sufficient.

We also examine the sensitivity of our design to the choice of the skeleton for $q_1, q_2, \dots, q_K$. Table 4 contains the operating characteristics for four different skeletons: q_1^* , the skeleton shown in Table 1, q_2^* , the true values for Setting 4, q_3^* , the true values for Setting 5, and q_4^* , the true values for Setting 2. Not surprisingly, we see that our method performs best in both

Settings 4 and 5 when the skeleton matches the true response rates. However, the gain in performance is quite small and overall, our method appears to have little sensitivity to the chosen skeleton of response rates.

4. Discussion

We first emphasize that ours is the only method besides that proposed by Ivanova for dose-finding based on DLT rates and response rates when response is only observed in subjects who do not experience DLT. The novelty of our design is not in the modeling of the DLT rates, which uses a model-based approach that Ivanova also uses implicitly to determine which doses are overly toxic. Rather, the novelty lies in the model used with the response rates, in which no parametric model is assumed for the association of dose with the probability of conditional response. We only assume each conditional response rate is a Beta random variable and smooth the observed response rates from neighboring doses with an auto-regressive covariance matrix. Implicit in our design is the assumption that subjects who experience DLT cannot receive further treatment at lower doses in hopes of eliciting a clinical response. Thus, our design is not appropriate for protocols that allow for intra-patient dose reductions after DLT, as response could now be observed in patients with DLT. Furthermore, the model for response for this situation would be more complex as patients may receive more than one dose, and the times of administration of later doses are conditional upon the time to DLT after the first dose. One could envision using extensions of the methods of Braun et al.¹⁶ and Liu and Braun¹⁷ which directly model multiple administrations into the time to DLT.

We also propose a dose-assignment algorithm that has not been proposed in other designs. Our approach is a compromise between a greedy dose-assignment approach, in which a single dose assignment is possible for each subject, and randomization, in which each subject can be assigned to any of the doses. Our two-stage approach combines the greedy dose assignment with one other dose deemed close to the MSD in order to further explore the doses and gain more confidence that we truly are in a neighborhood of the true MSD. Our method can be viewed as using accumulated information, i.e. variability, to determine dose assignments, which is in the spirit of existing work in optimal design for traditional Phase I designs that are based on maximizing the determinant of the information matrix for the model parameters.¹⁸ Via simulation, we have demonstrated that our proposed two-stage approach to dose assignments performs uniformly better than a traditional greedy approach, performs better than randomization among all doses when success rates are similar among all doses, and performs better than the design of Ivanova when response rates are low and non-monotonic with dose.

Our design is especially timely given the current emphasis on translational outcomes in Phase I designs. Specifically, the response measured in subjects without DLT is no longer the traditional measure of clinical response, e.g. shrinkage of a solid tumor, but is instead the change in a specific biomarker, as described in Le Tourneau et al.¹⁹ Such an outcome is common when studying a therapy or agent that is designed to target a specific biomarker known to be in the pathway in the development or termination of the disease. This change in biomarker would be measured after the period of observation for DLT as evidence for the

activity of the agent. Although our methods would require the biomarker change to be dichotomized, i.e. an indicator that the change is greater than some desired or clinically meaningful change, one interesting avenue of research would be the extension of our methods to use a continuous measure of response.

One issue of randomization and our two-stage approach is that subjects will always have positive probability of being assigned to doses other than the MSD, even near the end of the study, while traditional Phase I designs have focused on assigning patients to the same dose, i.e. the MSD, near the end of the study. However, recent work Azriel et al.²⁰ shows for the traditional setting of dose-finding based solely on toxicity, any design that eventually focuses upon one dose being assigned to subjects cannot asymptotically find the correct MTD with probability 1.0. We would expect these results would hold true for our setting as well.

Our approach to generating correlated random variables that have Beta marginal distributions, although easy to implement, had the drawback that the resulting correlation of the random variables did not equal the value of the parameter used to induce the correlation. An alternate approach is that proposed by Magnussen,²¹ in which correlation of the sample generated is quite close to that desired when the correlation is small. However, one would suspect little improvement to the performance of our method, especially given our finding that non-zero values of ρ tended to lead to operating characteristics similar to those when $\rho=0$.

Our design requires that each subject without DLT must be fully followed for response before the next subject can be enrolled, which could lead to delayed or forgone enrollment of patients if the patient inter-arrival times are short with respect to the length of follow-up for response. However, the likelihood computations in our design could be easily modified to incorporate the data of subjects without DLT who were still being followed for response using the methods proposed in the TITE-CRM²² or the EM-CRM²³ for dose-finding studies based solely on DLT.

Acknowledgments

Funding

The research is supported by the US National Institute of Health grant 5R01CA148713.

References

1. Thall PF, Russell KT. A strategy for dose finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics*. 1998; 54:251–264. [PubMed: 9544520]
2. Braun TM. The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Control Clin Trial*. 2002; 23:240–256.
3. Ivanova A. A new dose-finding design for bivariate outcomes. *Biometrics*. 2003; 59:1001–1007. [PubMed: 14969479]
4. Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*. 2004; 60:684–693. [PubMed: 15339291]
5. Yin G, Li Y, Ji Y. Bayesian dose-finding in phase I/II clinical trials using toxicity and efficacy odds ratios. *Biometrics*. 2006; 62:777–787. [PubMed: 16984320]

6. O'Quigley J, Hughes MD, Fenton T. Dose-finding designs for HIV studies. *Biometrics*. 2001; 57:1018–1029. [PubMed: 11764240]
7. 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]
8. Zohar S, O'Quigley J. Identifying the most successful dose (MSD) in dose-finding studies in cancer. *Pharmaceut Stat*. 2006; 5:187–199.
9. Zohar S, O'Quigley J. Optimal designs for estimating the most successful dose. *Stat Med*. 2006; 25:4311–4320. [PubMed: 16969893]
10. Li Z, Durham SD, Flournoy N. An adaptive design for maximization of a contingent binary response. *IMS Lecture Notes - Monograph Series*. 1995; 25:179–196.
11. Lissitchkov T, Arnaudov G, Peytchev D, et al. Phase-I/II study to evaluate dose limiting toxicity, maximum tolerated dose, and tolerability of bendamustine HCl in pre-treated patients with B-chronic lymphocytic leukaemia (Binet stages B and C) requiring therapy. *J Cancer Res Clin Oncol*. 2006; 132:99–104. [PubMed: 16292542]
12. Lee SM, Cheung YK. Calibration of prior variance in the Bayesian continual reassessment method. *Stat Med*. 2011; 30:2081–2089. [PubMed: 21413054]
13. Thall P, Inoue LYT, Martin TG. Adaptive decision making in a lymphocyte infusion trial. *Biometrics*. 2002; 58:560–568. [PubMed: 12229990]
14. O'Quigley J, Reiner E. A stopping rule for the continual reassessment method. *Biometrika*. 1998; 85:741–748.
15. Rosenberger WF, Stallard N, Ivanova A, et al. Optimal adaptive designs for binary response trials. *Biometrics*. 2001; 57:909–913. [PubMed: 11550944]
16. Braun TM, Yuan Z, Thall PF, et al. Determining a maximum tolerated schedule of a cytotoxic agent. *Biometrics*. 2005; 61:335–343. [PubMed: 16011679]
17. Liu C, Braun TM. Parametric non-mixture cure models for schedule finding of therapeutic agents. *J Roy Stat Soc Ser C*. 2009; 58:225–236.
18. Whitehead J, Brunier H. Bayesian decision procedures for dose determining experiments. *Stat Med*. 1995; 14:885–893. [PubMed: 7569508]
19. Le Tourneau C, Lee JJ, Siu LL. Dose escalation methods in phase I cancer clinical trials. *J Nat Cancer Inst*. 2009; 101:708–720. [PubMed: 19436029]
20. Azriel D, Mandel M, Rinott Y. The treatment versus experimentation dilemma in dose finding studies. *J Stat Plan Infer*. 2011; 141:2759–2768.
21. Magnussen S. An algorithm for generating positively correlated beta-distributed random variables with known marginal distributions and a specified correlation. *Comput Stat Data Anal*. 2004; 46:397–406.
22. Cheung YK, Chappell R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics*. 2000; 56:1177–1182. [PubMed: 11129476]
23. Yuan Y, Yin G. Robust EM continual reassessment method in oncology dose finding. *J Am Stat Assoc*. 2011; 106:818–831. [PubMed: 22375092]

Table 1

Skeleton and actual probabilities of DLT, response, and success for six settings. Boldface indicates MSD

Skeleton values						
Outcome	Dose					
	1	2	3	4	5	
DLT	0.01	0.03	0.05	0.10	0.20	
Response	0.20	0.20	0.20	0.20	0.20	
Success	0.20	0.19	0.19	0.18	0.16	
Actual values						
Setting	Outcome	Dose				
		1	2	3	4	5
1	DLT	0.01	0.03	0.05	0.10	0.20
	Response	0.10	0.20	0.30	0.20	0.10
	Success	0.10	0.19	0.29	0.18	0.08
2	DLT	0.10	0.20	0.30	0.40	0.50
	Response	0.20	0.25	0.25	0.20	0.10
	Success	0.18	0.20	0.18	0.12	0.05
3	DLT	0.01	0.03	0.05	0.10	0.20
	Response	0.03	0.08	0.10	0.15	0.25
	Success	0.03	0.08	0.10	0.14	0.20
4	DLT	0.01	0.05	0.10	0.15	0.20
	Response	0.10	0.21	0.33	0.47	0.44
	Success	0.10	0.20	0.30	0.40	0.35
5	DLT	0.15	0.25	0.40	0.60	0.80
	Response	0.20	0.40	0.60	0.70	0.80
	Success	0.17	0.30	0.36	0.28	0.16
6	DLT	0.01	0.03	0.05	0.15	0.23
	Response	0.40	0.50	0.60	0.70	0.80
	Success	0.40	0.49	0.57	0.60	0.62

DLT: dose-limiting toxicity; MSD: most successful dose.

Table 2

Operating characteristics of methods for identifying the MSD with various dose assignment algorithms. A = proposed design with two-stage dose assignment algorithm; B = proposed design with 'greedy' assignments; C = proposed design with fully randomized assignments; D = Ivanova³ design; E = Zohar and O'Quigley⁹ optimal design. Boldface indicates MSD. First five columns of values are the percentage of simulations in which each dose was selected as the MSD; the last five columns of values are the average proportion of subjects assigned to each dose

Setting	Method	MSD selection					Dose assignments				
		Dose					Dose				
		1	2	3	4	5	1	2	3	4	5
1	A	0.05	0.23	0.50	0.19	0.03	0.14	0.23	0.29	0.21	0.13
	B	0.07	0.25	0.43	0.19	0.06	0.11	0.26	0.38	0.18	0.08
	C	0.05	0.21	0.49	0.21	0.04	0.13	0.22	0.34	0.20	0.10
	D	0.00	0.00	0.03	0.49	0.41	0.03	0.04	0.08	0.21	0.64
	E	0.00	0.17	0.74	0.09	0.00	.	.	.	.	.
2	A	0.25	0.35	0.27	0.12	0.01	0.25	0.26	0.24	0.17	0.08
	B	0.28	0.33	0.27	0.10	0.03	0.31	0.31	0.24	0.11	0.04
	C	0.25	0.32	0.27	0.14	0.01	0.26	0.29	0.24	0.16	0.05
	D	0.05	0.25	0.37	0.26	0.05	0.09	0.18	0.26	0.29	0.19
	E	0.28	0.45	0.23	0.04	0.00	.	.	.	.	.
3	A	0.02	0.10	0.13	0.26	0.49	0.13	0.18	0.19	0.22	0.28
	B	0.03	0.12	0.15	0.25	0.45	0.09	0.16	0.18	0.23	0.34
	C	0.03	0.10	0.14	0.24	0.49	0.12	0.17	0.18	0.22	0.31
	D	0.00	0.00	0.00	0.25	0.75	0.03	0.04	0.06	0.20	0.67
	E	0.00	0.03	0.07	0.20	0.70	.	.	.	.	.
4	A	0.01	0.06	0.20	0.48	0.26	0.09	0.16	0.22	0.29	0.24
	B	0.03	0.15	0.30	0.36	0.17	0.08	0.19	0.28	0.30	0.14
	C	0.01	0.05	0.21	0.48	0.25	0.10	0.16	0.23	0.33	0.19
	D	0.00	0.00	0.07	0.52	0.41	0.04	0.05	0.11	0.28	0.52
	E	0.00	0.01	0.10	0.62	0.27	.	.	.	.	.
5	A	0.06	0.22	0.54	0.17	0.00	0.18	0.26	0.32	0.19	0.04
	B	0.20	0.40	0.35	0.06	0.00	0.25	0.38	0.30	0.06	0.00
	C	0.07	0.24	0.52	0.17	0.00	0.20	0.30	0.36	0.13	0.01

		MSD selection					Dose assignments				
Setting	Method	Dose					Dose				
		1	2	3	4	5	1	2	3	4	5
6	D	0.06	0.34	0.52	0.08	0.00	0.17	0.35	0.35	0.12	0.01
	E	0.02	0.33	0.61	0.04	0.00	.	.	.	.	.
	A	0.03	0.07	0.20	0.34	0.36	0.12	0.16	0.21	0.24	0.26
	B	0.37	0.31	0.20	0.09	0.03	0.39	0.31	0.19	0.08	0.03
	C	0.06	0.11	0.26	0.37	0.20	0.19	0.23	0.26	0.23	0.10
	D	0.01	0.03	0.21	0.39	0.36	0.05	0.08	0.19	0.35	0.32
	E	0.00	0.04	0.23	0.33	0.39	.	.	.	.	.

MSD: most successful dose.

Table 3

Operating characteristics of proposed two-stage dose-assignment design in Settings 4 and 5 with varying parameter values. Boldface indicates MSD. First row for each setting contains values presented in Table 1. First five columns of values are the percentage of simulations in which each dose was selected as the MSD; the last five columns of values are the average proportion of subjects assigned to each dose

Setting	(A, ρ)	MSD selection					Dose assignments				
		Dose					Dose				
		1	2	3	4	5	1	2	3	4	5
4	(1, 0.0)	0.01	0.06	0.20	0.48	0.26	0.09	0.16	0.22	0.29	0.24
	(2, 0.0)	0.01	0.05	0.22	0.48	0.25	0.10	0.16	0.22	0.28	0.24
	(5, 0.0)	0.01	0.06	0.21	0.49	0.24	0.11	0.16	0.23	0.28	0.22
	(1, 0.5)	0.00	0.04	0.19	0.54	0.23	0.08	0.13	0.21	0.32	0.26
	(2, 0.5)	0.01	0.06	0.20	0.54	0.19	0.09	0.13	0.22	0.32	0.25
	(5, 0.5)	0.01	0.05	0.20	0.54	0.20	0.09	0.13	0.23	0.31	0.23
	(1, 0.8)	0.02	0.06	0.18	0.52	0.23	0.09	0.12	0.21	0.33	0.26
	(2, 0.8)	0.01	0.06	0.18	0.52	0.22	0.09	0.12	0.22	0.32	0.25
	(5, 0.8)	0.02	0.06	0.21	0.54	0.18	0.09	0.13	0.24	0.32	0.22
	(1, 0.0)	0.06	0.22	0.54	0.17	0.00	0.18	0.26	0.32	0.19	0.04
5	(2, 0.0)	0.06	0.27	0.53	0.14	0.00	0.19	0.28	0.32	0.18	0.04
	(5, 0.0)	0.08	0.28	0.53	0.11	0.00	0.21	0.28	0.31	0.16	0.03
	(1, 0.5)	0.05	0.22	0.55	0.18	0.00	0.17	0.27	0.32	0.20	0.04
	(2, 0.5)	0.07	0.21	0.57	0.15	0.00	0.18	0.28	0.33	0.19	0.03
	(5, 0.5)	0.10	0.26	0.55	0.09	0.00	0.20	0.30	0.31	0.16	0.03
	(1, 0.8)	0.06	0.23	0.56	0.16	0.00	0.17	0.27	0.32	0.20	0.04
	(2, 0.8)	0.09	0.22	0.57	0.12	0.00	0.19	0.29	0.32	0.18	0.03
	(5, 0.8)	0.14	0.29	0.53	0.04	0.00	0.24	0.31	0.30	0.13	0.02

MSD: most successful dose.

Operating characteristics of proposed two-stage dose-assignment design in Settings 4 and 5 with varying skeleton values for conditional probability of response. q_1^* = skeleton shown in Table 1; q_2^* = true values for Setting 4; q_3^* = true values for Setting 5; q_4^* = true values for Setting 2. Boldface indicates MSD. First row for each setting contains values presented in Table 1. First five columns of values are the percentage of simulations in which each dose was selected as the MSD; the last five columns of values are the average proportion of subjects assigned to each dose

Table 4

		MSD selection					Dose assignments				
Setting	Skeleton	Dose					Dose				
		1	2	3	4	5	1	2	3	4	5
4	q_1^*	0.01	0.06	0.20	0.48	0.26	0.09	0.16	0.22	0.29	0.24
	q_2^*	0.00	0.03	0.18	0.56	0.23	0.05	0.12	0.22	0.34	0.28
	q_3^*	0.00	0.02	0.18	0.53	0.27	0.05	0.12	0.22	0.31	0.29
	q_4^*	0.01	0.08	0.25	0.49	0.18	0.10	0.18	0.25	0.29	0.18
	q_1^*	0.06	0.22	0.54	0.17	0.00	0.18	0.26	0.32	0.19	0.04
5	q_2^*	0.03	0.18	0.54	0.25	0.00	0.14	0.25	0.33	0.24	0.05
	q_3^*	0.04	0.18	0.56	0.23	0.01	0.13	0.25	0.32	0.24	0.06
	q_4^*	0.05	0.23	0.57	0.15	0.00	0.17	0.29	0.33	0.18	0.03
		MSD: most successful dose.									