

# NIH Public Access

Author Manuscript

*Stat Med.* Author manuscript; available in PMC 2012 October 24

## Published in final edited form as:

Stat Med. 2011 July 30; 30(17): 2062–2069. doi:10.1002/sim.4024.

# Extended model-based designs for more complex dose-finding studies

# John O'Quigley<sup>a,\*,†</sup> and Mark Conaway<sup>b</sup>

<sup>a</sup>Inserm, Université Paris VI, Place Jussieu, 75005 Paris, France

<sup>b</sup>University of Virginia, Charlottesville, 22901 VA, U.S.A.

## Abstract

We discuss extensions of model-based designs, such as the continual reassessment method, for use in dose-finding studies. Rather than work with a single model to carry out the design and analysis of a dose-finding study we indicate how the use of several models can greatly increase flexibility. We can appeal to established results on Bayesian model choice and this device makes the inferential problem essentially straightforward. The greater flexibility enables us to take on board many different kinds of added complexity. Examples include extended models to deal with subject heterogeneity, extended models to take account of different treatment schedules and extended models to tackle the problem of partial ordering.

#### Keywords

clinical trials; Bayesian model choice; continual reassessment method; dose escalation; dosefinding studies; extended models; phase 1; safety; toxicity

# 1. Background

Model-based designs, including the continual reassessment method (CRM) [1], have gained in popularity over the last 20 years. One of the specific purposes of the CRM was to meet the ethical requirement of a dose-finding study in which, aside from the usual requirements of statistical accuracy, it is felt necessary to treat every included patient at the current best estimate of some acceptable 'ideal' target dose. Many developments and innovations have followed, the basic method and variants having found a number of other potential applications. We assume that we have available k doses;  $d_1, \dots d_k$ , possibly multidimensional and ordered in terms of the probabilities,  $R(d_i)$ , for toxicity at each of the levels,

i.e.  $R(d_i) \le R(d_j)$  whenever i < j. The most appropriate dose, the 'target' dose in any study and the dose defined to be the 'maximum tolerated dose (MTD)', denoted  $d_0 \in \{d_1, \dots, d_k\}$  is that dose having an associated probability of toxicity,  $R(d_0)$ , as close as we can get to some target 'acceptable' toxicity rate  $\theta$ . Specifically we define  $d_0 \in \{d_1, \dots, d_k\}$  such that

$$|R(d_0) - \theta| < |R(d_\ell) - \theta|, \quad \ell = 1, \dots, k, \quad d_\ell \neq d_0.$$
(1)

The binary indicator  $Y_j$  takes the value 1 in the case of a toxic response for the *j*th entered subject (j = 1, ..., n) and 0 otherwise. The dose for the *j*th entered subject,  $X_j$  is viewed as random taking values  $x_j \in \{d_1, ..., d_k\}; j = 1, ..., n$ . Thus,  $\Pr(Y_j = 1 | X_j = x_j) = R(x_j)$ .

Copyright © 2011 John Wiley & Sons, Ltd.

<sup>&</sup>lt;sup>\*</sup>Correspondence to: Inserm, Université Paris VI, Place Jussieu, 75005 Paris, France.. <sup>†</sup>jmoquigley@gmail.com.

O'Quigley and Conaway

Little is known about  $R(\cdot)$  and, given accumulating observations, we have two goals; first to identify with as great a precision as possible  $d_0$ , second to ensure that each individual is treated in some optimal way given the current state of knowledge at the time of treatment. Usually we take 'optimal' to mean that every patient is given the best dose level, this being defined as the one as close as we can get to  $d_0$ . Of course if we knew  $d_0$  the problem no longer exists and, in practice, we can only hope to treat at the level we believe, in the light of all available information, to be our best bet of being  $d_0$ . In statistical terms the two goals are; (1) estimate  $d_0$  consistently and efficiently and (2) during the course of the study, concentrate as many experiments as possible at and around  $d_0$ , more precisely treat the *j*th included patient at the same level we would have estimated as being  $d_0$  had the study ended after the inclusion of j-1 patients. We model  $R(x_j)$ , the true probability of toxic response at  $X_j = x_j$ ,  $x_j \in \{d_1, \dots, d_k\}$  by

$$R(x_j) = \Pr(Y_j = 1 | X_j = x_j) = E(Y_j | x_j) = \psi(x_j, a) \quad (2)$$

for some one parameter model  $\psi(x_j, a)$  and *a* defined on the set  $\mathscr{A}$ . For every *a*,  $\psi(x,a)$  should be monotone increasing in *x* and, for any *x*,  $\psi(x,a)$  should be monotone in *a*. For every  $d_i$  there exists some  $a_i \in \mathscr{A}$  such that  $R(d_i) = \psi(d_j, a_i)$ , i.e. the one parameter model is rich enough, at each dose, to exactly reproduce the true probability of toxicity at that dose. We have a lot of flexibility in our choice for  $\psi(x,a)$ . The simple choice:  $\psi(d_j,a) = a_i^{\exp(a)}$ , where *i* runs from 1 to *k*, and where  $0 < a_1 < \ldots < a_k < 1$  and  $-\infty < a < \infty$ , has worked well in our experience. The true mechanism generating the observations can be quite removed from our working model overall, but, close to our target, the true situation and our working model coincide. For the six levels studied in the simulations by O'Quigley *et al.* [1] the working model had  $a_1 = 0.05$ ,  $a_2 = 0.10$ ,  $a_3 = 0.20$ ,  $a_4 = 0.30$ ,  $a_5 = 0.50$  and  $a_6 = 0.70$ . Once a model has been chosen and we have data in the form of the set  $\Omega j = \{y_1, x_1, \ldots, y_j, x_j\}$ , the outcomes of the first *j* experiments we obtain estimates  $\hat{R}(d_i)$ ,  $(i = 1, \ldots, k)$  of the true unknown probabilities  $R(d_i)$ ,  $(i = 1, \ldots, k)$  at the *k* dose levels (see below). The target dose level is that level having associated with it a probability of toxicity as close as we can get to  $\theta$ . The dose or dose level  $x_j$  assigned to the *j*th included patient is such that

$$|\widehat{R}(x_j) - \theta| < |\widehat{R}(d_\ell) - \theta|, \quad \ell = 1, \dots, k, \quad d_\ell \neq x_j. \quad (3)$$

This equation should be compared with equation (1). It translates the idea that the overall goal of the study is also the goal for each included patient. The CRM is then an iterative sequential design, the level chosen for the *n*+1th patient, who is hypothetical, being also our estimate of  $d_0$ . After having included *j* subjects, we can calculate a posterior distribution for *a* which we denote by  $f(a,\Omega_j)$ . We then induce a posterior distribution for  $\psi(d_{ji}a)$ , i = 1,...,k, from which we can obtain summary estimates of the toxicity probabilities at each level so that;

$$\widehat{R}(d_i) = \int_{a \in \mathcal{A}} \psi(d_i, a) f(a, \Omega_j) da, \quad i = 1, \dots, k.$$
(4)

Using equation (3) we can now decide which dose level to allocate to the (j+1)th patient. In this context the starting level  $d_i$  should be such that  $\int_{\mathscr{A}} \psi(d_i, u) g(u) du = \theta$ . This may be a difficult integral equation to solve and, practically, we might take the starting dose to be obtained from  $\psi(d_i, \mu_0) = \theta$  where  $\mu_0 = \int_{\mathscr{A}} ug(u) du$ . It is also common practice to reduce the number of integrals we need to evaluate by working with an alternative estimate  $\tilde{R}$   $(d_i) = \psi(d_i, \mu)$ , i = 1, ..., k, where  $\mu = \int_{\alpha \in \mathscr{A}} af(\alpha, \Omega_j) d\alpha$ . There is no obvious value in this apart

O'Quigley and Conaway

from intensive simulation studies where = the second estimate, which approximates the first, reduces the amount of calculation although not even by an order of magnitude in most cases. Given the set  $\Omega_i$  and the log-likelihood  $\mathcal{L}_i(a)$  the posterior density for *a* is;

 $f\left(a|\Omega_{j}\right) = H_{j}^{-1} \exp\left\{\mathscr{L}_{j}\left(a\right)\right\} g\left(a\right), \quad H_{j} = \int_{-\infty}^{\infty} \exp\left\{\mathscr{L}_{j}\left(u\right)\right\} g\left(u\right) du.$ (5)

The dose  $x_{j+1} \in \{d_1, \dots, d_k\}$  assigned to the (j+1)th included patient is the dose minimizing

the Euclidean distance between  $\theta$  and  $\int \psi \{x_j+1, u\} f(u|\Omega_j) du$ . Often it will make little difference if, rather than work with the of the toxicities, we work with the expectation of *a*, thereby eliminating the need for k-1 integral calculations. Thus we treat the (j+1)th included patient at level  $x_{j+1} \in \{d_1, \dots, d_k\}$  such that  $|\theta-\psi\{x_{j+1}, \mu_j\}|$  is minimized where  $\mu_j = \int uf(u|\Omega_j) du$ .

#### 2. Extended model-based designs

Suppose now that, instead of the single model of equation (2), we have some class of models of interest and we denote these models as  $\psi_m(x_j, a)$  for m = 1, ..., M where there are a total of M possible models.

In particular, we might consider

$$\psi_m(d_i, a) = \alpha_{mi}^{\exp(a)}$$
 (i=1,...,k;m=1,...,M) (6)

where  $0 < a_{m1} < ... < a_{mk} < 1$  and  $-\infty < a < \infty$ , as an immediate generalization of the single model described the beginning of the section. Further, we may wish to take account of any prior information concerning the plausibility of each model and so introduce  $\pi(m), m = 1, ..., M$ , where  $\pi(m) \ge 0$  and where  $\sum_{m} \pi(m) = 1$ . In the simplest case where each model is weighted equally, we would take  $\pi(m)=1/m$ .

If the data are to be analyzed under model *m* then, following the inclusion of *j* patients, the logarithm of the likelihood, can be written as

$$\mathscr{L}_{mj}(a) = \sum_{\ell=1}^{j} y_{\ell} \log \psi_m(x_{\ell}, a) + \sum_{\ell=1}^{j} (1 - y_{\ell}) \log (1 - \psi_m(x_{\ell}, a))$$
(7)

where any terms not involving the parameter *a* have been equated to zero. Under model *m* we obtain a summary value of the parameter *a*, in particular the maximum of the posterior mode and we refer to this as  $\hat{a}_{mj}$ . Given the value of  $\hat{a}_{mj}$  under model *m* we have an estimate of the probability of toxicity at each dose level  $d_i$  via:  $\hat{R}(d_i) = \psi_m(d_i, \hat{a}_{mj})$  (i = 1, ..., k). On the basis of this formula, and having taken some value for *m*, the dose to be given to the (j+1)th patient,  $x_{j+1}$ , is determined. Thus, we need some value for *m* and we make use of the posterior probabilities of the models given the data  $\Omega j$ . Denoting these posterior probabilities by  $\pi(m|\Omega j)$  then:

$$\pi\left(m|\Omega_{j}\right) = \frac{\pi\left(m\right)\int_{-\infty}^{\infty}\exp\left\{\mathscr{L}_{mj}\left(u\right)\right\}g\left(u\right)\,\mathrm{d}u}{\sum_{m=1}^{M}\pi\left(m\right)\int_{-\infty}^{\infty}\exp\left\{\mathscr{L}_{mj}\left(u\right)\right\}g\left(u\right)\,\mathrm{d}u} \tag{8}$$

In some cases the  $\pi(m|\Omega_j)$  are only of very indirect interest, such as using several models and then averaging to decide on the best current, running, estimate of the MTD. In other

cases the  $\pi(m|\Omega_j)$  can play a more central role and we will want to say something about *m* itself as we make progress. Once *m*, our indicator over potential models takes values greater than one, then we consider that we are dealing with 'extended' models. There are a large number of potential extensions of the CRM and here we focus on three of them;

- 1. Extensions to deal with patient heterogeneity. A closely related problem from the methodological standpoint is when different patients receive different treatment schedules which may have a potential impact on the probability of encountering a DLT. Although in this case the whole group of treated patients may be homogeneous, we can make use of the methods which account for heterogeneity.
- 2. Extensions that enable us to relax the monotonicity assumption and, in particular, tackle the dose-finding problem using two drugs in which the toxic ordering may not be known.
- **3.** Extension that enable us to make use of several rather than a single CRM model and thereby address to some extent the issue of arbitrariness in model choice.

Of these three problems the first two are very practical ones. In many actual dose-finding studies it is more likely to be the rule than the exception that there exists significant heterogeneity among the patients. The goal of fully accounting for all sources of heterogeneity would coincide with that of individualized dose finding, a laudable although, at this stage, an unrealistic goal. A lesser goal is that of combining patients into rough prognostic groups and some of these immediately suggest themselves. Common examples are dividing the patients into heavily pre-treated and less heavily pre-treated patient groups, or, possibly, dividing the groups on the basis of adult and adolescent groups. The second problem concerns dosing involving more than a single agent when, in order to maintain caution, as we increase one of the component compounds we may decrease the other. The result of this can be that only a partial ordering of the toxic probabilities can be claimed. The exact ordering of the various combinations is unknown and an extended CRM model would allow for this extra flexibility. Problem 3 has both a theoretical and a practical side although essentially the question is a theoretical one: to what extent dose the arbitrariness in any model choice impact final conclusions and, by appealing to a whole class rather than a single arbitrarily chosen model, do we increase robustness or can we be more confident in the final result. We consider the first problem in the following section and then, in subsequent sections, we consider the other two problems.

#### 3. Patient heterogeneity

As in other types of clinical trials we are essentially looking for an average effect. Patients naturally differ in the way they may react to a treatment and, although hampered by small samples, we may sometimes be in a position to specifically address the issue of patient heterogeneity. One example occurs in patients with acute leukemia where it has been observed that children will better tolerate more aggressive doses (standardized by their weight) than adults. Similarly, heavily pre-treated patients are more likely to suffer from toxic side effects than lightly pre-treated patients. In such situations, we may wish to carry out separate trials for the different groups in order to identify the appropriate MTD for each group. Otherwise we run the risk of recommending an 'average' compromise dose level, too toxic for a part of the population and suboptimal for the other. Usually, clinicians carry out two separate trials or split a trial into two arms after encountering the first DLTs when it is believed that there are two distinct prognostic groups. This has the disadvantage of failing to utilize information common to both groups. The most common situation is that of two samples where we aim to carry out a single trial keeping in mind potential differences between the two groups. A multi-sample CRM is a direct generalization although we must remain realistic in terms of what is achievable in the light of the available sample sizes.

The papers [2, 3] focus mostly on models for the two group case, since this case is the most common and there are not usually enough resources, in terms of patient numbers, to deal with more complex structures. Elaborating higher dimensional models, at least conceptually, is straightforward. The dose toxicity model is written;

$$\Pr(Y=1|d_i, z) = \psi(d_i, a, b)$$
 (9)

where the parameter b measures to some extent the difference between the groups. An obvious example that has been used successfully is

$$\psi(d_i, a, b) = \alpha_i^{\exp(a+bz)} \quad (i=1,\dots,k) \quad (10)$$

where, again,  $0 < a_1 < ... < a_k < 1$ ;  $-\infty < a < \infty$ ,  $-\infty < b < \infty$  and *z* is a binary group indicator. Asymptotic theory is cumbersome for these models, but consistency can be shown under restrictive assumptions [2].

An alternative approach, in harmony with the underlying CRM idea of exploiting underparametrized models, is to be even more restrictive than allowed by the above regression models. Rather than allow for a large, possibly infinite, range of potential values for the second parameter *b*, measuring differences between the groups, the differences themselves are taken from a very small finite set. Since, in any event, if the first group finishes with a recommendation for some level,  $d_0$  say, then the other group will be recommended either the same level or some level, one, two or more, steps away from it. The idea is to parameterize these steps directly. The indices themselves are modeled and the model is less cluttered if we work with  $\log \psi(d_{j,a})$  rather than  $\psi(d_{j,a})$  writing;

$$\log\psi(d_i, a) = \exp(a)\log\alpha_{\phi(i)}, \quad \phi(i) = i + zh_i(t) \quad (11)$$

where

$$h_i(t) = tI (1 \le i + t \le k) + kI (i + t > k) + I (i + t < 1), \quad t = 0, 1, 2, \dots, \quad (12)$$

the second two terms in the above expression taking care of edge effects. It is easy to put a discrete prior on t, possibly giving the most weight to t = 0 and only allowing one or two dose level shifts if the evidence of the accumulating data points strongly in that direction. No extra work is required in order to generalize to several groups. Under a condition, analogous to condition 7 (see [2]), applied to both groups separately, consistency of the model in terms of identifying the correct level, can be demonstrated. This is of interest but it is more relevant to study small sample properties, often via the use of simulations, since, for dose-finding studies, samples are invariably rather small.

1. Model 1: *m*=1

$$\Pr(Y=1|d_i, z=0) = \psi(d_i, a), \quad i=1,...,k$$
  
$$\Pr(Y=1|d_i, z=1) = \psi(d_i, a), \quad i=1,...,k$$

2. Model 2: *m=2* 

$$\Pr(Y=1|d_i, z=0) = \psi(d_i, a), \quad i=1, ..., k$$
  

$$\Pr(Y=1|d_i, z=1) = \psi(d_{i+1}, a), \quad i=1, ..., k-1$$
  

$$\Pr(Y=1|d_i, z=1) = \psi(d_k, a), \quad i=k$$

$$\Pr(Y=1|d_i, z=0) = \psi(d_i, a), \quad i=1,...,k$$
  

$$\Pr(Y=1|d_i, z=1) = \psi(d_{i-1}, a), \quad i=2,...,k$$
  

$$\Pr(Y=1|d_i, z=1) = \psi(d_1, a), \quad i=1$$

The above models allow up to a single difference in dose levels between the groups. This difference can be in either direction, corresponding to a situation in which we do not know, or have any reasonably solid knowledge about which of the two groups is likely to fare the worst. At the same time we rule out the possibility that any difference, should one exist, be greater than a single level. It is obviously very straightforward to construct models which would allow for differences up to two or more levels, again in either direction. In addition, nothing hinders us from allowing differences in one direction to be limited to say one level at most whereas, in the other direction, we may allow greater differences than one level. Indeed, we could decide that we will not allow any difference greater than zero in one direction while allowing differences of one or more in the other. This would correspond to the case where we know that should any difference exist it can only be in a given direction. In practice this is likely to be the most common situation, a well-known example being heavily pre-treated and lightly pre-treated patients. The MTD for the heavily pre-treated patients will be no higher than that for the lightly pre-treated patients.

#### 4. Different treatment schedules

Within the context of a Phase I dose-finding study it will sometimes be the case that there is more than one treatment schedule. The dose may be broken down in different ways; given say once a day for 3 days in a given week, given in a single dose for that week, or perhaps given twice weekly at half of the single dose. For a given total amount of dose per week the schedule itself may have an impact upon the probability of observing toxicity at any given dose. One of the authors (J. O. Q.) was involved in the design of a study involving 3 schedules and 2 prognostic groups (heavily and more lightly pre-treated patients). This meant that there were a total of 6 possibly groups and the promoters of the study wanted to carry out 6 Phase I studies in parallel. Including no more than around 20 patients per group means a total number of patients well over 100 and it became clear to the sponsors of the study that this was not feasible. Sharing information across groups can allow for a more efficient use of resources. This information can take the form of a parameter quantifying possible shifts in the MTD between groups. The amount of savings can only be quantified via deeper comparative study but, in the light of preliminary work on the heterogeneity problem [2], it would be possible to reduce the sample size of around 120 described just above to little more than one half of that.

In the practical context of modeling, we would introduce indicator variables, in a way similar to those used to specify the problem of patient heterogeneity. The variable 'schedule' on 3 levels could be represented by 2 binary covariates,  $z_2$  and  $z_3$ . Together with the variable  $z_1$  indicating the degree of prior treatment that would allow for 6 groups in all. We can write the model as,

 $\log\psi(d_{i}, a) = \exp(a)\log\alpha_{\phi(i)}, \quad \phi(i) = i + z_{1}h_{i}(t) + z_{2}h_{i}(s) + z_{3}h_{i}(u) \quad (13)$ 

where, in a way, entirely analogous to the definition of  $h_t(t)$  in equation (12), we allow the variables, *t*, *s* and *u* to take integer values beginning with zero (no effect) but keeping within the range of allowed doses, i.e. we use indicator variables like those in equation (12) to cater for edge effects. Although the problem is becoming quickly more involved than that concerned with two sample heterogeneity, it may still be worth writing out all of the models

as done there. This may be slightly tedious but is straightforward and would be useful in the discussion stage of trial development.

As much as possible of what is known of the physical problem can be incorporated into the design. For example, it may be argued that the shorter the time interval over which the treatment is given, the greater the potential for any increase in the probability of toxicity at any given dose. In this case the above 3 treatment schedules would be ordered and we could reduce the use of two covariates  $z_2$  and  $z_3$  to a single covariate. Again, given the clinicians precise knowledge of the situation, it may be possible to eliminate the notion that the 3 schedules might have 3 distinct MTDs. They may only differ by at most a single level and this can be explicitly expressed via our model construction.

#### 5. Partially ordered dose toxicity models

Conaway *et al.* [4] and Wages *et al.* (2010; under review) proposed methods for phase I trials involving multiple agents in which some of the orderings of the toxicity probabilities between combinations of agents are not known prior to the study. As an example, these papers cite a study [5] involving pacitaxel and carboplatin administered in the combinations in Table I. The ordering for combinations 3 and 5 is not known since combination 3 has a greater dose of paclitaxel but a lower dose of carboplatin than combination 5. Many of the orderings are known. For example, combination 2 has a greater probability of a toxicity than combination 1 because combination 2 has the same dose of pacitaxel and the same dose of carboplatin as combination 1.

The papers [4] (Wages *et al.*, 2010; under review) consider all possible 'simple orders' consistent with the known orderings. A simple order is one in which all orderings between pairs of treatment combinations are known. In the Patnaik (2001) study [5], there are 6 possible simple orders for the toxicity probabilities associated with the treatment combinations.

$$\begin{split} m &= 1: R \ (x_1) \leq R \ (x_2) \leq R \ (x_3) \leq R \ (x_4) \leq R \ (x_5) \leq R \ (x_6) \\ m &= 2: R \ (x_1) \leq R \ (x_2) \leq R \ (x_3) \leq R \ (x_5) \leq R \ (x_6) \leq R \ (x_4) \\ m &= 3: R \ (x_1) \leq R \ (x_2) \leq R \ (x_5) \leq R \ (x_6) \leq R \ (x_3) \leq R \ (x_4) \\ m &= 4: R \ (x_1) \leq R \ (x_2) \leq R \ (x_5) \leq R \ (x_3) \leq R \ (x_4) \leq R \ (x_6) \\ m &= 5: R \ (x_1) \leq R \ (x_2) \leq R \ (x_3) \leq R \ (x_5) \leq R \ (x_4) \leq R \ (x_6) \\ m &= 6: R \ (x_1) \leq R \ (x_2) \leq R \ (x_5) \leq R \ (x_3) \leq R \ (x_6) \leq R \ (x_4) \\ \end{split}$$

Each of the simple orders can be thought of as one of M = 6 possible models.

Using the accumulated data from *j* patients,  $\Omega j$ , the maximum likelihood estimate  $\hat{a}_m$  of the parameter  $a_m$  in 6 can be computed for each ordering m, m = 1, ..., M, along with the value of the log-likelihood (7) at  $\hat{a}_m$ . Wages *et al.* (2010; under review) propose an escalation method that first chooses the ordering with the largest maximized log-likelihood value,  $\mathscr{L}_m(\hat{a}_m)$ . If we denote this ordering by  $m^*$ , the authors use the estimate of  $a_m^*$  to estimate the toxicity probabilities for each treatment combination under ordering  $m^*$ ,

 $\widehat{R}(d_i) = \psi_m^*(d_i, \widehat{a}_m^*)$ , (i = 1, ..., k). The next patient is then allocated to the dose combination with the estimated toxicity probability closest to the target. Wages *et al.* (2010; under review) investigate several variations of this basic design, including two-stage designs and designs that incorporate randomization among the different possible orderings and describe the operating characteristics of their proposed design.

#### 6. Bayesian averaging and maximization for working model selection

The choice of the working model, i.e. the  $a_i$  in the setting up of any CRM design is largely arbitrary. Cheung and Chappell [6] describe how operating characteristics can be less sensitive to certain working model choices. O'Quigley and Zohar [7] indicate that an 'unreasonable' choice may have a negative impact on operating characteristics. Unfortunately it is not easy to provide a sharp and precise definition as to what we mean by 'reasonable' and the only operationally useful definition of a reasonable model would be one that exhibits good robustness properties. Some working models, while respecting the constraints of Shen and O'Quigley [8] required for convergence, might be anticipated to be not reasonable in this sense. Lee and Cheung [9] provide algorithms that can furnish a satisfactory, if not optimal, working model. Their approach is based on that of indifference intervals described in Cheung and Chappell [6]. A somewhat different strategy for tackling the same question was adopted by Yin and Yuan [10]. These authors suggested that, rather than identifying a single working model, we work with a class of working models and make progress by appealing to the technique of Bayesian model averaging (BMA). This technique makes use of the posterior estimates for the relevant toxic probabilities and these are then weighted with respect to the corresponding posterior model probabilities. Daimon and colleagues [11] also considered making use of several working models, selecting via an sequentially adaptive technique based on different criteria. In particular they studied the posterior predictive loss (PPL) [12], the deviance information criterion (DIC) [13] and the posterior model probability (PMP) [10, 14].

To overcome the arbitrariness in pre-specification of a single working model, especially for a phase I trial in which initial guess of the toxicity probabilities are rarely accurate, as well as to avoid poor pre-specification, our proposal consists of the following procedures: (1) to use all elicited or possible working models corresponding to initial guesses of the toxicity probabilities given by investigators before the start of the trial, (2) to update each of them by the CRM simultaneously, (3) to select one working model out of them, automatically and adaptively by using some criterion, during the course of the trial, and (4) to estimate the MTD based on the selected working model and allocate the estimated MTD to each included patient. Different Bayesian model selection criteria can be used, for example, the PPL [12], the DIC [13] or the PMP [10, 14].

Yin and Yuan [10] argue that their approach leads to greater robustness. However, as long as we work with 'reasonable' models (see O'Quigley and Zohar [7] for a definition of reasonable) then it is not likely that we will gain very much in terms of robustness. One intuitive explanation is that we can see the Bayesian averaging as being a process of taking the mean according to a distribution of recommendations based on different models. If, instead, we take a mean based on a distribution across the parameterizations of the models, then this results in a single simple model. In general the mean of a function is not the same as the function of the mean but, under assumptions of local linearity, they are likely to be very close. Thus, the Bayesian averaging will behave in a way close to that of using a particular single model. This argument would also support the idea that, in order to obtain comparable performance, we would not anticipate encountering any penalty in terms of sample size by using Bayesian averaging as opposed to working with a single model. Lee and Cheung [9] tackle the issue in a slightly different way, taking the view that the best approach is via a single simple model but that, given certain operational objectives, we can strive to obtain a particular single model from the available class of models which can effectively meet these objectives.

#### Acknowledgments

We thank the editors and reviewers for suggestions for improving the clarity of the presentation. Partial support for this work was provided by the National Institute for Health grant # NIH/NIAAA RC1 AA019274. This work was also supported in part by NIH/NCI grant 1R01CA142859-01A1 Designs for phase I trials of combinations of agents.

#### References

- 1. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase 1 clinical trials in cancer. Biometrics. 1990; 46:33–48. [PubMed: 2350571]
- O'Quigley J, Shen L, Gamst A. Two sample continual reassessment method. Journal of Biopharmaceutical Statistics. 1999; 9:17–44. [PubMed: 10091908]
- O'Quigley J, Paoletti X. Continual reassessment method for ordered groups. Biometrics. 2003; 59:430–440. [PubMed: 12926728]
- Conaway M, Dunbar S, Peddada S. Designs for single-or multiple-agent phase I trials. Biometrics. 2004; 45:661–669. [PubMed: 15339288]
- Patnaik A, Warner E, Michael M, Egorin M, Moore M, Siu L, Fracasso P, Rivkin S, Kerr I, Litchman M, Oza A. Phase I dose-finding and pharmokinetic study of paclitaxel and carboplatin with oral valspodar in patients with advanced solid tumors. Journal of Clinical Oncology. 2000; 18:3677–3689. [PubMed: 11054441]
- 6. Cheung YK, Chappell R. A simple technique to evaluate model sensitivity in the continual reassessment method. Biometrics. 2002; 58:671–674. [PubMed: 12230003]
- 7. O'Quigley J, Zohar S. Retrospective robustness of the continual reassessment method. Journal of Biopharmaceutical Statistics. 2009; 20:1013–1025.
- O'Quigley J, Shen LZ. Continual reassessment method: a likelihood approach. Biometrics. 1996; 52:673–684. [PubMed: 8672707]
- 9. Lee SM, Cheung YK. Model calibration in the continual reassessment method. Clinical Trials. 2009; 6:227–238. [PubMed: 19528132]
- Yin G, Yuan Y. Bayesian model averaging continual reassessment method in phase I clinical trials. Journal of the American Statistical Association. 2009; 104:954–968.
- 11. Daimon S, Zohar S, O'Quigley J. Bayesian adaptive model-selecting continual reassessment methods in phase I dose-finding clinical trials. Statistics in Medicine. DOI: 10.1002/sim.4054.
- 12. Gelfand A, Ghosh S. Model choice: a minimum posterior predictive loss approach. Biometrika. 1998; 85:1–11.
- Spiegelhalter D, Best N, Carlin B, van der Linde A. Bayesian measures of model complexity and fit. Journal of the Royal Statistical Society, Series B. 2002; 64:583–639.
- Kass RE, Raftery AE. Bayes factors and model uncertainty. Journal of the American Statistical Association. 1995; 90:773–795.

NIH-PA Author Manuscript

Phase I study of paclitaxel and carboplatin in solid tumors.

Combination

67.5

94.5 6

54 6

6

9

9

Agent Paclitaxel Carboplatin

9

**5** 67.5 7.5

4

-

Table I

NIH-PA Author Manuscript