Stochastic approximation with virtual observations for dose-finding on discrete levels

BY YING KUEN CHEUNG

*Department of Biostatistics, Columbia University, 722 West 168th Street, New York,
New York 10032, U.S.A.
yc632@columbia.edu*

AND MITCHELL S. V. ELKIND

*Department of Neurology, Columbia University, 710 West 168th Street, New York,
New York 10032, U.S.A.
mse13@columbia.edu*

SUMMARY

Phase I clinical studies are experiments in which a new drug is administered to humans to determine the maximum dose that causes toxicity with a target probability. Phase I dose-finding is often formulated as a quantile estimation problem. For studies with a biological endpoint, it is common to define toxicity by dichotomizing the continuous biomarker expression. In this article, we propose a novel variant of the Robbins–Monro stochastic approximation that utilizes the continuous measurements for quantile estimation. The Robbins–Monro method has seldom seen clinical applications, because it does not perform well for quantile estimation with binary data and it works with a continuum of doses that are generally not available in practice. To address these issues, we formulate the dose-finding problem as root-finding for the mean of a continuous variable, for which the stochastic approximation procedure is efficient. To accommodate the use of discrete doses, we introduce the idea of virtual observation that is defined on a continuous dosage range. Our proposed method inherits the convergence properties of the stochastic approximation algorithm and its computational simplicity. Simulations based on real trial data show that our proposed method improves accuracy compared with the continual re-assessment method and produces results robust to model misspecification.

Some key words: Continual re-assessment method; Dichotomized data; Discrete barrier; Heteroscedasticity; Robust estimation; Semiparametric mean-variance relationship.

1. INTRODUCTION

The early phase clinical development of a new drug typically involves testing several doses of the new drug, with safety as the primary concern. A specific objective is to identify the maximum tolerated dose, defined as the largest dose that causes toxicity with a prespecified probability. Traditionally, the maximum tolerated dose is approached from a low dose according to a 3 + 3 algorithm whereby escalation occurs after cohorts of three nontoxic observations. As this traditional algorithm is criticized for its arbitrary statistical properties, several authors propose novel designs, such as the continual re-assessment method (O’Quigley et al., 1990) and rigorous extension of the 3 + 3 algorithm (Lin & Shih, 2001; Cheung, 2007) to address

Table 1. *Summary of the liver function data in NeuSTART*

Dose k	Cohort size	Patients with toxicity	log(ALT)		Isotonic		Toxicity rate	
			Mean	SD	Estimate	SD	CRM	Isotonic
1	3	0	3.24	0.23	3.24	0.23	0.01	0.00
2	10	0	3.25	0.42	3.25	0.42	0.03	0.00
3	12	2	3.78	0.72	3.63	0.66	0.06	0.04
4	8	0	3.42	0.54	3.63	0.66	0.13	0.04
5	0	—	—	—	—	—	0.24	—

ALT, peak alanine aminotransferase; CRM, continual re-assessment method; SD, standard deviation.

dose-finding as a quantile estimation problem. These designs, utilizing binary toxicity data for estimation, provide clinicians with viable options for testing drugs in general.

In clinical studies with a biological safety endpoint, the toxicity outcome is often defined by dichotomizing a quantitative biomarker expression. Take for example the Neuroprotection with Statin Therapy for Acute Recovery Trial, NeuSTART, a recent phase I trial of lovastatin in acute ischemic stroke patients conducted at Columbia University (Elkind et al., 2008). In this trial, high dose lovastatin was administered to patients for three days followed by a standard dose for 30 days. A primary safety concern for administering high dose lovastatin is toxicity measured by the elevation in liver enzyme alanine aminotransferase. In NeuSTART, each patient was evaluated up to day 30 after the start of treatment, and toxicity was said to occur if the patient's post-treatment peak transferase exceeded three times the upper normal limit of 123 U/L. The objective of the trial was to estimate the dose with toxicity probability closest to 10%. A two-stage continual re-assessment method (Cheung, 2005) was used for dose escalation in 33 subjects at five high-dose levels. Table 1 summarizes the liver function outcomes by dose group in the study: two patients at dose level 3 exhibited elevated liver function, and the trial did not reach the highest dose. According to the continual re-assessment method, dose level 4 was estimated to be the maximum tolerated dose with a 13% toxicity rate. However, an isotonic fit to the continuous transferase measurements suggested a much smaller rate of 4% for dose level 4. Details of the continual re-assessment method and the isotonic fit are described in § 5.

In retrospect, we realized that statistical efficiency might have been lost by using only the dichotomized data, and that the original measurements might be utilized to retain information. In this article, we propose a variant of the Robbins–Monro (1951) stochastic approximation that uses the quantitative measurements for dose-finding in phase I trials. The Robbins–Monro procedure is a stochastic root-finding method, and is a natural method of choice for the quantile estimation objective in dose-finding (Anbar, 1977, 1984). While their paper has generated a voluminous statistical and engineering literature, the stochastic approximation has seldom been considered for use in clinical trials; exceptions are Anbar (1984) and O'Quigley & Chevret (1991). Two considerations render application of the Robbins–Monro procedure unsuited to many phase I trial settings. First, from a statistical viewpoint, the method has been demonstrated to be inferior to parametric procedures such as the maximum likelihood recursion for binary data (Wu, 1985). In our current application where the binary outcome is defined by a continuous measurement, this difficulty can be resolved if we make appropriate use of the original continuous data. Second, the method entails the availability of a continuum of doses. This is neither feasible in practice nor preferred for reporting purposes in publications. In addition, in trials involving combination of treatments, each subsequently higher regimen may involve incrementing doses of different treatments and hence there is no natural scale of dosage. This important point has been noted in O'Quigley & Chevret (1991) and Shen & O'Quigley (2000), both of whom address the problem

by rounding the stochastic approximation output to its closest dose at each step. We shall return to this when we introduce a novel concept called virtual observation.

2. STOCHASTIC APPROXIMATION WITH VIRTUAL OBSERVATIONS

2.1. Problem formulation

Consider a trial in which patients are enrolled in small cohorts of size m . Let X_i^* and X_i respectively denote the dose assigned to the i th cohort and the actual dose given to patients. The assigned dose X_i^* can take on any real number from a conceptual scale that represents an ordering of doses or regimens. Suppose for the moment that X_i can take on any real number and that $X_i = X_i^*$ for all i . Let Y_{ij} be the safety measurement of the j th patient in the i th cohort; the patient is said to experience toxicity if Y_{ij} exceeds some threshold t_0 . The objective of a dose-finding study is the identification of θ such that $\pi(\theta) = p$, where $\pi(x) = P(Y_{ij} > t_0 \mid X_i = x)$ is the probability of toxicity at dose x .

Assume that $Y_{ij} = M(X_i) + \sigma(X_i)\epsilon_{ij}$, where the ϵ s are independently distributed according to a common distribution G with $E(\epsilon) = 0$ and $\text{var}(\epsilon) = 1$, so that $M(x)$ and $\sigma(x)$, respectively, denote the mean and the standard deviation of the safety outcome at dose x . Both $M(x)$ and $\sigma(x)$ are unspecified. Under this regression model, the toxicity probability at dose x ,

$$\pi(x) = 1 - G\left\{\frac{t_0 - M(x)}{\sigma(x)}\right\}, \quad (1)$$

and the target dose θ can be re-expressed as the solution to the equation $f(x) \equiv M(x) + z_p\sigma(x) = t_0$, where z_p is the upper p th percentile of G . For brevity, suppose that $\pi(x)$ and $f(x)$ are strictly increasing on the relevant dose range so that the root θ exists uniquely. This monotone dose-toxicity assumption is reasonable in phase I trials, although the main results in this article hold under the much weaker assumptions stated in the Appendix. In practice, the actual dose given is confined to a discrete set of K levels, denoted by $\{1, \dots, K\}$, with $\pi(1) < \dots < \pi(K)$ and $f(1) < \dots < f(K)$. As such, it is possible that $\pi(k) \neq p$ and $f(k) \neq t_0$ for all k . One may then define the maximum tolerated dose by $v_1 = \arg \min_k |\pi(k) - p|$ as in O'Quigley et al. (1990). Alternatively, under the current formulation, it is natural to define the maximum tolerated dose by $v_2 = \arg \min_k |f(k) - t_0|$. The objectives v_1 and v_2 represent the closest doses to θ on two different scales, and are not necessarily identical.

LEMMA 1. Define c_k such that $G(c_k) = 1 - \pi(k)$ for $k = 1, \dots, K$. (a) Assume $\pi(v_1) \geq p$. If $\pi(k) < 1 - G\{z_p + \sigma(v_1)(z_p - c_{v_1})/\sigma(k)\}$ for all $k < v_1$, then $v_2 = v_1$; else, $v_2 = v_1 - 1$.

(b) Assume $\pi(v_1) \leq p$. If $\pi(k) > 1 - G\{z_p + \sigma(v_1)(z_p - c_{v_1})/\sigma(k)\}$ for all $k > v_1$, then $v_2 = v_1$; else, $v_2 = v_1 + 1$.

Therefore, $v_2 = v_1$ if $\pi(v_1) = p$. Lemma 1 also implies that $v_2 = v_1$ when the dose-toxicity curve is steep around v_1 . In other words, when $v_2 \neq v_1$, $\pi(v_1)$ and $\pi(v_2)$ will be quite close to each other, thus rendering both objectives similar. For clarity in presentation, we will focus on the estimation of v_1 in accordance with the conventional phase I objective.

2.2. The procedure

Let $U_i = \bar{Y}_i + [E\{S_i/\sigma(X_i)\}]^{-1}z_p S_i$, where $\bar{Y}_i$ and S_i are the respective sample mean and standard deviation of the measurements in the m subjects of cohort i , so that $E(U_i \mid X_i) = f(X_i)$. Since $S_i/\sigma(X_i)$ is a pivotal quantity that depends on the ϵ s but not M and σ , we can derive $E\{S_i/\sigma(X_i)\}$ for any given distribution G either by analytical calculation or by simulation. Define, for the i th cohort, a virtual observation $V_i = U_i + \beta(X_i^* - X_i)$ made at the assigned dose

X_i^* for some $\beta > 0$. Then the next assigned dose X_{i+1}^* is calculated via a stochastic approximation recursion based on the virtual observation:

$$X_{i+1}^* = X_i^* - \frac{1}{i\beta}(V_i - t_0) \quad (2)$$

and the next actual dose $X_{i+1} = C(X_{i+1}^*)$, where $C(x) = k$ if $l_k \leq x < l_{k+1}$ with $l_1 = -\infty$, $l_{K+1} = \infty$, and $l_k = k - 0.5$ for $k = 2, \dots, K$. That is, $C(x)$ is the rounded value of x if $0.5 \leq x < K + 0.5$.

There are various ways to initiate the recursion. One could start the first cohort at the lowest dose, i.e. $X_i^* = X_i = 1$, or at the prior maximum tolerated dose as in O'Quigley et al. (1990). One may also institute a predetermined escalation sequence until the first sign of toxicity appears. Specifically, let $X_i^* = X_i = x_{0,i}$ for the i th cohort, where $x_{0,i}$ is a nondecreasing sequence of dose levels with $x_{0,i} \leq x_{0,i+1}$, and switch to recursion (2) once a toxicity is observed. While a two-stage strategy is not necessary for the method's implementation, it may be practical to dictate the pace of escalation in accordance with the clinician's preference when there is no toxicity (Cheung, 2005).

Another way to avoid aggressive escalation is by restricting the trial from skipping an untested dose. That is, the next assigned dose would be $\min(X_{i+1}^*, \max_{1 \leq n \leq i} X_n + 2)$ instead of X_{i+1}^* . If this restriction is to be applied, then it should be only to the first few cohorts because the increment diminishes as i increases.

2.3. Discrete barrier: an illustration

As an alternative to the virtual observation recursion (2), we may generate a discretized stochastic approximation of the given dose X_i based on the actual observations U_i by taking the closest dose to the value obtained from a recursion output; that is,

$$X_{i+1} = C \left\{ X_i - \frac{1}{i\beta}(U_i - t_0) \right\}. \quad (3)$$

O'Quigley & Chevret (1991) and Shen & O'Quigley (2000) take a similar discretization approach to accommodate the use of discrete doses in the context of different dose-finding designs. The discretized stochastic approximation (3) is straightforward and at first looks reasonable, but may indefinitely confine the design sequence to a wrong dose. Consider a two-stage design for five dose levels with $t_0 = 4.81$. In the event of no toxicity, i.e. $Y_{ij} \leq 4.81$, escalation will proceed according to an initial sequence: $x_{0,1} = 1$, $x_{0,2} = 2$, $x_{0,3} = x_{0,4} = 3$, $x_{0,5} = x_{0,6} = x_{0,7} = 4$ and $x_{0,i} = 5$ for $i \geq 8$. Once the first toxicity is seen, recursion (3) with $\beta = 0.05$ will be used to assign doses. Take a simple instance with negligible variability so that $U_i = f(X_i)$ with $f(1) = 4.20$, $f(2) = 4.67$, $f(3) = 4.80$, $f(4) = 4.93$ and $f(5) = 5.30$; hence, the correct dose is level 3. The outcome sequence will follow the initial design and yield $(X_1, U_1) = (1, 4.20)$, $(X_2, U_2) = (2, 4.67)$, $(X_3, U_3) = (X_4, U_4) = (3, 4.80)$ and $(X_5, U_5) = (4, 4.93)$, after which recursion (3) will come into effect and give $X_6 = C\{4 - (5 \times 0.05)^{-1}(4.93 - 4.81)\} = 4$, and $U_6 = 4.93$. It is easy to see that the remaining patients will receive dose 4. If, instead, we use the virtual observation recursion (2) after the initial design, then $X_6^* = 3.52$ which gives $X_6 = 4$, $U_6 = 4.93$ and $V_6 = 4.906$, followed by $X_7^* = 3.14$ and $X_7 = 3$. The remaining process will continue at the correct dose 3.

To see how the virtual observation recursion corrects the problem, consider another instance with $f(x) = t_0 + \beta(x - \theta)$, where θ is an integer, and U_i is observed with noise Z_i , i.e. $U_i = f(X_i) + Z_i$. If the i th dose X_i is an overdose dose, namely, $X_i = \theta + \Delta_i$ for some $\Delta_i > 0$, then the update according to (3) is $X_{i+1} = X_i + [C\{(1 - i^{-1})\Delta_i - (i\beta)^{-1}Z_i\} - \Delta_i]$. As i grows,

the second term is likely to equal zero and does not contribute to future updates. A discrete barrier is thus built by rounding. The virtual observation recursion (2) updates a dose with $X_{i+1}^* = X_i^* - i^{-1}(\Delta_i^* + Z_i/\beta)$, where $\Delta_i^* = X_i^* - \theta$. The second term, which is of the order $O(i^{-1})$, can be carried over to future updates and provides continuation to overcome the discrete barrier.

3. DESIGN PROPERTIES

The idea of virtual observation is to create an objective function $h(x) = E(V_i \mid X_i^* = x)$ with a local slope at the available doses $\{1, \dots, K\}$. To be precise, $h(x) = f\{C(x)\} + \beta\{x - C(x)\}$ is piecewise continuous with jumps at $\{0.5, 1.5, \dots, K - 0.5\}$ and is linearly increasing with slope β except at the jumps. Following from the standard results of stochastic approximation (Sacks, 1958; Lai & Robbins, 1979), under the condition that $h(x) = t_0$ has a unique root at θ_β , the virtual dose X_n^* generated by (2) is consistent for θ_β , and X_n for $C(\theta_\beta)$. Furthermore, since the objective $h(x)$ is equal to the target $f(x)$ on the actual doses, i.e. $h(k) = f(k)$, we hope that the root θ_β will be close to θ , and $C(\theta_\beta)$ to v_1 .

PROPOSITION 1. (a) Assume $v_1 = v_2$. If $B_{v_1} < \beta < \min_{k \neq v_1} B_k$, then $X_n = v_1$ eventually with probability one, where $B_k = 2\sigma(k)|c_k - z_p|$.

(b) Assume $v_1 \neq v_2$. If $\beta < \min_{k \neq v_1} B_k$, then $\text{pr}(X_n = v_1 \text{ or } v_2 \text{ eventually}) = 1$ and

$$|\pi(X_n) - p| \leq \left| G \left\{ z_p + \frac{\sigma(v_1)}{\sigma(v_2)}(z_p - c_{v_1}) \right\} - G(c_{v_1}) \right| - |\pi(v_1) - p|$$

eventually with probability one.

Supposing $c_{v_1} \approx z_p$, i.e. $B_{v_1} \approx 0$, and the parameter β is adequately small, Proposition 1 implies that X_n will eventually recommend a neighbouring dose of v_1 , if not v_1 itself. The specific choice of β can be aided by the knowledge about π and σ . While these quantities are often unknown prior to the trial, we illustrate in § 4 how β may be determined with only mild prior inputs under a semiparametric mean-variance assumption.

Proposition 1(b) prescribes how far X_n may deviate from θ on the probability scale if β is adequately small. For example, if $p = 0.1$ and $\pi(v_1) = 0.08$, applying Proposition 1(b) gives $|\pi(X_n) - 0.1| \leq 0.024$ eventually by assuming $\sigma(v_1) = \sigma(v_2)$. The maximum deviation is only slightly larger than $|\pi(v_1) - p| = 0.02$. The deviation will be bounded above by 0.037 if we conservatively assume $\sigma(v_1) = 1.5\sigma(v_2)$, while we would expect $\sigma(v_1)/\sigma(v_2) \approx 1$ as v_1 and v_2 are neighbours.

4. DESIGN CALIBRATION

4.1. Consistency

The virtual observation recursion is specified by a tuning parameter β . Assuming that the condition for consistency in Proposition 1 holds, we can further show that X_n^* is asymptotically normal with mean θ_β and variance $\beta^{-2}\text{var}(V_i \mid X_i^* = \theta_\beta)$. To minimize the asymptotic variance, therefore, we should set β at its largest possible value that guarantees consistency. However, since the upper bound of β depends on the unknown $\pi(k)$ and $\sigma(k)$, to simplify the calibration process, we restrict our attention to scenarios in which

$$\pi(v_1) = \pi(v_2) = p, \quad \pi(k) \leq p_L < p < p_U \leq \pi(k'), \quad (4)$$

for $k < v_1 < k'$ and some prespecified limits p_L, p_U . This class of scenarios consists of toxicity configurations in which there is an unambiguous target with $\theta = v_1 = v_2 = \theta_\beta$, and the toxicity probabilities at the other doses lie outside the interval (p_L, p_U) . By not considering scenarios with $\pi(k) \in (p_L, p_U)$, we are in a sense indifferent to whether such a dose should be selected as the maximum tolerated dose. Thus, the interval (p_L, p_U) is called the indifference interval (Cheung & Chappell, 2002).

PROPOSITION 2. Assume that $\sigma(k)$ depends on dose k only via $M(k)$ as follows: $s\{\sigma(k)\} = s\{\sigma(\theta)\} + \phi\{M(k) - M(\theta)\}$ for some $\phi \geq 0$ and some smooth function $s(\sigma)$ with $s(\sigma) \geq 0$ and $s'(\sigma) > 0$. Under (4), X_n is consistent for θ if $0 < \beta/\{2\sigma(\theta)\} < w \equiv \min(z_p - z_p^2/z_L, z_p - z_U)$, where $z_L = G^{-1}(1 - p_L)$ and $z_U = G^{-1}(1 - p_U)$.

The consistency upper bound for β depends on the unknown $\sigma(\theta)$. On the other hand, we can often obtain clinical data to estimate the standard deviation σ_0 of the safety measurements in a comparable population that is untreated or treated at lower doses. Since $\sigma_0 \leq \sigma(\theta)$ by the assumption that variance increases with doses, we may use $2w\sigma_0$ as a conservative consistency upper bound for β .

4.2. Robustness

Our proposed method works under a fairly general dose-toxicity relationship as described in Proposition 2, but requires us to specify a distribution G . To provide some assurance in an application, it is important to evaluate the method's robustness against misspecification. Suppose we assume ϵ_{ij} follows $\tilde{G}$ instead of the true G , and generate recursion (2) based on $\tilde{V}_i = \tilde{Y}_i + [E_{\tilde{G}}\{S_i/\sigma(X_i)\}]^{-1}\tilde{z}_p S_i + \beta\{X_i^* - C(X_i^*)\}$ where $E_{\tilde{G}}(\cdot)$ denotes expectation computed under $\tilde{G}$ and $\tilde{G}(\tilde{z}_p) = 1 - p$. Furthermore, let

$$\tilde{z}_p^* = \frac{E_G\{S_i/\sigma(X_i)\}}{E_{\tilde{G}}\{S_i/\sigma(X_i)\}}\tilde{z}_p$$

and $\tilde{h}(x) \equiv E_G(\tilde{V}_i | X_i^* = x) = M\{C(x)\} + \tilde{z}_p^* \sigma\{C(x)\} + \beta\{x - C(x)\}$. If β is chosen such that $\tilde{h}(x) = t_0$ has a unique solution at $\tilde{\theta}_\beta$, then it can be shown that the virtual dose X_n^* generated based on $\tilde{V}_i$ s will converge to $\tilde{\theta}_\beta$ and X_n to $C(\tilde{\theta}_\beta)$.

PROPOSITION 3. Define $\tilde{v} = \arg \min_k \tilde{B}_k$, where $\tilde{B}_k = 2\sigma(k)|c_k - \tilde{z}_p^*|$. (a) Under (4), $\tilde{v} = \theta$ if $p_L < 1 - G\{\tilde{z}_p^* + \sigma(\theta)|z_p - \tilde{z}_p^*|/\sigma(1)\} \equiv \tilde{p}_L$ and $p_U > 1 - G\{\tilde{z}_p^* - |z_p - \tilde{z}_p^*|\} \equiv \tilde{p}_U$.

(b) Assume that the mean-variance relationship in Proposition 2 holds, and that $p_L < \tilde{p}_L$ and $p_U > \tilde{p}_U$ under (4). Then, X_n is consistent for $C(\tilde{\theta}_\beta) = \tilde{v} = \theta$ if $|z_p - \tilde{z}_p^*| < \beta/\{2\sigma(\theta)\} < \tilde{w} \equiv \min(z_p - z_p\tilde{z}_p^*/z_L, \tilde{z}_p^* - z_U)$.

Proposition 2 is a special case of Proposition 3 when $\tilde{G} = G$. Generally, Proposition 3 states that the false objective $\tilde{\theta}_\beta$ will coincide with the true θ under a steep dose-toxicity curve, where steepness is quantified by $\tilde{p}_L$ and $\tilde{p}_U$. For example, Table 2 displays the key robustness quantities for recursion (2) with standard normal $\tilde{G}$, $m = 3$ and $p = 0.1$ evaluated under the logistic, t_5 , and Gumbel distributions standardized so that they have mean zero and variance one; the logistic distribution has location 0 and scale 0.55; the noise under t_5 is generated as a t -variate with 5 degrees of freedom multiplied by 0.77; the distribution function for the Gumbel distribution is $G(z) = \exp[-\exp\{(-0.45 - z)/0.78\}]$. To evaluate $\tilde{p}_L$ we set $\sigma(1) = 0.25\sigma(\theta)$. Table 2 shows that $\tilde{p}_L$ and $\tilde{p}_U$ are within 0.05 of the target p , suggesting mild conditions on the dose-toxicity curve for robust estimation of θ . In addition, $|z_p - \tilde{z}_p^*|$ is generally small, and is much smaller

Table 2. *Robustness analysis of the virtual observation recursion with standard normal $\tilde{G}$, $m = 3$ and $p = 0.1$*

G	z_p	$E_G\{S_i/\sigma(X_i)\}$	$\tilde{z}_p^*$	$\tilde{p}_L$	$\tilde{p}_U$	$ z_p - \tilde{z}_p^* $	$\tilde{w}$
$N(0, 1)$	1.282	0.8862	1.282	0.100	0.100	0.000	0.245
Logistic	1.211	0.8663	1.253	0.071	0.100	0.041	0.277
t_5	1.143	0.8438	1.220	0.053	0.100	0.077	0.250
Gumbel	1.305	0.8514	1.231	0.076	0.119	0.073	0.265

than $\tilde{w}$, when computed with $p_L = 0.05$, $p_U = 0.15$. We thus anticipate that a wide range of β will satisfy the consistency conditions under misspecification of the model.

5. APPLICATION TO NEUSTART

5.1. Trial design and data

In NeuSTART, we adopted a two-stage continual re-assessment method for dose escalation in 33 subjects at five dose levels using the initial sequence given in § 2.3. Once the first toxicity was seen, the dose-toxicity curve would be updated after an observation was made from the most recent patient, and the next patient would be given the dose with toxicity probability estimated to be closest to $p = 0.10$. The toxicity probability for dose k was modelled as d_k^ψ with $d_1 = 0.02$, $d_2 = 0.06$, $d_3 = 0.10$, $d_4 = 0.18$ and $d_5 = 0.30$, where ψ was a priori lognormal with location 0 and scale 1.34 (O'Quigley & Shen, 1996). Enrolment into the trial began in October 2005 and was completed in August 2007. The first sign of elevated liver function test was seen in the eighth subject who was given dose level 3. The subsequent dose assignments were then model-based.

Table 1 reports the sample mean and the standard deviation of peak transferase levels (on log scale) for each dose. Assuming monotone dose-toxicity, we estimate $M(k)$ and $\sigma(k)$ by pooling data in dose levels 3 and 4, because the sample mean of dose 3 is larger than that of dose 4. Substituting these estimates for the corresponding parameter values in (1) gives an isotonic fit of $\pi(k)$, which is shown in the last column of Table 1 and plotted in Fig. 1(a). The estimate of $\pi(4)$ based on the isotonic fit is much lower than the target 10%. In contrast, the continual re-assessment method estimates that $\pi(4)$ exceeds the target, even though no toxicity was observed at dose level 4. This is indeed a feature of the continual re-assessment method that allows strong influence of observations at lower doses on estimation at the higher doses via parametric extrapolation.

Figure 1(b) displays the liver function data by dose, and reveals that the two toxic outcomes were results of peak transferase levels exceeding the threshold $t_0 = \log(123)$ by a slight margin. Were these measurements to drop a few units, the trial outcomes would have been quite different. Such sensitivity in dichotomized analysis speaks favourably for using the original continuous measurements as the basis of estimation, especially when the data are subject to measurement errors.

The liver function data demonstrate heteroscedasticity in Table 1 and Fig. 1(b). Specifically, the variance increases with dose as the mean increases. Such a monotone mean-variance relationship is typical for biological laboratory values, although the exact form of the relationship is usually hard to estimate. Finally, to identify the noise distribution, we calculate the standardized residual for each observation via centring the transferase measurement by the isotonic estimate of its mean and standardization by its standard deviation estimate. A normal Q–Q-plot of these standardized residuals shows that the noise ϵ_{ij} fits well to the standard normal distribution.

Fig. 1. Model fits for the NeuSTART data. (a) Fitted dose-toxicity curve by the continual re-assessment method (dashed) and the isotonic fit (solid); the target probability is also indicated (thicker solid line). (b) Liver function outcomes. Each observation is indicated (o), as well as the isotonic fit (solid) and the toxicity threshold (dashed).

5.2. Redesign of NeuSTART

Having checked all the model components, we may consider the virtual observation recursion assuming standard normal noise, so that $E\{S_i/\sigma(X_i)\} = 0.886$ and $z_p = 1.28$ for a trial with $m = 3$ and $p = 0.1$. To determine β , we approximate σ_0 by $\hat{\sigma}_{0,\text{low}} = 0.59$, which is the 80% confidence lower limit based on the 20 observations in the combined dose levels 3 and 4. Then the upper bound for β is 0.29 according to Proposition 2 if we set $p_L = 0.05$ and $p_U = 0.15$, and it is less than $2\tilde{w}\hat{\sigma}_{0,\text{low}}$ under various noise distributions, see Table 2. This estimate is conservative as far as consistency and robustness are concerned, because $\hat{\sigma}_{0,\text{low}}$ tends to underestimate σ_0 and hence $\sigma(\theta)$. Also, $0.29/\{2\sigma(\theta)\} > |z_p - \tilde{z}_p^*|$ under the distributions in Table 2, unless the true $\sigma(\theta) \geq 3.77$, which is very unlikely. Therefore, recursion (2) will be generated based on $V_i = \bar{Y}_i + 1.446S_i + 0.29(X_i^* - X_i)$.

5.3. Simulation study

Simulations were run to compare the performance of the virtual observation recursion and the continual re-assessment method. In the simulation, we used the same two-stage strategies for both methods: start at the lowest dose, follow the initial sequence as in § 2.3 and switch to the continual re-assessment method or recursion (2) upon the first toxic outcome. Furthermore, we applied the restriction that no untested level would be skipped in escalation. For the purposes of theoretical comparisons, we also ran simulations using the nonparametric optimal design described in O'Quigley et al. (2002). The design assumed that the binary toxicity outcomes at every test dose were observable for each given patient and used the complete toxicity profile to evaluate the sample toxicity proportion at each dose. It can be readily shown that this design is optimal in that the sample proportion for $\pi(k)$ achieves the Cramér–Rao lower bound with variance proportional to $G(c_k)\{1 - G(c_k)\}$. While the nonparametric optimal design cannot be implemented in practice where each patient is treated at a single dose, the design's performance can be evaluated in simulations and can be used as a benchmark for efficiency.

In each simulated trial, a true dose-toxicity curve $\pi(k)$ was first generated from the posterior distribution under the assumed model in the continual re-assessment method, i.e. $\pi(k) = d_k^\psi$ with ψ drawn from the posterior given the NeuSTART data. Continuous outcomes were generated from a normal distribution with mean and standard deviation determined so that, for a given $\pi(k)$, $\sigma^4(k) = \log\{\alpha + 0.52M(k)\}$ for some $\alpha \in [-1, 1]$. With $p = 0.1$ and $t_0 = \log(123)$, the

Fig. 2. (a) and (c) The proportion of selecting v_1 vs. $\sigma(\theta)$ by the virtual observation recursion assuming normal noise. (b) and (d) The proportion of selecting an acceptable dose by our method. Our method was run with $\beta = 0.15$ (dot-dashed), 0.29 (solid), 0.45 (dotted) under normal noise in (a) and (b), and with $\beta = 0.29$ and noise generated from the logistic (solid), t_5 (dotted) and Gumbel (dot-dashed) distributions in (c) and (d). The selection probabilities of the continual re-assessment method (heavier solid) and the nonparametric optimal design (heavier dashed) are indicated.

values of $\sigma(\theta)$ represent a wide range of variability from 0.57 to 1. For each generated $\pi(k)$, we recorded the true v_1 ; ran the continual re-assessment method, the nonparametric optimal design and the virtual observation recursion with $\beta = 0.15, 0.29$ or 0.45 under various values of α ; and recorded their recommended doses.

Figure 2(a) plots the proportions of selecting v_1 against $\sigma(\theta)$ in 10 000 generated dose-toxicity curves. Since $\pi(v_1) \neq 0.1$ in general in the simulated configurations, it may be acceptable to select a dose with $\pi(k) \in (0.05, 0.15)$. Figure 2(b) displays the proportions of selecting such an acceptable dose; these proportions were calculated after discarding dose-toxicity curves with $\pi(v_1) \notin (0.05, 0.15)$. The virtual observation recursion with $\beta = 0.29$ is uniformly better than the continual re-assessment method and exceeds the nonparametric optimal design. Using $\beta = 0.15$ leads to inferior performance especially as $\sigma(\theta)$ increases; this is in accordance with the fact that a small β is associated with a large asymptotic variance of X_n^* . Using $\beta = 0.45$ exhibits a reverse trend where performance improves as $\sigma(\theta)$ increases and is generally better than the continual re-assessment method. This suggests not only that we can use data to form a good choice of β , but also that improvement over the continual re-assessment method can be achieved on a wide range of β .

The efficiency gain is achieved through the use of the continuous data with an additional normality assumption. Figures 2(c) and 2(d) display the operating characteristics of the virtual

observation recursion with $\beta = 0.29$ under G other than normal. The accuracies of the method are quite comparable and remain superior to the continual re-assessment method under model misspecification, thus offering assurance of robustness in our application.

6. DISCUSSION

In many dose-finding studies, the experimental drug is available only at a discrete set of doses. This is an important feature that distinguishes the phase I dose-finding literature from the large literature of stochastic approximation and its well-studied descendants, although both deal with quantile estimation. [Shen & O’Quigley \(1996, 2000\)](#) point out the difficulty in the theoretical investigation of dose-finding methods due to the discrete barrier. In this article, we introduce the idea of virtual observations to bridge the gap between stochastic approximation for continuous dosage ranges and practical situations in phase I trials. On a technical note, if we apply (2) by replacing β with some other value $0 < b < 2\beta$, then, assuming consistency, the sequence X_n^* will be asymptotically normal with mean θ_β and variance $\{b(2\beta - b)\}^{-1} \text{var}(V_i | X_i^* = \theta_\beta)$. This leads to the well-established fact that $b = \beta$ is the optimal choice that achieves the minimum asymptotic variance. Thus, by creating an objective function $h(x)$ using virtual observation, we set up a root-finding problem with a known slope β around the root. This is a nice design feature that makes the application of stochastic approximation stable and reliable in small-sample settings.

A specific contribution of this article is the use of continuous data in phase I trials, in which the safety outcomes are typically dichotomized. We demonstrate via simulation based on the NeuSTART data that our procedure does better than the continual re-assessment method and the nonparametric optimal design. At first glance, it appears inconceivable that any method can improve on the Cramér–Rao lower bound. A closer look reveals that this is to be expected, as we set out to retrieve information through the use of the original continuous data. Precisely, an efficient estimator for $\pi(k)$ based on Y_{ij} s observed at dose k has asymptotic variance proportional to $a_k^\top \Sigma a_k$, where $a_k = \{G'(c_k), c_k G'(c_k)\}^\top$ and Σ is the inverse of the information matrix for $\{M(k), \sigma(k)\}$ with diagonals $I_{11} = \int \{G''(z)/G'(z)\}^2 dG(z)$, $I_{22} = \int \{zG''(z)/G'(z) + 1\}^2 dG(z)$, and off-diagonals $I_{12} = I_{21} = \int z\{G''(z)/G'(z)\}^2 dG(z)$; see [Lehmann \(1983\)](#). For standard normal G , for instance, the ratio of $a_k^\top \Sigma a_k$ to the Cramér–Rao lower bound for $\pi(k)$ is $\{G'(c_k)\}^2(1 + c_k^2/2)[G(c_k)\{1 - G(c_k)\}]^{-1}$, and attains a maximum 0.663 when $\pi(k) = 0.21$ or 0.79, and converges to 0 as $\pi(k)$ approaches 0 or 1. With a target $p = 0.10$, this variance ratio at θ is 0.623. In other words, 33 continuous measurements observed at θ contain about the same amount of information as 42 dichotomized outcomes. To verify this intuition, we ran the nonparametric optimal design with a sample size of 42 under the simulation set-up in § 5.3 and it was uniformly superior to the virtual observation recursion with 33 subjects: the proportion of selecting v_1 was 0.58 and that of selecting an acceptable dose was 0.75. This might be because the optimal design had observations from all patients at the true v_1 but the recursion did not. The efficiency advantage should be considered in light of the computational ease and operational transparency of the stochastic approximation. First, dose assignment for the next cohort can be obtained by a hand calculator using (2). Second, the update rule (2) is intuitive and easy to explain: de-escalation occurs if V_i is large because of either a large sample mean $\bar{Y}_i$ or a large variance S_i^2 observed in the current cohort. As shown in Fig. 1(b), an increase in drug level raises not only the mean level of the liver function test but also its variability. For biological measurements, the drug impact on the variance is usually larger than that on the mean. Therefore, it is sensible for a dose-finding method to curb its escalation pace in the presence of increased variability.

ACKNOWLEDGEMENT

This work was supported by grants from the U.S. National Institute of Neurological Disorders and Stroke, part of the National Institutes of Health.

APPENDIX

Technical details

In this appendix, we first state and discuss the conditions on $\pi(x)$, $f(x)$ and $h(x)$ required by the lemma and the propositions. Recall that $v_1 = \arg \min_k |\pi(k) - p|$.

Condition 1. The functions $\pi(x)$ and $f(x)$ are assumed to be weakly monotone in x such that (a) $\pi(k') < \pi(v_1 - 1) < \pi(v_1) < \pi(v_1 + 1) < \pi(k)$, (b) $f(k') < f(v_1 - 1) < f(v_1) < f(v_1 + 1) < f(k)$, (c) $\pi(k') < \pi(v_1 - 1) < p < \pi(v_1 + 1) < \pi(k)$ and (d) $f(k') < f(v_1 - 1) < t_0 < f(v_1 + 1) < f(k)$ for all $k' < v_1 - 1$ and $k > v_1 + 1$.

To account for scenarios with $v_1 = 1$ or K , define $\pi(0) = 0$, $\pi(K + 1) = 1$, $f(0) = -\infty$ and $f(K + 1) = \infty$. Condition 1 does not require π and f to be monotone for doses below $v_1 - 1$ and above $v_1 + 1$ as long as doses below (above) v_1 are less (more) toxic than the target level. It is easy to see that strict monotonicity satisfies Condition 1. Condition 1(c) is an easy consequence of Condition 1(a) and the definition of v_1 , and Condition 1(d) follows Condition 1(b): because $f(k) = M(k) + c_k \sigma(k) + \sigma(k)(z_p - c_k) = t_0 + \sigma(k)(z_p - c_k)$, we have $\pi(v_1 - 1) < p$ implies $f(v_1 - 1) < t_0$, and $\pi(v_1 + 1) > p$ implies $f(v_1 + 1) > t_0$.

Condition 2. One and only one of the following statements is true:

- (a) The equation $h(x) = t_0$ has a unique root, denoted by θ_β , not among the jump points $\{l_2, \dots, l_K\}$.
- (b) The root of $h(x) = t_0$ does not exist.
- (c) The equation $h(x) = t_0$ has multiple roots, with the smallest root denoted by θ' and the largest root by θ'' . Both θ' and θ'' are not among the jump points.

We require that the roots (θ_β , θ' and θ'') are not among the jump points so as to satisfy the Lipschitz condition around the roots. For example, Condition 2(c) guarantees that there exist $K' > 0$ and $K'' > 0$ such that $|h(x) - h(\theta')| \leq K'|x - \theta'|$ and $|h(x) - h(\theta'')| \leq K''|x - \theta''|$ for all x .

Proof of Lemma 1. For Lemma 1(a), suppose that $\pi(v_1) \geq p$. Because $f(k) = t_0 + \sigma(k)(z_p - c_k)$, $\pi(v_1) \geq p$ implies $f(v_1) \geq t_0$. Thus, by Condition 1(d) and definition of v_2 , we know that $v_2 = v_1$ or $v_1 - 1$. More precisely, $v_1 = v_2$ if and only if $f(v_1) - t_0 < t_0 - f(k)$ if and only if $\sigma(v_1)(z_p - c_{v_1}) < \sigma(k)(c_k - z_p)$ if and only if $c_k > z_p + \sigma(v_1)(z_p - c_{v_1})/\sigma(k)$ for all $k < v_1$, giving the desired result. The proof of Lemma 1(b) is analogous. $\square$

Proof of Proposition 1(a). First, the assumption $\beta > 2\sigma(v_1)|c_{v_1} - z_p|$ implies that $h(\theta_\beta) = t_0$ for some θ_β with $C(\theta_\beta) = v_1$. When $\pi(v_1) = p$, $h(v_1) = f(v_1) = t_0$ and hence $\theta_\beta = v_1$. When $\pi(v_1) > p$, we have $h(v_1) = f(v_1) > t_0$ and $h(v_1 - 0.5) = h(v_1) - 0.5\beta = t_0 + \sigma(v_1)(z_p - c_{v_1}) - 0.5\beta < t_0$; the last inequality holds by the assumption. By continuity, there exists $\theta_\beta \in (v_1 - 0.5, v_1)$ so that $h(\theta_\beta) = t_0$ and $C(\theta_\beta) = v_1$. Similar arguments can be made for the case $\pi(v_1) < p$.

Second, the assumption $\beta < \min_{k \neq v_1} 2\sigma(k)|c_k - z_p|$ precludes the scenarios where $h(x) = t_0$ has multiple roots, because the assumption implies $h(k' + 0.5) = t_0 + \sigma(k')(z_p - c_{k'}) + 0.5\beta < t_0$ for all $k' < v_1$ and $h(k - 0.5) = t_0 + \sigma(k)(z_p - c_k) - 0.5\beta > t_0$ for all $k > v_1$.

Consequently, $h(x) = t_0$ has a unique root at θ_β . Using standard convergence results of the stochastic approximation gives $X_n^* \rightarrow \theta_\beta$ with probability one, and $X_n = C(X_n^*) = C(\theta_\beta) = v_1$ by the Lipschitz assumption; cf. Condition 2. $\square$

Proof of Proposition 1(b). When $\pi(v_1) > p$, $v_2 = v_1 - 1$ by Lemma 1(a) and the assumption $v_1 \neq v_2$. Furthermore, we have $h(v_1) - t_0 > t_0 - h(v_2) > 0.5\beta$; the first inequality holds by definition of v_2 , the

second by the assumption $\beta < \min_{k \neq v_1} 2\sigma(k)|c_k - z_p|$. As a result, the root of $h(x) = t_0$ does not exist with $h(v_2 + 0.5) < t_0$ and $h(v_1 - 0.5) > t_0$. More precisely, we have $(x - l_{v_1})\{h(x) - t_0\} > 0$ for all $x \neq l_{v_1}$ and $h(l_{v_1}) > t_0$ and $h(l_{v_1}^-) < t_0$, where $l_{v_1}^-$ indicates a number that is arbitrarily close to l_{v_1} from the left. Define $Z_i = V_i - h(X_i^*)$ and let $\mathcal{F}_i$ denote the σ -field generated by $\{X_1^*, Z_1, \dots, Z_i\}$ so that V_i is $\mathcal{F}_i$ -measurable and X_i^* is $\mathcal{F}_{i-1}$ -measurable. Following the algebraic steps similar to those of [Robbins & Monro \(1951\)](#), we obtain

$$(X_{i+1}^* - l_{v_2})^2 = (X_i^* - l_{v_1})^2 - \frac{2(X_i^* - l_{v_1})\{h(X_i^*) - t_0\}}{i\beta} - \frac{2(X_i^* - l_{v_1})Z_i}{i\beta} + O(i^{-2}).$$

Taking conditional expectation on both sides with respect to $\mathcal{F}_{i-1}$, we have

$$E\{(X_{i+1}^* - l_{v_1})^2 \mid \mathcal{F}_{i-1}\} \leq (X_i^* - l_{v_1})^2 - \frac{2(X_i^* - l_{v_1})\{h(X_i^*) - t_0\}}{i\beta} + \frac{K'''}{i^2}$$

for some $K''' > 0$. Applying Theorem 1 of [Robbins & Siegmund \(1971\)](#) for nonnegative almost supermartingale, we can show that $\lim_{i \rightarrow \infty} (X_i^* - l_{v_1})^2$ exists and $\sum_{i=1}^{\infty} (X_i^* - l_{v_1})\{h(X_i^*) - t_0\} < \infty$ almost surely. Also, since $(X_i^* - l_{v_2})\{h(X_i^*) - t_0\} > 0$, we can conclude that $X_i^* \rightarrow l_{v_1}$ and $\text{pr}(X_n = v_1 \text{ or } v_2 \text{ eventually}) = 1$. By definition of v_1 , $|\pi(v_1) - p| < |\pi(v_2) - p|$. Thus, for sufficiently large n ,

$$\begin{aligned} |\pi(X_n) - p| &\leq |\pi(v_2) - p| = \{\pi(v_1) - \pi(v_2)\} - \{\pi(v_1) - p\} \\ &\leq \left| G\left\{z_p + \frac{\sigma(v_1)}{\sigma(v_2)}(z_p - c_{v_1})\right\} - G(c_{v_1}) \right| - |\pi(v_1) - p|. \end{aligned} \quad (\text{A1})$$

The inequality (A1) is due to Lemma 1(a) under the assumptions $v_1 \neq v_2$ and $\pi(v_1) > p$. The proof is completed by analogous arguments for the case $\pi(v_1) < p$. $\square$

Proof of Proposition 3(a). First, it is easy to see that $\sigma(k) < \sigma(\theta) < \sigma(k')$ for $k < \theta < k'$ under the assumed mean-variance relationship and Condition 1. Since $s\{\sigma(k)\} - s\{\sigma(\theta)\} = \phi\{M(k) - M(\theta)\} = \phi\{f(k) - f(\theta)\} - \phi\{\sigma(k) - \sigma(\theta)\}$, we have $f(k) < f(\theta)$ if and only if $\sigma(k) < \sigma(\theta)$ because $s(\sigma)$ is increasing. Now, the assumption $p_L < \tilde{p}_L$ implies

$$c_k > \tilde{z}_p^* + \frac{\sigma(\theta)}{\sigma(1)}|z_p - \tilde{z}_p^*| \geq \tilde{z}_p^* + \frac{\sigma(\theta)}{\sigma(k)}|z_p - \tilde{z}_p^*|$$

because $\sigma(k) < \sigma(\theta)$. Hence $\tilde{B}_k = 2\sigma(k)(c_k - \tilde{z}_p^*) > 2\sigma(\theta)|z_p - \tilde{z}_p^*| = \tilde{B}_\theta$ for all $k < \theta$. Using similar arguments, we can show $p_U > \tilde{p}_U$ implies $\tilde{B}_{k'} > \tilde{B}_\theta$ for $k' > \theta$. Thus, $\tilde{v} = \theta$ by definition. $\square$

Proof of Proposition 3(b). First, following similar arguments in the proof of Proposition 1(a), we can prove that $\tilde{h}(x) = t_0$ has a unique root at $\tilde{\theta}_\beta$ with $C(\tilde{\theta}_\beta) = \tilde{v}$ if $\tilde{B}_{\tilde{v}} < \beta < \min_{k \neq \tilde{v}} \tilde{B}_k$. Applying standard convergence results of stochastic approximation gives the consistency of X_n for $\tilde{v}$, which is equal to θ by Proposition 3(a).

Next, because $c_k \leq z_U < \tilde{z}_p^*$ by the assumption $p_U < \tilde{p}_U$ and $\sigma(\theta) \leq \sigma(k)$ for $k > \theta$, we have $2\sigma(\theta)(\tilde{z}_p^* - z_U) \leq 2\sigma(k)(\tilde{z}_p^* - c_k) = B_k$ for $k > \theta$. Hence,

$$2\sigma(\theta)(\tilde{z}_p^* - z_U) \leq \min_{k > \theta} B_k. \quad (\text{A2})$$

For the case $k < \theta$, we have $c_k \geq z_L > \max(z_p, \tilde{z}_p^*)$ and $\sigma(k) \leq \sigma(\theta)$. Substituting $M(k) = M(\theta) + [s\{\sigma(k)\} - s\{\sigma(\theta)\}]/\phi$ into $c_k = \{t_0 - M(k)\}/\sigma(k)$ gives

$$c_k\sigma(k) = t_0 - M(\theta) - \frac{s\{\sigma(k)\} - s\{\sigma(\theta)\}}{\phi} \quad (\text{A3})$$

for $\phi > 0$. Fix $M(\theta)$ and $\sigma(\theta)$. For any given c_k , we differentiate both sides of (A3) with respect to ϕ and get

$$\left[\frac{s'\{\sigma(k)\}}{\phi} + c_k \right] \frac{\partial \sigma(k)}{\partial \phi} = \frac{s\{\sigma(k)\} - s\{\sigma(\theta)\}}{\phi^2}$$

for $\phi > 0$. Therefore, $\partial\sigma(k)/\partial\phi < 0$ as long as $c_k \geq 0$ which holds for most practical purposes; i.e. it suffices to have $z_p \geq 0$, and $\sigma(k)$ is minimized as ϕ grows large upon which $M(k) \rightarrow M(\theta)$. Specifically, its infimum is achieved by taking the limit $\phi \rightarrow \infty$ in (A3):

$$\inf_{\phi > 0} \sigma(k) = \frac{t_0 - M(\theta)}{c_k} < \frac{t_0 - M(\theta)}{z_p} = \sigma(\theta).$$

Also, since $\sigma(k) = \sigma(\theta)$ when $\phi = 0$, the above expression provides the minimum value of $\sigma(k)$ over all possible $\phi \geq 0$ for a given c_k . Thus,

$$\min_{k < \theta} B_k \geq \min_{k < \theta} \inf_{\phi \geq 0} 2\sigma(k)(c_k - \tilde{z}_p^*) = 2\{t_0 - M(\theta)\} \min_{k < \theta} (1 - \tilde{z}_p^*/c_k) = 2\sigma(\theta)z_p(1 - \tilde{z}_p^*/z_L).$$

This, together with (A2), implies that $|z_p - \tilde{z}_p^*| < \beta/\{2\sigma(\theta)\} < \tilde{w}$ suffices the consistency condition, $\tilde{B}_\theta < \beta < \min_{k \neq \theta} \tilde{B}_k$. $\square$

REFERENCES

- ANBAR, D. (1977). The application of stochastic methods to the bioassay problem. *J. Statist. Plan. Inference* **1**, 191–206.
- ANBAR, D. (1984). Stochastic approximation methods and their use in bioassay and phase I clinical trials. *Commun. Statist.* **13**, 2451–67.
- CHEUNG, Y. K. (2005). Coherence principles in dose-finding studies. *Biometrika* **92**, 863–73.
- CHEUNG, Y. K. (2007). Sequential implementation of stepwise procedures for identifying the maximum tolerated dose. *J. Am. Statist. Assoc.* **102**, 1448–61.
- CHEUNG, Y. K. & CHAPPELL, R. (2002). A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics* **58**, 671–4.
- ELKIND, M. S., SACCO, R. L., MACARTHUR, R. B., FINK, D. J., PEERSCHKE, E., ANDREWS, H., NEILS, G., STILLMAN, J., CHONG, J., CONNOLLY, S., CORPORAN, T., LEIFER, D. & CHEUNG, K. (2008). The neuroprotection with statin therapy for acute recovery trial (NeuSTART): an adaptive design phase I dose-escalation study of high-dose lovastatin in acute ischemic stroke. *Int. J. Stroke* **3**, 210–8.
- LAI, T. L. & ROBBINS, H. (1979). Adaptive design and stochastic approximation. *Ann. Statist.* **7**, 1196–221.
- LEHMANN, E. L. (1983). *Theory of Point Estimation*. New York: Wiley.
- LIN, Y. & SHIH, W. J. (2001). Statistical properties of the traditional algorithm-based designs for phase I cancer trials. *Biostatistics* **2**, 203–15.
- O'QUIGLEY, J. & CHEVRET, S. (1991). Methods for dose finding studies in cancer clinical trials: a review and results of a monte carlo study. *Statist. Med.* **19**, 1647–64.
- O'QUIGLEY, J., PAOLETTI, X. & MACCARIO, J. (2002). Non-parametric optimal design in dose finding studies. *Biostatistics* **3**, 51–56.
- O'QUIGLEY, J., PEPE, M. & FISHER, L. (1990). Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* **46**, 33–48.
- O'QUIGLEY, J. & SHEN, L. Z. (1996). Continual reassessment method: a likelihood approach. *Biometrics* **52**, 673–84.
- ROBBINS, H. & MONRO, S. (1951). A stochastic approximation method. *Ann. Math. Statist.* **22**, 400–7.
- ROBBINS, H. & SIEGMUND, D. (1971). A convergence theorem for non-negative almost supermartingales and some applications. In *Optimizing Methods in Statistics*, Ed. J. S. Rustagi, pp. 233–57. New York: Academic Press.
- SACKS, J. (1958). Asymptotic distribution of stochastic approximation procedures. *Ann. Math. Statist.* **29**, 373–405.
- SHEN, L. Z. & O'QUIGLEY, J. (1996). Consistency of continual reassessment method under model misspecification. *Biometrika* **83**, 395–405.
- SHEN, L. Z. & O'QUIGLEY, J. (2000). Using a one-parameter model to sequentially estimate the root of a regression function. *Comp. Statist. Data Anal.* **34**, 357–69.
- WU, C. F. J. (1985). Efficient sequential designs with binary data. *J. Am. Statist. Assoc.* **80**, 974–84.

[Received October 2008. Revised June 2009]