Nonparametric Estimation of Maximum Tolerated Dose (MTD) in the Context of Phase I Clinical Trials

Anup Dewanji*
Applied Statistics Unit
Indian Statistical Institute
203, B. T. Road, Kolkata 700 108, India
(dewanjia@isical.ac.in)

Gopal K Basak
Theoretical Statistics and Mathematics Unit
Indian Statistical Institute
203, B. T. Road, Kolkata 700 108, India
(gkb@isical.ac.in)

Atanu Biswas
Applied Statistics Unit
Indian Statistical Institute
203, B. T. Road, Kolkata 700 108, India
(atanu@isical.ac.in)

Abstract: With the advent of technology, as there seems to be no reason not to implement a trial with a continuous domain of doses and varying step sizes, we provide a non-parametric methodology for such trials that aims to estimate Maximum Tolerated Dose (MTD) in the phase I clinical trials. We further provide extensive simulations for logistic and probit model with various parameter values to obtain a pseudo sample size for better design of the trials. We also illustrate a bootstrap method to obtain an estimate of the variance and the bias of our estimator.

Keywords: Continuous domain of dose, Performance characteristics, Robbins-Monro procedure, Pseudo sample size, Bootstrap estimate.

*Corresponding author

1 Introduction

Dose-response design and studies are conducted in phase I clinical trials to study toxicity in the drug development process. Sequential response-driven designs are usually carried out for this purpose. The patients are sequentially assigned to various dose levels of a drug, one or a group at a time, starting with the lowest dose. Then it proceeds by increasing or decreasing the dose level, or keeping it the same, depending on the particular design followed in the experiment. Skipping a dose is usually not allowed for ethical reason (see Faries, 1994). The goal is to find a dose level, x_α, say, which has a prescribed probability, α, of toxic response, where the responses are binary (toxic/non-toxic). This dose is often referred to as the maximum tolerated dose (MTD). Usually, the specified value of α lies between 0 and 0.5; in phase I trials for oncology, α is usually taken between 0.2 and 0.3. For non-fatal diseases the α is much smaller. The MTD is commonly estimated by determining the most acceptable dose from a pre-specified set of k increasing dose levels $x_1 < \cdots < x_k$ (See, for example, Dixon and Mood, 1954, and, more recently, Storer, 1989). Bueon et al. (2005) recently reviewed designs that have been used in some major studies since 1995.

Using the discreteness of the dose space, one group of designs, called up-and-down design, employs a Markovian technique to move up or down by one level, or stay at the same level, to eventually converge to the estimated MTD. See, for example, Durham and Flournoy (1995), Durham et al. (1995), Giovagnoli and Pintacuda (1998) and Bortot and Giovagnoli (2005). Note that these designs are essentially nonparametric in nature. There have been few other nonparamteric designs based on the isotonic estimate of the dose response relationship, usually called isotonic designs. See, for example, Leung and Wang (2001) and Ivanova et al. (2007).

O'Quigley et al. (1990) have introduced the popular Continuous Reassessment Method (CRM) with a Bayesian rule to choose the next dose level from $\{x_1, \cdots, x_k\}$ which leads to minimum posterior error in estimating the MTD. See also, for example, O'Quigley and Chevret (1991), O'Quigley (1992), Shen and O'Quigley (1996), Gasparini and Eisele (2000), Leung and Wang (2002) and Cheung and Chappel (2002). Another Bayesian design, given by Babb et al. (1998), uses a loss function to minimize the predicted amount by which any given patient is overdosed. Some other Bayesian designs are described in McLeish and Tosh (1990), Whitehead and Brunier (1995) and Whitehead and Williamson (1998).

Now-a-days, with the advent of technology, there seems to be no reason why a trial cannot be implemented with a continuous domain of dose and varying step sizes. The estimation of MTD in such case can be made more accurate. Mats et al. (1998) have

considered a continuous dose space for the purpose of finding locally optimal designs to estimate MTD assuming a general class of parametric dose response models including logistic, probit, complementary log-log, etc.. According to their findings, the support points of the optimal designs may not be contained in the considered dose levels if they form a finite lattice. Rosenberger et al. (2001) have considered a similar problem in the Bayesian context, while Liu et al. (2006) have studied sequential designs for logistic dose response model for both continuous and discrete dose spaces. Storer (1989, 1993) has considered logistic dose response model with discrete dose space to estimate the model parameters which are then used to estimate the MTD; this estimate may not be one of the considered dose levels. Li et al. (1995) have discussed the possibility of continuous dose space in the context of phase II trial assuming some parametric models.

In the present paper, we propose a method of estimating MTD with continuous domain of dose and variable step size using a procedure similar to Robbins and Monro (1951) without making any particular parametric assumption for the dose response model. We also demonstrate by means of a simulation study that, using variable step size, the estimated MTD can be very close to the true value, so that some efficiency can be achieved in the subsequent phase II trial. In practice, this procedure takes bigger step size at the beginning while tapering down gradually. In order to achieve nearness to the target dose faster, we calibrate the step size with a number of, say k, previous responses by making it dependent on the number of similar responses; see the next section for details.

The rest of the paper is organized as follows. In Section 2, we discuss our methodology. Section 3 describes some performance characteristics which are examined through a simulation study in Section 4. Section 5 ends with some concluding remarks.

2 Methodology

Let Y be the binary response variable representing the toxic response of interest due to application of a drug at a certain dose with $Y = 1$ meaning the toxic response. Let $P(x)$ be the probability of toxic response at dose $X = x$; that is, $P(x) = P[Y = 1|X = x]$. For a given α lying between 0 and 1, the objective is to find the MTD x_α such that $P(x_\alpha) = \alpha$ with the form of $P(\cdot)$ being unknown.

Let x^* be a level of dose which, from previous studies or animal trials, is known to be highly toxic. In practice, x^* is greater than x_α. Ideally, this x^* should be chosen as close to x_α as possible. We assume x^* to be, say, the qth percentile of $P(x)$. It is sufficient to know x^*, although it helps to have information on q as well. At present, we do not use

any information on q; but knowledge of x^* will be used. We assume that the participants are subjected to the drug sequentially at different doses (chosen adaptively depending on the previous doses and responses). Let x_i be the dose for the ith participant, for $i = 1, \cdots, n$, where n, the number of participants, is considered fixed in advance for the time. Although, in practice, there is a limit on the number of patients, n, we shall see later that n need not be fixed in advance and a reasonable stopping rule can be deviced to end the trial. It is a good idea to choose x_1, dose of the first participant, in the neighbourhood of x_α; however, since x_α is not known, we suggest choosing x_1, from our previous knowledge on animal trials, etc., in such a way that it provides a small risk of toxic response. If no such information is available, one may choose $x_1 = 0$ conservatively.

A generalization of Robbins-Monro (1951) procedure suggests the following adaptive formula for x_{i+1}, for $i = 1, \cdots, n$:

$$x_{i+1} = \max\left[x_i - C_i a_i (y_i - \alpha), 0\right], \tag{1}$$

where y_i is the observed value of Y for the ith participant, a_i's are positive constants such that $\sum_{i=1}^{\infty} a_i = \infty$ and $\sum_{i=1}^{\infty} a_i^2 < \infty$ (see Joseph, 2004), and C_i's are some bounded positive constants as described below. Under such conditions, this procedure has been proved to be consistent in the sense that the resulting estimate converges to the true MTD in probability as $n \to \infty$.

Initially, we take $C_i = C$ for $i = 1, \cdots, k$, where k (=5, say) is a suitably chosen integer. This constant C represents the length of jump in successive dose levels in either direction, modulated by a_i, for the first k participants. The idea is to observe and note the trend of recent k changes in dose level to make a decision about the next dose level. Then, for $i > k$, depending on the trend in the last k dose levels, the next jump length is determined as a multiple of C (to preserve the unit of dose measurements), given by

$$C_i = C(1 + \delta_i), \tag{2}$$

where

$$\delta_i = \sum_{l=i-k}^{i-1} |I\{d_l \geq 0\} - I\{d_l < 0\}|$$

with $d_1 = x_1$ and $d_l = x_l - x_{l-1}$. That is, if the successive dose levels (for the recent k participants) either increase or decrease (corresponding to non-toxic or toxic responses, respectively), the δ_i will be k (the largest value) leading to a big jump; on the other hand, if there are several direction changes in dose levels (for the recent k participants), the δ_i will be smaller (at least being 0) leading to a small jump. In order to choose a value of C, we make the assumption that if, theoretically, starting with x_1, all the n^* responses (with a suitably chosen n^* satisfying some optimality criterion) are non-toxic,

the selected dose level for the (n^*+1)st participant will be x^*. Then, from (1), we have

$$x^* = x_{n^*+1} = x_1 + \alpha \sum_{l=1}^{n^*} C_l a_l = x_1 + \alpha C \left[\sum_{l=1}^{k} a_l + \sum_{l=k+1}^{n^*} (1+k) a_l \right].$$

This gives us

$$C = \frac{x^* - x_1}{\alpha \left[\sum_{l=1}^{k} a_l + \sum_{l=k+1}^{n^*} (1+k) a_l \right]}. \tag{3}$$

Since x_1 is chosen conservatively and x^* is chosen to have high probability of toxic response, we always have $x^* > x_1$ and, hence, $C > 0$. The pseudo sample size n^* is chosen in an optimal manner in some sense, as described in the following. It is to be noted that the value of C, hence the successive jump lengths, decreases as n^* increases leading to possible under-estimation. Therefore, larger value of n^* makes the estimate more conservative, whereas smaller value of n^*, leading to larger jump lengths, runs the risk of over-estimation. In practice, since there may be both toxic and non-toxic responses, larger jump lengths (due to smaller n^*) lead to more variance and, on the other hand, smaller jump lengths (due to larger n^*) lead to more bias. We see this pattern in our simulation study also. Figure 1 illustrates this phenomenon by means of three typical paths of the adaptive procedure (1) with $n = 50$ and corresponding to $n^* = 125, 50$ and 20, respectively, in the left, middle and right panel. These paths are generated from a fixed logistic dose response curve with $a = -2$ and $b = 0.1$ (see Section 4) using $x^* = 33.86$ with $q = 0.8$ and $k = m = 5$. From the corresponding entry in Table 1, the 'optimal' (see the next paragraph and also Section 4) n^* lies between 40 and 75; therefore, the figure in the middle panel corresponds to a proper choice of n^*.

We suggest finding the 'optimal' value of n^* by minimizing the mean squared error (MSE) of the estimated x_α. The 'optimal' value of n^* may depend on the parameters of the dose response curve (such as, intercept and slope) in addition to the values of α and q. It is seen, by means of an extensive simulation study in Section 4, that the dependence of 'optimal' value of n^* on q and the slope parameter is rather weak. We give, in Table 1, a guideline for the choice of 'optimal' n^*, for different values of α, q, intercept and slope of $P(x)$, which minimizes the corresponding estimated MSE. The experimenter is likely to have some idea about the intercept parameter, from previous studies or animal trials, to help in choosing n^* from our guideline, for a given α.

Once the constant C is chosen, as a function of n^*, x_1, x^* and α, the adaptive procedure of determining the successive dose levels for the n participants can be carried out using (1). The final dose x_{n+1} can be taken as an estimate of x_α. We, however, suggest the mean of the last m dose levels (that is, $x_{n-m+2}, \cdots, x_{n+1}$) as an estimate of x_α, for some suitably chosen m, to reduce the variability. Let us denote this estimate by $\hat{x}_\alpha$.

In order to obtain an estimate of the variance of this estimate, we suggest a bootstrap sampling method as described in the following. Using the observed values of the dose levels $x_1, \cdots, x_n$ and the corresponding responses $y_1, \cdots, y_n$, we find a smooth parametric estimate of the dose response curve $P(x)$ by assuming a specific model, say, logistic. The difficulty lies in the fact that the successive x_i's are not independent. Therefore, the estimated $P(x)$, denoted by $\hat{P}(x)$, can be considered to be an approximation to be used only for the purpose of bootstrap sampling. For each bootstrap sample, say the bth, we start with $x_{b1} = x_1$ and simulate y_{b1} from the *Bernoulli* distribution with success probability $\hat{P}(x_{b1})$. Then, choose the successive bootstrap dose levels by using the formula, similar to (1),

$$x_{b,i+1} = \max\left(x_{bi} - C_i a_i (y_{bi} - \alpha), 0\right),$$

for $i = 1, \cdots, n$. Then, the bth bootstrap estimate of x_α is given by the mean of $x_{b,n-m+2}, \cdots, x_{b,n+1}$, which is denoted by $\hat{x}_{b\alpha}$. We carry out this bootstrap sampling B (say, 200) number of times to obtain $\hat{x}_{1\alpha}, \cdots, \hat{x}_{B\alpha}$. The sample variance of these bootstrap estimates is used as an estimate of variance of $\hat{x}_\alpha$. This estimated variance has been observed, in the simulation study in Section 4, to be less than the sample variance for about 80-90% of the times. One possible explanation of this is the assumption of independence between the observed x_i's, when they are not. The bootstrap estimates $\hat{x}_{1\alpha}, \cdots, \hat{x}_{B\alpha}$ also provide an estimate of bias of $\hat{x}_\alpha$, as given by

$$\frac{1}{B}\sum_{b=1}^{B} \hat{x}_{b\alpha} - \hat{x}_\alpha.$$

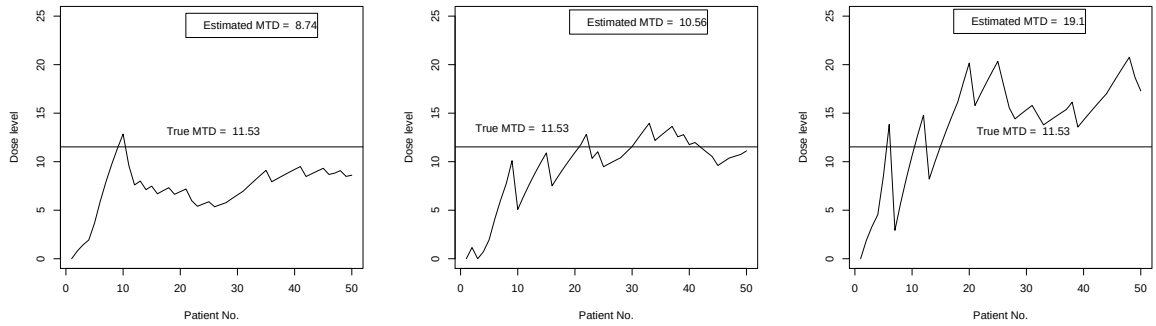

Figure 1. A typical path of the generalized Robbin-Monro procedure with $n = 50$ using the logistic model with a=-2 and b=0.1. The horizontal line gives the true MTD. The left, middle and right panels correspond to $n^* = 125, 50$ and 20, respectively.

3 Some Performance Characteristics

Although the objective of this work is to estimate the MTD, x_α, using continuous dose regimen and variable step size, as closely as possible to the true value, it should not be at the cost of too many adverse toxic responses during the trial. The idea is to be able to estimate MTD from below rather than moving back and forth on either side of the true value. The construction of the C_i's, as in (2), representing the variable step size, and the algorithm for updating the dose, given by (1), have been suggested to minimize this 'loss' due to likely toxic responses at the selected dose levels above the MTD. The purpose of this section is to introduce some performance characteristics in terms of which this particular behavior of our methodology can be studied.

One standard choice is bias, given by $E[\hat{x}_\alpha - x_\alpha]$. Ideally, we like this to be zero; if not, this should be negative so that the estimate is conservative (that is, not an over-estimate). Along with bias, we also consider variance and MSE of $\hat{x}_\alpha$ and like them to be small as a mark of good performance.

One natural choice to measure the performance of our methodology is the expected proportion of toxic responses out of the n observations, denoted by $PTOX$. In order to be ethically proper, we like this $PTOX$ to be small. The next one is the expected proportion of selected dose levels (excluding x_1) which are larger than the MTD, x_α. Let us denote this by $PROP$, which we also like to be small. We also consider the average (over the n observations) of the expected amount of overdoses, $(x_i - x_\alpha)$'s, given $x_i \geq x_\alpha$, and denote it by $MDIFF$. We like this to be small to ensure that, even if we select a dose larger than x_α, it is not too large. In practice, however, we are more concerned with the expected additional risk of toxicity associated with a dose larger than x_α, rather than the actual dose difference. Therefore, we consider $PDIFF$, the average of the probability differences, $(P(x_i) - \alpha)$'s (additional risk), for those dose levels which are larger than x_α. We like this average additional risk to be small as well.

Formally, these four characteristics can be written as

$$\begin{aligned}
PTOX &= \frac{1}{n}\sum_{i=1}^{n} P[Y_i = 1], \\
PROP &= \frac{1}{n}\sum_{i=2}^{n+1} P[X_i > x_\alpha], \\
MDIFF &= \frac{1}{n}\sum_{i=2}^{n+1} E[X_i - x_\alpha | X_i > x_\alpha], \quad \text{and} \\
PDIFF &= \frac{1}{n}\sum_{i=2}^{n+1} \left(P[Y_i = 1 | X_i > x_\alpha] - \alpha\right),
\end{aligned}$$

where Y_i denotes the random variable representing the response of the ith participant, for $i = 1, \cdots, n$, and X_i represents the ith selected dose, for $i = 2, \cdots, n+1$. These quantities may be derived in theory in terms of the unknown dose response function $P(x)$. This derivation, however, becomes increasingly difficult with large n because of the complicated dependence structure among the selected doses X_i's and the responses Y_i's. Nevertheless, based on one simulated sample with known x_α and dose response $P(\cdot)$, these can be readily estimated as

$$\begin{aligned}
\widehat{PTOX} &= \frac{1}{n}\sum_{i=1}^{n} I\left(y_i = 1\right), \\
\widehat{PROP} &= \frac{1}{n}\sum_{i=2}^{n+1} I\left(x_i > x_\alpha\right), \\
\widehat{MDIFF} &= \frac{1}{n}\sum_{i=2}^{n+1} \left(x_i - x_\alpha\right) I\left(x_i > x_\alpha\right), \quad \text{and} \\
\widehat{PDIFF} &= \frac{1}{n}\sum_{i=2}^{n+1} \left(P(x_i) - \alpha\right) I\left(x_i > x_\alpha\right).
\end{aligned}$$

The averages of these estimates over a number of simulated samples can be taken as the estimates of the corresponding quantities, as reported in the next section.

4 A Simulation Study

In our simulation study, we assess the properties of the suggested estimate of x_α, developed in Section 2, in terms of the performance characteristics described in the previous section. We consider both logistic and probit dose response model for $P(x)$, given by $P(x) = [1 + e^{-(a+bx)}]^{-1}$ and $P(x) = \Phi(a + bx)$, with $b > 0$, respectively, where $\Phi(\cdot)$ denotes the cumulative distribution function for a standard normal variate. We need to choose a such that $\alpha > P(0)$. This results in $a < \log(\alpha/(1-\alpha))$ for the logistic model and $a < \Phi^{-1}(\alpha)$ for the probit model. We consider two values of α, namely, 0.2 and 0.3. We also consider different values for the intercept and slope parameters a and b, respectively. We take $a = -2,\ -5,\ -10$ and $b = 0.05,\ 0.1,\ 0.5,\ 2.0$ to represent different degrees of background rate and dose response. We choose x^* as the qth percentile, x_q, of the dose response curve, given by $x^* = (\log(q/(1-q)) - a)/b$ for the logistic model and $x^* = (\Phi^{-1}(q) - a)/b$ for the probit model, with $q = 0.5$ and 0.8. We take $a_i = (1+i)^{-r}$ with $r = 0.9$ (see (1)) and $k = m = 5$.

One simulation run starts with $x_1 = \max(0, x_\epsilon)$, where x_ϵ denotes the ϵth percentile of the dose response curve, with $\epsilon = 0.01$. This is to reflect the choice that information about the minimum risk, say ϵ, leads to $x_1 = x_\epsilon$ and no information leads to $x_1 = 0$. We

generate the response variables Y_i's from a *Bernoulli* trial with probability of success $P(x_i)$ and obtain the subsequent dose levels x_{i+1}'s using the formula (1), for $i = 1, \cdots, n$, with $n = 30,\ 50$ and 100. In order to obtain the successive jump lengths, we derive C using (3) with n^* given by $[nw]$, where $[x]$ denotes the largest integer less than or equal to x. We take $w = 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4, 1.5, 2.0$ for our simulation study to investigate which value of n^* leads to minimum estimated MSE.

We carry out 1000 such simulation runs for a fixed set of parameters (a, b), values of $\alpha,\ q$, sample size (n) and w. For each simulation run, we record $\hat{x}_\alpha$, the estimate of x_α, as described in Section 2. Since computation of variance estimate of $\hat{x}_\alpha$ by bootstrap sampling method (see Section 2) is somewhat extensive, we do not perform this task for all values of w in our simulation study. Instead, we adopt the following strategy. We first estimate the mean and variance of $\hat{x}_\alpha$ by the average and sample variance based on the 1000 estimates, $\hat{x}_\alpha$'s, from the 1000 simulation runs. This leads to estimates of bias and MSE of $\hat{x}_\alpha$ based on simulation. For a fixed set of dose response parameters (a, b), values of $\alpha,\ q$ and sample size n, we identify the value of w, or n^*, that gives the minimum value of estimated MSE of $\hat{x}_\alpha$. The bias-variance trade-off, as discussed in Section 2, gives such a minimum value. However, in order to deal with the sampling variation, we consider all those values of w, or n^*, which lead to values of estimated MSE within 10% of the minimum estimated MSE (treating them as 'plausibly optimal' values of n^*). Then, we take the mid-value of these n^* values to be considered as the 'optimal' value of n^*. In Table 1, we present this range of 'plausibly optimal' n^* values for different values of $\alpha,\ q,\ (a, b)$ and n, and for both logistic and probit model, to be considered as a guideline for choosing n^*. Since larger value of n^* leads to conservative estimate (as discussed in Section 2), one may want to choose the largest one from this range.

The 'plausibly optimal' n^* values seem to depend weakly on n and, in general, there is an increasing trend with n. Since larger value of n^* means smaller jump, this increasing trend with n makes the method more meaningful. The dependence of the 'plausibly optimal' n^* values on α is strong, as expected. But the dependence on q is not so strong; this is desirable as the practitioners often choose the value of x^*, for a particular value of q, from the previous experience. Another desirable outcome is that the dependence on the slope (b) is apparently weak; the practitioners often will not have any information on b. There seems to be some dependence on the intercept (a); the case with $a = -2$ giving larger values for 'plausibly optimal' n^*. Therefore, when there is higher background response probability, the jump size is smaller; this is a desirable feature acting as a safeguard against possible over-estimation. Often, based on prior information or experimental data, there may be some information on the background response probability,

or intercept a. The logit and probit models seem to give similar values, except for the case with $a = -2$ (that is, when the background response probability is relatively high). This background response probability is about 12% for logit model and about 2% for probit model. Based on the prior information on the background response probability, as mentioned before, one may be able to choose one model from the other for the purpose of selecting the 'plausibly optimal' n^* values. It may be noted that the logit model giving higher background response probability (with $a = -2$) also desirably leads to larger values of 'plausibly optimal' n^*. In general, when the background response probability is not high ($a = -5$ or -10), the dependence of the 'plausibly optimal' n^* values on the other parameters (except α) is rather weak, giving the methodology more flexibility.

Table 1. 'Plausibly optimal' values of n^*.

				Logit Model			Probit Model		
a	b	q	α	$n = 30$	$n = 50$	$n = 100$	$n = 30$	$n = 50$	$n = 100$
-2	0.05	0.5	0.2	39-60	65-100	80-200	1-30	20-40	20-50
	0.05	0.5	0.3	12-39	20-55	20-70	1-15	15-20	10-20
	0.05	0.8	0.2	33-60	60-100	100-200	21-60	30-60	50-90
	0.05	0.8	0.3	27-45	45-100	40-60	12-18	15-35	20-40
-2	0.1	0.5	0.2	36-60	60-100	70-200	15-36	25-35	20-40
	0.1	0.5	0.3	21-36	25-50	30-50	1-15	15	20
	0.1	0.8	0.2	45-60	55-100	100-200	27-45	25-75	60-100
	0.1	0.8	0.3	21-60	40-75	40-90	1-24	10-35	20-40
-2	0.5	0.5	0.2	33-60	50-100	50-200	1-45	20-35	20-50
	0.5	0.5	0.3	1-60	15-50	20-60	1-12	2-20	20-30
	0.5	0.8	0.2	30-60	65-100	80-200	24-60	25-100	40-100
	0.5	0.8	0.3	21-60	30-100	30-100	12-21	20-30	20-40
-2	2	0.5	0.2	39-60	55-100	60-200	21-30	15-40	20-50
	2	0.5	0.3	1-60	20-50	30-50	1-15	2-15	20
	2	0.8	0.2	33-60	40-100	120-200	27-45	30-75	40-90
	2	0.8	0.3	27-60	25-75	30-110	1-24	20-30	20
-5	0.05	0.5	0.2	1-21	20-30	20-40	1-21	20-25	20-40
	0.05	0.5	0.3	1-9	2-15	10-20	1-15	2-15	20
	0.05	0.8	0.2	15-33	25-60	30-60	18-39	25-55	30-70
	0.05	0.8	0.3	1-21	15-25	20-30	1-18	15-30	20-30
-5	0.1	0.5	0.2	1-24	15-30	20-30	1-30	20-30	20-40
	0.1	0.5	0.3	1-15	2-15	20	1-15	2-15	20
	0.1	0.8	0.2	15-39	25-40	40-50	18-39	25-60	40-70
	0.1	0.8	0.3	1-18	15-25	20-30	1-21	20-25	20-30
-5	0.5	0.5	0.2	1-21	20	20-30	12-30	20-25	20-40
	0.5	0.5	0.3	9-15	2-15	10-20	1-12	2-15	10-20
	0.5	0.8	0.2	15-30	30	30-50	24-45	30-45	30-80
	0.5	0.8	0.3	1-21	20	20-30	1-18	15-25	20-40
-5	2	0.5	0.2	1-24	15-20	20-30	1-24	15-25	20-40
	2	0.5	0.3	1-15	15	20	1-9	10-15	10-20
	2	0.8	0.2	18-39	25-40	30-60	18-39	25-60	30-90
	2	0.8	0.3	12-18	15-20	20-30	12-21	15-30	20-40
-10	.05	0.5	0.2	1-24	15-25	20-30	1-27	15-30	30
	.05	0.5	0.3	9-15	15	10-20	1-15	2-15	10-20
	.05	0.8	0.2	18-39	25-45	30-60	18-45	35-50	30-60
	.05	0.8	0.3	1-21	15-25	20-30	1-18	20-25	20-40
-10	.1	0.5	0.2	18-21	15-25	20-40	1-27	15-30	20-30
	.1	0.5	0.3	9-15	2-15	20	1-15	10-15	10-20
	.1	0.8	0.2	15-39	25-40	30-50	18-45	25-55	40-80
	.1	0.8	0.3	12-15	15-25	20-30	1-21	20	20-40
-10	.5	0.5	0.2	15	15-25	20-40	1-24	15-30	20-30
	.5	0.5	0.3	1-12	2-15	10-20	1-12	10-15	10-20
	.5	0.8	0.2	15-27	20-50	30-60	24-39	30-65	30-70
	.5	0.8	0.3	1-18	15-20	20-30	12-24	15-20	20-40
-10	2	0.5	0.2	12-15	15-25	20-40	1-24	15-30	20-40
	2	0.5	0.3	1-12	2-10	20	1-12	2-15	20
	2	0.8	0.2	21-24	20-40	40-60	18-60	20-65	30-70
	2	0.8	0.3	1-21	15-25	20-30	12-21	20-25	30

The next phase of the simulation study considers only the optimal n^*, as discussed before. For each simulation run, we obtain $\hat{x}_\alpha$ along with its bias and variance estimate by a bootstrap sampling method, as discussed in Section 2, with $B = 200$ bootstrap samples. We also record $\widehat{PTOX}$, $\widehat{PROP}$, $\widehat{MDIFF}$ and $\widehat{PDIFF}$, as described in Section 3. We then take averages of these quantities over the 1000 simulation runs. The sample variances of the last four quantities are also computed. The bias and variance of $\hat{x}_\alpha$ are also estimated based on the sample of 1000 estimates of x_α. The averages of the bootstrap estimates of bias and variance over the 1000 simulation runs tend to be lower than the corresponding sample versions and the difference reduces with increasing sample size (n). While the mean of the bootstrap estimates of bias is much lower than the sample estimates, the corresponding variance estimate is lower by a factor, generally ranging between 0.6 to 1.0. Also, as remarked in Section 2, about 80 to 90% of the 1000 bootstrap estimates of the variance, depending on the sample size, tend to be smaller than the corresponding sample versions.

In Tables $2a, b$ and $3a, b$ for logit and probit models, respectively, we present the simulation based estimates of the quantities x_α, $PTOX$, $PROP$, $MDIFF$ and $PDIFF$, along with the corresponding standard errors in parentheses. The standard error of $\hat{x}_\alpha$, obtained by the bootstrap method, is reported. We have calculated this bootstrap estimates by assuming *logit*, *probit* and *complementary log−log* models for dose response for the purpose of bootstrap sampling (see Section 2). These estimates turn out to be very similar and we report only the ones corresponding to the *logit* model. The simulation based variance estimates (not reported here) are generally 1 to 1.7 times larger than the bootstrap estimates depending on the sample size. For limitation of space, we report only the estimates corresponding to $a = -2,\ -5$, $b = 0.05,\ 0.5,\ 2.0$ and $n = 30,\ 100$. The results corresponding to $a = -10$ are qualitatively similar to those for $a = -5$ and those with $n = 50$ lie between the corresponding results for $n = 30$ and $n = 100$. The results for $b = 0.1$ similarly lie between those for $b = 0.05$ and $b = 0.5$.

Table 2a. Estimated performance characteristics with the optimal n^* for the logit model with $a = -2$.

			$\alpha = 0.2$			$\alpha = 0.3$		
	Esti-		$b = 0.05$	$b = 0.5$	$b = 2.0$	$b = 0.05$	$b = 0.5$	$b = 2.0$
q	mates	n	x_α=12.27	x_α=1.23	x_α=0.31	x_α=23.05	x_α=2.31	x_α=0.58
0.5	$\hat{x}_\alpha$	30	8.562	0.865	0.214	14.507	1.376	0.351
			(3.533)	(0.377)	(0.091)	(4.739)	(0.443)	(0.111)
		100	8.391	0.865	0.212	16.361	1.679	0.414
			(2.792)	(0.279)	(0.066)	(3.778)	(0.383)	(0.101)
	$\widehat{PTOX}$	30	0.160	0.162	0.159	0.188	0.183	0.181
			(0.051)	(0.052)	(0.050)	(0.049)	(0.052)	(0.053)
		100	0.164	0.164	0.161	0.208	0.212	0.212
			(0.027)	(0.027)	(0.029)	(0.032)	(0.032)	(0.032)
	$\widehat{PROP}$	30	0.157	0.160	0.152	0.046	0.033	0.036
			(0.240)	(0.239)	(0.234)	(0.133)	(0.112)	(0.120)
		100	0.104	0.116	0.106	0.032	0.044	0.038
			(0.237)	(0.242)	(0.232)	(0.127)	(0.152)	(0.142)
	$\widehat{MDIFF}$	30	2.201	0.241	0.057	1.860	0.134	0.044
			(1.882)	(0.187)	(0.048)	(1.678)	(0.126)	(0.041)
		100	1.292	0.123	0.029	1.011	0.105	0.029
			(1.281)	(0.125)	(0.030)	(1.127)	(0.098)	(0.030)
	$\widehat{PDIFF}$	30	0.019	0.021	0.020	0.020	0.014	0.019
			(0.017)	(0.017)	(0.017)	(0.019)	(0.014)	(0.019)
		100	0.011	0.010	0.010	0.011	0.011	0.012
			(0.011)	(0.011)	(0.010)	(0.012)	(0.011)	(0.013)
0.8	$\hat{x}_\alpha$	30	10.871	1.115	0.274	17.108	1.679	0.404
			(5.338)	(0.536)	(0.130)	(5.414)	(0.559)	(0.136)
		100	10.587	1.038	0.261	19.803	1.899	0.470
			(3.406)	(0.367)	(0.089)	(4.221)	(0.418)	(0.100)
	$\widehat{PTOX}$	30	0.178	0.177	0.175	0.204	0.205	0.205
			(0.048)	(0.049)	(0.048)	(0.050)	(0.050)	(0.051)
		100	0.178	0.179	0.178	0.241	0.232	0.231
			(0.028)	(0.027)	(0.028)	(0.035)	(0.035)	(0.033)
	$\widehat{PROP}$	30	0.292	0.306	0.298	0.129	0.110	0.106
			(0.283)	(0.292)	(0.281)	(0.226)	(0.205)	(0.198)
		100	0.280	0.269	0.266	0.175	0.123	0.113
			(0.334)	(0.326)	(0.329)	(0.285)	(0.246)	(0.238)
	$\widehat{MDIFF}$	30	3.960	0.422	0.101	3.054	0.260	0.058
			(3.013)	(0.324)	(0.073)	(2.545)	(0.203)	(0.050)
		100	2.197	0.223	0.057	1.938	0.174	0.041
			(1.929)	(0.200)	(0.047)	(1.741)	(0.162)	(0.042)
	$\widehat{PDIFF}$	30	0.035	0.038	0.036	0.034	0.029	0.025
			(0.029)	(0.032)	(0.028)	(0.029)	(0.023)	(0.023)
		100	0.019	0.019	0.019	0.021	0.019	0.018
			(0.017)	(0.018)	(0.017)	(0.020)	(0.018)	(0.019)

Table 2b. Estimated performance characteristics with the optimal n^* for the logit model with $a = -5$.

			$\alpha = 0.2$			$\alpha = 0.3$		
	Esti-		$b = 0.05$	$b = 0.5$	$b = 2.0$	$b = 0.05$	$b = 0.5$	$b = 2.0$
q	mates	n	x_α=72.27	x_α=7.23	x_α=1.81	x_α=83.05	x_α=8.31	x_α=2.08
0.5	$\hat{x}_\alpha$	30	66.348	6.528	1.644	76.545	7.767	1.966
			(11.655)	(1.106)	(0.269)	(12.907)	(0.889)	(0.236)
		100	69.687	7.035	1.756	81.303	8.155	2.015
			(5.360)	(0.568)	(0.133)	(4.767)	(0.476)	(0.117)
	$\widehat{PTOX}$	30	0.134	0.133	0.131	0.224	0.179	0.210
			(0.034)	(0.032)	(0.033)	(0.032)	(0.043)	(0.042)
		100	0.145	0.152	0.151	0.243	0.242	0.224
			(0.025)	(0.026)	(0.025)	(0.031)	(0.032)	(0.033)
	$\widehat{PROP}$	30	0.255	0.239	0.220	0.282	0.204	0.280
			(0.207)	(0.197)	(0.200)	(0.144)	(0.220)	(0.215)
		100	0.219	0.264	0.250	0.282	0.291	0.207
			(0.246)	(0.254)	(0.248)	(0.258)	(0.262)	(0.248)
	$\widehat{MDIFF}$	30	8.752	0.812	0.209	16.353	0.575	0.202
			(4.504)	(0.411)	(0.108)	(5.282)	(0.346)	(0.099)
		100	3.175	0.374	0.091	3.522	0.366	0.067
			(2.336)	(0.235)	(0.061)	(2.319)	(0.235)	(0.053)
	$\widehat{PDIFF}$	30	0.085	0.078	0.080	0.185	0.065	0.093
			(0.049)	(0.044)	(0.046)	(0.055)	(0.041)	(0.047)
		100	0.028	0.033	0.032	0.039	0.041	0.030
			(0.022)	(0.022)	(0.023)	(0.027)	(0.027)	(0.024)
0.8	$\hat{x}_\alpha$	30	65.240	6.607	1.639	78.982	7.911	1.958
			(9.536)	(0.981)	(0.235)	(9.727)	(0.947)	(0.227)
		100	69.768	7.017	1.744	81.427	8.136	2.037
			(5.735)	(0.559)	(0.132)	(4.658)	(0.468)	(0.118)
	$\widehat{PTOX}$	30	0.118	0.121	0.112	0.205	0.204	0.185
			(0.032)	(0.033)	(0.033)	(0.044)	(0.043)	(0.044)
		100	0.151	0.154	0.151	0.237	0.237	0.236
			(0.025)	(0.025)	(0.025)	(0.032)	(0.034)	(0.032)
	$\widehat{PROP}$	30	0.186	0.216	0.181	0.260	0.253	0.224
			(0.200)	(0.213)	(0.204)	(0.222)	(0.217)	(0.221)
		100	0.248	0.265	0.250	0.272	0.267	0.272
			(0.251)	(0.251)	(0.254)	(0.262)	(0.255)	(0.261)
	$\widehat{MDIFF}$	30	6.604	0.638	0.153	8.851	0.890	0.145
			(3.748)	(0.412)	(0.094)	(4.456)	(0.450)	(0.090)
		100	3.510	0.382	0.087	3.349	0.338	0.083
			(2.295)	(0.251)	(0.058)	(2.241)	(0.232)	(0.056)
	$\widehat{PDIFF}$	30	0.062	0.060	0.057	0.102	0.103	0.066
			(0.039)	(0.043)	(0.039)	(0.053)	(0.053)	(0.043)
		100	0.031	0.034	0.031	0.037	0.038	0.037
			(0.022)	(0.024)	(0.021)	(0.026)	(0.027)	(0.026)

Table 3a. Estimated performance characteristics with the optimal n^* for the probit model with $a = -2$.

			$\alpha = 0.2$			$\alpha = 0.3$		
	Esti-		$b = 0.05$	$b = 0.5$	$b = 2.0$	$b = 0.05$	$b = 0.5$	$b = 2.0$
q	mates	n	x_α=23.17	x_α=2.32	x_α=0.58	x_α=29.51	x_α=2.95	x_α=0.74
0.5	$\hat{x}_\alpha$	30	20.457	1.944	0.484	27.707	2.732	0.697
			(4.753)	(0.421)	(0.100)	(5.443)	(0.567)	(0.136)
		100	21.508	2.162	0.540	28.431	2.742	0.696
			(2.713)	(0.260)	(0.070)	(2.581)	(0.246)	(0.064)
	$\widehat{PTOX}$	30	0.133	0.118	0.113	0.214	0.230	0.216
			(0.036)	(0.037)	(0.035)	(0.047)	(0.040)	(0.047)
		100	0.150	0.150	0.150	0.240	0.214	0.228
			(0.027)	(0.027)	(0.027)	(0.034)	(0.034)	(0.036)
	$\widehat{PROP}$	30	0.223	0.163	0.142	0.266	0.289	0.280
			(0.229)	(0.214)	(0.199)	(0.235)	(0.218)	(0.242)
		100	0.202	0.208	0.211	0.258	0.138	0.197
			(0.260)	(0.264)	(0.267)	(0.278)	(0.235)	(0.267)
	$\widehat{MDIFF}$	30	3.463	0.237	0.056	4.316	0.657	0.114
			(2.053)	(0.192)	(0.041)	(2.482)	(0.318)	(0.061)
		100	1.528	0.152	0.038	1.796	0.113	0.038
			(1.188)	(0.124)	(0.029)	(1.318)	(0.105)	(0.030)
	$\widehat{PDIFF}$	30	0.054	0.037	0.034	0.080	0.123	0.085
			(0.035)	(0.032)	(0.027)	(0.048)	(0.060)	(0.047)
		100	0.023	0.023	0.023	0.033	0.020	0.027
			(0.019)	(0.020)	(0.018)	(0.025)	(0.019)	(0.022)
0.8	$\hat{x}_\alpha$	30	19.765	1.967	0.500	27.040	2.670	0.687
			(4.200)	(0.396)	(0.104)	(4.740)	(0.457)	(0.129)
		100	21.634	2.180	0.545	28.380	2.839	0.716
			(2.650)	(0.276)	(0.065)	(2.744)	(0.259)	(0.066)
	$\widehat{PTOX}$	30	0.120	0.120	0.126	0.198	0.200	0.214
			(.035)	(0.035)	(0.035)	(0.045)	(0.047)	(0.047)
		100	0.152	0.152	0.154	0.235	0.237	0.252
			(0.027)	(0.028)	(0.027)	(0.035)	(0.034)	(0.036)
	$\widehat{PROP}$	30	0.169	0.172	0.203	0.228	0.221	0.279
			(0.217)	(0.221)	(0.232)	(0.242)	(0.235)	(0.242)
		100	0.211	0.234	0.242	0.245	0.252	0.305
			(0.259)	(0.278)	(0.277)	(0.284)	(0.285)	(0.281)
	$\widehat{MDIFF}$	30	2.822	0.277	0.069	3.494	0.293	0.108
			(1.915)	(0.200)	(0.048)	(2.179)	(0.204)	(0.056)
		100	1.566	0.163	0.043	1.606	0.166	0.053
			(1.204)	(0.131)	(0.031)	(1.266)	(0.129)	(0.038)
	$\widehat{PDIFF}$	30	0.044	0.043	0.043	0.065	0.054	0.080
			(0.032)	(0.033)	(0.032)	(0.042)	(0.039)	(0.043)
		100	0.023	0.024	0.026	0.029	0.030	0.039
			(0.019)	(0.021)	(0.020)	(0.024)	(0.024)	(0.028)

Table 3b. Estimated performance characteristics with the optimal n^* for the probit model with $a = -5$.

			$\alpha = 0.2$			$\alpha = 0.3$		
	Esti-		$b = 0.05$	$b = 0.5$	$b = 2.0$	$b = 0.05$	$b = 0.5$	$b = 2.0$
q	mates	n	x_α=83.17	x_α=8.32	x_α=2.08	x_α=89.51	x_α=8.95	x_α=2.24
0.5	$\hat{x}_\alpha$	30	80.092	7.931	2.007	87.425	8.705	2.165
			(5.860)	(0.520)	(0.139)	(5.308)	(0.630)	(0.176)
		100	81.883	8.187	2.042	88.114	8.854	2.219
			(2.579)	(0.294)	(0.069)	(2.532)	(0.252)	(0.065)
	$\widehat{PTOX}$	30	0.136	0.112	0.133	0.212	0.221	0.227
			(0.035)	(0.034)	(0.036)	(0.046)	(0.041)	(0.034)
		100	0.151	0.150	0.151	0.225	0.241	0.242
			(0.026)	(0.027)	(0.027)	(0.034)	(0.033)	(0.035)
	$\widehat{PROP}$	30	0.242	0.164	0.248	0.279	0.282	0.308
			(0.211)	(0.206)	(0.216)	(0.228)	(0.199)	(0.178)
		100	0.242	0.243	0.227	0.212	0.277	0.295
			(0.264)	(0.255)	(0.259)	(0.261)	(0.268)	(0.275)
	$\widehat{MDIFF}$	30	5.042	0.305	0.107	4.908	0.731	0.216
			(2.514)	(0.197)	(0.059)	(2.420)	(0.300)	(0.083)
		100	1.781	0.177	0.043	1.546	0.189	0.048
			(1.281)	(0.125)	(0.033)	(1.229)	(0.128)	(0.035)
	$\widehat{PDIFF}$	30	0.082	0.047	0.069	0.092	0.137	0.160
			(0.044)	(0.033)	(0.041)	(0.046)	(0.056)	(0.058)
		100	0.027	0.027	0.026	0.028	0.034	0.035
			(0.020)	(0.020)	(0.021)	(0.023)	(0.024)	(0.026)
0.8	$\hat{x}_\alpha$	30	79.909	7.957	2.000	87.316	8.712	2.176
			(5.082)	(0.469)	(0.126)	(5.652)	(0.551)	(0.134)
		100	82.057	8.216	2.044	88.588	8.825	2.206
			(2.863)	(0.270)	(0.070)	(2.654)	(0.265)	(0.061)
	$\widehat{PTOX}$	30	0.119	0.110	0.119	0.218	0.217	0.187
			(0.035)	(0.034)	(0.036)	(0.042)	(0.044)	(0.046)
		100	0.153	0.150	0.148	0.241	0.232	0.231
			(0.028)	(0.026)	(0.027)	(0.036)	(0.035)	(0.034)
	$\widehat{PROP}$	30	0.205	0.180	0.208	0.302	0.293	0.212
			(0.225)	(0.217)	(0.220)	(0.218)	(0.217)	(0.227)
		100	0.261	0.251	0.222	0.277	0.234	0.230
			(0.262)	(0.265)	(0.256)	(0.278)	(0.271)	(0.272)
	$\widehat{MDIFF}$	30	3.282	0.261	0.079	5.203	0.516	0.090
			(2.322)	(0.197)	(0.055)	(2.484)	(0.251)	(0.057)
		100	1.925	0.186	0.044	1.955	0.174	0.041
			(1.345)	(0.130)	(0.030)	(1.383)	(0.136)	(0.031)
	$\widehat{PDIFF}$	30	0.052	0.041	0.050	0.097	0.096	0.067
			(0.040)	(0.033)	(0.038)	(0.051)	(0.048)	(0.044)
		100	0.029	0.028	0.026	0.035	0.032	0.030
			(0.021)	(0.021)	(0.019)	(0.026)	(0.025)	(0.023)

Note that the proposed $\hat{x}_\alpha$ under-estimates x_α, a desirable feature while estimating the MTD. The negative bias and the corresponding standard error tend to be lower in magnitude with increasing sample size, as expected. The magnitude of bias and the corresponding standard error seem to be little affected by the choice of q, which is also a desirable feature. Whatever little difference exists for the different choices of q seems to disappear when mean squared error is considered. Generally, about 12 to 25% of the responses are toxic (given by $\widehat{PTOX}$). Again, about 10 to 30% of the selected doses happen to exceed the true MTD (given by $\widehat{PROP}$). The amount of overdose (given by $\widehat{MDIFF}$) generally does not reach beyond 30% of the true MTD, but often is as low as 3-4% above the MTD. This amount of overdose generally decreases with n. The additional risk (given by $\widehat{PDIFF}$) is also generally small, often lying within 5%.

5 Concluding Remarks

The main feature of this paper is the model-free approach for estimation of maximum tolerated dose (MTD) in the context of phase I clinical trials with a continuous domain of dose and variable step size using a generalization of Robbins-Monro (1951) procedure. The step sizes are chosen adaptively by incorporating a larger amount of information (that can be fixed, as given by k in (2), by the user) than only the previous response, as is usually done. The method is quick to produce the next updated dose using as much information as deemed suitable and, in principle, may vary at each step by changing the value of k. A desirable outcome is the apparent under-estimation of MTD. The suggested bootstrap variance estimate also seems to be an under-estimate, but the extent of under-estimation is limited to about 60% and reduces with increasing sample size. This, however, leads to conservative estimation of MTD. Derivation of standard error for the estimate of MTD in this context has generally not been attempted, which is another added feature of our method.

In order to apply the method, as described in Section 2, one needs to have some information on x_1, the starting dose, x^*, a highly toxic dose and n^*, to be chosen optimally in some sense. Intuitively, choice of x_1 may not have much impact if the sample size n is sufficiently large. In practice, we have little control over the choice except insisting that it be close to x_α, or choose $x_1 = 0$ conservatively. Choice of x^* and optimal n^* are related to each other. We have expressed x^* as the qth percentile of the dose response curve and studied the impact of this choice through q in Section 4. As remarked in the previous section, the estimate of MTD and the corresponding standard error seem to be quite robust against the two choices of q (see Tables $2a, b$ and $3a, b$). The choice of

optimal n^* seems to be very important. Table 1 gives a guideline for this choice. It is seen there that this choice depends weakly on q and the slope of the dose response curve; however, there seems to be some dependence on the intercept. It is likely to have some information on the background response giving some idea about the choice of n^* from Table 1; this information also helps in the choice of x_1.

Note that, once the optimal n^* is chosen resulting in the value of C, the adaptive procedure of updating the next dose level proceeds without any reference to the sample size n. Therefore, the method does not have to stop at nth participant and, as a matter of fact, the updating can be continued till a suitable stopping rule. This stopping rule may be in terms achieving some pre-assigned bias and/or precision which are seen to improve with sample size. In application, after few responses are obtained, one can simulate the trial with an assumed dose response curve and make an assessment of the required sample size for given bias and precision.

The present approach can be implemented adaptively to a more general set up where, after a few patients are observed, the experimenter decides to change the value of α. The already available data can be used to estimate the dose response curve, which can be used to choose the new values of x_1 and x^* (and, hence, C) corresponding to the new α to suggest a procedure similar to (1) for the subsequent patients. The details are under study.

References

1. Babb, J., Rogatko, A. and Zacks, S. (1998). Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine* **17**, 1103-1120.

2. Bortot, P. and Giovagnoli, A. (2005). Up-and-down designs of the first and second order. *Journal of Statistical Planning and Inference* **134**, 236-253.

3. Bueon C., Bjerrum, O. J. and Thomsen, M. S. (2005). How first-time-in-human studies are being performed: A survey of phase I dose-escalation trials in healthy volunteers published between 1995 and 2004. *Journal of Clinical Pharmacology* **45**, 1123-1136.

4. Cheung, Y. K. and Chappell, R. (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56**, 1177-1182.

5. Dixon, W. J. and Mood, A. M. (1954). A method for obtaining and analyzing sensitivity data. *Journal of American Statistical Association* **43**, 109-126.

6. Durham, S. D. and Flournoy, N. (1995). Up-and-down designs I: Stationary treatment distributions. In *Adaptive Designs. Lecture-Notes Monograph Series*, N. Flournoy and W. F. Rosenberger (Eds.), Institute of Mathematical Statistics, Hayward, CA, Vol. 25, 139-157.

7. Durham, S. D., Flournoy, N. and Haghighi, A. A. (1995). Up-and-down designs II: Exact treatment moments. In *Adaptive Designs. Lecture-Notes Monograph Series*, N. Flournoy and W. F. Rosenberger (Eds.), Institute of Mathematical Statistics, Hayward, CA, Vol. 25, 158-178.

8. Faries, D. (1994). Practical modifications of the continual reassessment method for phase I cancer clinical trials. *Journal of Biopharmaceutical Statistica* **4**, 147-164.

9. Gasparini and Eisele (2000). A curve-free method for phase I clinical trials. ıBiometrics **56**, 609-615.

10. Giovagnoli, A. and Pintacuda, N. (1998). Properties of frequency distributions induced by general up-and-down methods for estimating quantiles. *Journal of Statistical Planning and Inference* **74**, 51-63.

11. Ivanova, A., Flournoy, N. and Chung, Y. (2007). Cumulative cohort design for dose-finding, *Journal of Statistical Planning and Inference* **137**, 2316-2327.

12. Joseph V. R. (2004). Efficient Robbins-Monro procedure for binary data. *Biometrika* **91**, 461-470.

13. Li, Z., Durham, S.D. and Flournoy, N. (1995). An adaptive design for maximization of a contingent binary response. In *Adaptive Designs, IMS Lecture Notes Monograph Series, Volume 25.* in Flournoy, N. and Rosenberger W. F. (eds.), Adaptive Designs, Hayward , CA : Institute of Mathematical Statistics, pp. 179-196.

14. Liu, G., Rosenberger, W.F. and Haines, L.M. (2006). Sequential designs for logistic phase I clinical trials. *Journal of Biopharmaceutical Statistics* **16**, 605-621.

15. Leung, D. H. and Wang, Y. G. (2001). Isotonic designs for phase I trials. *Controlled Clinical Trials* **22**, 126-138.

16. Leung, D. H. and Wang, Y. G. (2002). An extension of the continual reassessment method using decision theory. *Statistics in Medicine* **21**, 51-63.

17. Mats, V. A., Rosenberger, W. F., and Flournoy, N. (1998). Restricted optimality for phase I clinical trials. In Flournoy, N., Rosenberger, W. F., and Wong, W.

K. (eds.), New Developments and Applications in Experimental Design, Hayward, CA: Institute of Mathematical Statistics, pp. 50-61.

18. McLeish, D. L. and Tosh, D. (1990). Sequential designs in bioassay. *Biometrics* **46**, 103-116.

19. O'Quigley, J. (1992). Estimating the probability of toxicity at recommended dose following a phase I clinical trial in cancer. *Biometrics* **48**, 853-862.

20. O'Quigley, J., Pepe, M. and Fisher, L. (1990). Continual reassessment method: A practical design for phase I clinical trial in cancer. *Biometrics* **46**, 33-48.

21. O'Quigley, J. and Chevret, S. (1991). Methods for dose-finding studies in cancer clinical trials: A review and results of a Monte Carlo study. *Statistics in Medicine* **10**, 1647-1664.

22. Robbins, H. and Monro, S. (1951). A stochastic approximation method. *Annals of Mathematical Statistics* **29**, 373-405.

23. Rosenberger, W. F., Haines, L. M., Perevozskaya, I. (2001). Constrained Bayesian optimal designs for phase I clinical trials: continuous dose space. In Atkinson, A. C., Hackl, P., and Mueller, W. G. (eds.), mODa6 Advances in Model Oriented Design and Data Analysis, Heidelberg: Physica-Verlag, pp. 209-217.

24. Shen, L. Z. and O'Quigley, J. (1996). Consistency of continual reassessment method under model misspecification. *Biometrika* **83**, 395-405.

25. Storer, B. E. (1989). Design and analysis of phase I clinical trials. *Biometrics* **45**, 925-937.

26. Storer, B. E. (1993). Small-sample confidence sets for the MTD in a phase I clinical trial. *Biometrics* **49**, 1117-1125.

27. Storer, B. E. (2001). An evaluation of phase I clinical trial designs in the continuous dose-response setting. *Statistics in Medicine* **20**, 2399-2408.

28. Whitehead, J. and Brunier, H. (1995). Bayesian decision procedures for dose determining experiments. *Statistics in Medicine* **14**, 885-893.

29. Whitehead, J. and Williamson, D. (1998). Bayesian decision procedures based on logistic regression models for dose-finding studies. *Journal of Biopharmaceutical Statistica* **8**, 445-467.