

Efficient Robbins-Monro Procedure for Binary Data

V. Roshan Joseph

School of Industrial and Systems Engineering

Georgia Institute of Technology

Atlanta, GA 30332-0205, USA

roshan@isye.gatech.edu

Accepted by *Biometrika*

SUMMARY

It is well known that Robbins-Monro procedure does not perform well in the estimation of extreme quantiles. This is because the procedure is implemented using some asymptotic results, which are not suitable to deal with binary data. Here we propose a modification of Robbins-Monro procedure and derive the optimal procedure for binary data under some reasonable approximations. The improvement obtained by using the optimal procedure for the estimation of extreme quantiles is substantial.

Some key words: Quantile estimation; Sequential designs; Stochastic approximation.

1. INTRODUCTION

In many applications, interest centers around finding the threshold of a variable that will cause certain amount of successes (or failures) in the output. For examples, an explosive designer will be interested to find the level of shock necessary to make 99.99% of the explosives to fire (Neyer, 1994), a quality engineer will be interested to find the level of the stack force necessary to make 0.1% of multifeeds in the paper feeder of a copier machine (Joseph and Wu, 2002), and so on. The problem can be formally stated as follows. Let Y be a binary response with probability of success $M(x)$, where x is the variable whose threshold is of interest. The objective is to find the value of $x = \theta$ that will make the probability of success equal to α , i.e. $M(\theta) = \alpha$. Usually $M(x)$ is a distribution function obtained by assuming some distribution for a latent variable underlying the binary response. Therefore θ can be regarded as the α -quantile of this distribution. The problem is that the function $M(x)$ is unknown to the experimenter. However, the experimenter can observe Y at different values of x in order to find θ . The objective is to devise an experimental strategy that will help to estimate θ with minimum number of observations and with great accuracy. The experiment can be performed by using a *sequential design* in which the values of x 's are chosen sequentially which are allowed to depend on the already observed data.

One sequential design strategy known as stochastic approximation is to choose x_1, x_2, \dots such that $x_n \rightarrow \theta$ in probability. Robbins and Monro (1951) proposed the following sequential procedure:

$$x_{n+1} = x_n - a_n(y_n - \alpha), \tag{1}$$

where y_n is the binary response observed at x_n and $\{a_n\}$ is a pre-specified sequence of positive constants. Robbins and Monro showed that if a_n satisfies the conditions

$$\sum_{n=1}^{\infty} a_n = \infty, \quad \text{and} \quad \sum_{n=1}^{\infty} a_n^2 < \infty, \quad (2)$$

then $x_n \xrightarrow{p} \theta$. Robbins and Monro did not explain how to choose the sequence $\{a_n\}$ optimally. They recommended the simple choice $a_n = c/n$ for some constant c . Based on the results of Chung (1954), Hodges and Lehmann (1956), and Sacks (1958), the procedure is fully asymptotically efficient with $a_n = \{n\dot{M}(\theta)\}^{-1}$ under some conditions, where \dot{M} denotes the first derivative of M . This choice became the standard choice for the practical implementation of the Robbins-Monro procedure (see, for example, Wetherill and Glazebrook, 1986). The optimal choice of a_n in small samples has not been investigated, although in most experiments the small sample size is the most interesting case.

Wetherill (1963) studied the performance of the Robbins-Monro procedure through simulations and found that it works quite well for $\alpha = 0.5$ (LD50) but performs very poorly for extreme quantiles. He proposed several modifications to the procedure but again they were found to be inadequate for extreme quantiles and concluded that the Robbins-Monro and related procedures should be used only for the estimation of quantiles in the immediate neighborhood of LD50. Cochran and Davis (1965) and Young and Easterling (1994) also obtained similar conclusions based on extensive simulation results.

The Robbins-Monro procedure is a very general procedure which can be used for several kinds of data and is not restricted to binary data. Binary data on the other hand possesses some peculiar properties, which can be exploited to improve the procedure. Two such properties

are relevant: the variance of Y is a function of x given by $M(x)\{1 - M(x)\}$ and $M(x)$ is a distribution function. In this article we will make use of these properties to produce optimal Robbins-Monro procedures for binary data. We will also propose a slight modification to the Robbins-Monro procedure to get better convergence properties.

Several methods for quantile estimation based on binary data are proposed in the literature. See for example Wu (1985), McLeish and Tosh (1990), Kalish (1990), Neyer (1994), Sitter and Wu (1999), and Voelkel (2003) among many others. The objective of this study is to find the optimal sequential procedure within the class of the Robbins-Monro type procedures, not to find the best overall method.

2. OPTIMAL ROBBINS-MONRO PROCEDURE

In most applications, the distribution function $M(x)$ is from a location family with location parameter θ . Therefore hereafter we denote $M(x)$ by $M(x - \theta)$. Thus $M(0) = \alpha \in (0, 1)$ is specified. We also assume that $\dot{M}(0) > 0$ is known. The experimenter starts the experiment at some value x_1 , which is believed to be close to θ based on some prior knowledge. Therefore we may choose a prior distribution for θ say with $E(\Theta) = x_1$ and $var(\Theta) = \tau_1^2 < \infty$, where τ_1 represents the initial uncertainty of θ with respect to x_1 . Let $Z_n = x_n - \Theta$. Note that although x_1 is a fixed quantity, x_2, \dots, x_n are random due to their dependence on past data. Consider a *modified* Robbins-Monro process given by

$$x_{n+1} = x_n - a_n(y_n - b_n).$$

For binary data, $\{b_n\}$ is a sequence of constants in $(0, 1)$. They need not be equal to α , but is expected to get close to α as n gets larger. We will see that using a b_n different from α ,

the performance of the Robbins-Monro procedure can be greatly improved. Thus our model is,

$$\begin{aligned} y_n|Z_n &\sim \text{Bernoulli}\{M(Z_n)\}, \\ Z_{n+1} &= Z_n - a_n(y_n - b_n), \\ E(Z_1) &= 0 \text{ and } E(Z_1^2) = \tau_1^2. \end{aligned}$$

The objective is to find sequences $\{a_n\}$ and $\{b_n\}$ such that $Z_n \xrightarrow{p} 0$ at the fastest rate. First we investigate the conditions under which the desired convergence can be obtained.

Suppose the sequence $\{b_n\}$ satisfies the condition

$$\sum_{n=2}^{\infty} a_n |b_n - \alpha| \sum_{j=1}^{n-1} a_j < \infty, \quad (3)$$

then we have the following convergence result whose proof closely follows that of Robbins and Monro (1951). The above condition together with (2) ensures that b_n converges to α . Moreover, because $\sum_{j=1}^{n-1} a_j$ increases with n , the convergence of b_n to α should be fast enough for (3) to hold.

THEOREM 1 *If (2) and (3) hold, then $Z_n \xrightarrow{p} 0$.*

Proof: We have

$$E(Z_{n+1}^2|Z_n) = Z_n^2 - 2a_n Z_n \{M(Z_n) - b_n\} + a_n^2 E\{(y_n - b_n)^2|Z_n\}.$$

Let $\tau_n^2 = E(Z_n^2)$, $d_n = E[Z_n \{M(Z_n) - b_n\}]$, and $e_n = E\{(y_n - b_n)^2\}$. Then

$$\tau_{n+1}^2 = \tau_n^2 - 2a_n d_n + a_n^2 e_n, \quad (4)$$

which gives

$$\tau_{n+1}^2 = \tau_1^2 - 2 \sum_{j=1}^n a_j d_j + \sum_{j=1}^n a_j^2 e_j. \quad (5)$$

Since $b_n \in (0, 1)$ and $y_n = 0$ or 1 , we have $0 < e_n < 1$ for all n . Thus $0 < \sum_{j=1}^n a_j^2 e_j < \sum_{j=1}^n a_j^2$.

Therefore by condition (2) the positive-term series $\sum_{j=1}^n a_j^2 e_j$ converges.

We have

$$\sum_{j=1}^n a_j d_j = \sum_{j=1}^n a_j E[Z_j \{M(Z_j) - b_j\}] = \sum_{j=1}^n a_j E[Z_j \{M(Z_j) - \alpha\}] + \sum_{j=1}^n a_j (\alpha - b_j) E(Z_j).$$

Since M is a distribution function and $\dot{M}(0) > 0$, there exists δ and δ' such that $0 < \delta \leq (M(z) - \alpha)/z \leq \delta' < \infty$ for all z . This implies $E[Z_j \{M(Z_j) - \alpha\}] \geq \delta E(Z_j^2) = \delta \tau_j^2$. Thus

$$\sum_{j=1}^n a_j d_j \geq \delta \sum_{j=1}^n a_j \tau_j^2 + \sum_{j=1}^n a_j (\alpha - b_j) E(Z_j).$$

Consider the magnitude of the second term on the right side,

$$\left| \sum_{j=1}^n a_j (\alpha - b_j) E(Z_j) \right| \leq \sum_{j=1}^n a_j |\alpha - b_j| |E(Z_j)| < \sum_{j=2}^n a_j |\alpha - b_j| \sum_{i=1}^{j-1} a_i,$$

because $Z_n = Z_1 - \sum_{j=1}^{n-1} a_j (y_j - b_j) \Rightarrow |E(Z_n)| < \sum_{j=1}^{n-1} a_j$. Also since $\tau_{n+1}^2 \geq 0$ for all n , from

(5) we get $\sum_{j=1}^n a_j d_j \leq (\tau_1^2 + \sum_{j=1}^n a_j^2 e_j)/2$. Thus

$$\begin{aligned} \delta \sum_{j=1}^n a_j \tau_j^2 &\leq \sum_{j=1}^n a_j d_j - \sum_{j=1}^n a_j (\alpha - b_j) E(Z_j) \\ &< \frac{1}{2} (\tau_1^2 + \sum_{j=1}^n a_j^2 e_j) + \sum_{j=2}^n a_j |\alpha - b_j| \sum_{i=1}^{j-1} a_i. \end{aligned}$$

The right side converges by condition (3) and therefore the positive-term series $\sum_{j=1}^n a_j \tau_j^2$

converges. Now

$$\begin{aligned} \sum_{j=1}^n |a_j d_j| &\leq \sum_{j=1}^n a_j |E[Z_j \{M(Z_j) - \alpha\}]| + \sum_{j=1}^n a_j |\alpha - b_j| |E(Z_j)| \\ &< \delta' \sum_{j=1}^n a_j \tau_j^2 + \sum_{j=2}^n a_j |\alpha - b_j| \sum_{i=1}^{j-1} a_i, \end{aligned}$$

Because the right side converges, the series $\sum_{j=1}^n a_j d_j$ converges absolutely and hence the series converges. Thus from (5), $\lim \tau_n^2$ exists and is equal to $\tau_1^2 - 2 \sum_{j=1}^{\infty} a_j d_j + \sum_{j=1}^{\infty} a_j^2 e_j$. But since $0 \leq \sum_{n=1}^{\infty} a_n \tau_n^2 < \infty$ and $\sum_{n=1}^{\infty} a_n = \infty$, this limit must be 0 and therefore $Z_n \rightarrow^p 0$.

◇

There are infinite number of sequences for $\{a_n\}$ and $\{b_n\}$ satisfying the conditions of the Theorem 1. We are interested to find the particular sequence that gives the best convergence properties. We propose to choose a_n and b_n such that $E(Z_{n+1}^2)$ is a minimum subject to the condition that $E(Z_{n+1}) = 0$. Similar ideas of choosing two sequences to minimize the conditional mean square error in sequential designs was employed by Hu (1997, 1998) in a Bayesian framework.

We have

$$E(Z_{n+1}) = E(Z_n) - a_n[E\{M(Z_n)\} - b_n] = 0.$$

Because the sequence a_1, \dots, a_{n-1} and b_1, \dots, b_{n-1} are obtained such that $E(Z_2) = \dots = E(Z_n) = 0$, we obtain

$$b_n = E\{M(Z_n)\}. \quad (6)$$

From (4), we have

$$\tau_{n+1}^2 = \tau_n^2 - 2a_n E[Z_n\{M(Z_n) - b_n\}] + a_n^2 \left([E\{M(Z_n)\} - b_n]^2 + E\{M(Z_n)\}[1 - E\{M(Z_n)\}] \right). \quad (7)$$

Minimizing τ_{n+1}^2 with respect to a_n and using (6), we obtain

$$a_n = \frac{E\{Z_n M(Z_n)\}}{E\{M(Z_n)\}[1 - E\{M(Z_n)\}]}. \quad (8)$$

Unfortunately, the optimal sequences depend on the function M which is unknown to the experimenter. Therefore the best we can do is to choose a function G that can closely approximate the true function M and derive the sequences. The resulting Robbins-Monro procedure is optimal only when $M = G$ and approximately optimal otherwise. It can be considered as an efficient procedure as long as the deviation of G from the true function is not severe. The choice of $a_n = \{n\dot{M}(0)\}^{-1}$ is based on a linear approximation of M around 0, which is not good for a distribution function particularly in the tail areas.

Consider an approximation for $M(z)$ given by,

$$G(z) = \Phi\{\Phi^{-1}(\alpha) + \beta z\}, \quad (9)$$

where $\beta = \dot{M}(0)/\phi\{\Phi^{-1}(\alpha)\}$, Φ is the standard normal distribution function, and ϕ its density function. Now a_n and b_n can be obtained from (6) and (8) using $G(z)$ instead of $M(z)$. We also need the distribution of Z_n to evaluate the expectations in (6) and (8). It is easy to show that the density function of Z_{n+1} is

$$f_{Z_{n+1}}(z) = \{1 - G(z - a_n b_n)\}f_{Z_n}(z - a_n b_n) + G\{z + a_n(1 - b_n)\}f_{Z_n}\{z + a_n(1 - b_n)\}.$$

The $f_{Z_n}(z)$ can be recursively computed starting with $f_{Z_1}(z)$. Let $Z_1 \sim N(0, \tau_1^2)$, so $f_{Z_1}(z) = 1/\tau_1\phi(z/\tau_1)$. Clearly the expectations in (6) and (8) are difficult to compute with this exact distribution. Therefore we resort to approximation. It is quite natural to approximate the distribution of Z_n by $N(0, \tau_n^2)$ as the first two moments of the two distributions match exactly. It turns out that this is a very good approximation as verified by plotting these functions

for various values of α , τ_1 , and n . Using this approximation we obtain

$$b_n = E\{G(Z_n)\} = \int_{-\infty}^{\infty} \Phi\{\Phi^{-1}(\alpha) + \beta z\} 1/\tau_n \phi(z/\tau_n) dz = \Phi\left\{\frac{\Phi^{-1}(\alpha)}{(1 + \beta^2 \tau_n^2)^{\frac{1}{2}}}\right\},$$

and

$$\begin{aligned} a_n &= \frac{E\{Z_n G(Z_n)\}}{E\{G(Z_n)\}[1 - E\{G(Z_n)\}]} = \frac{1}{b_n(1 - b_n)} \int_{-\infty}^{\infty} z \Phi\{\Phi^{-1}(\alpha) + \beta z\} 1/\tau_n \phi(z/\tau_n) dz \\ &= \frac{1}{b_n(1 - b_n)} \frac{\beta \tau_n^2}{(1 + \beta^2 \tau_n^2)^{\frac{1}{2}}} \phi\left\{\frac{\Phi^{-1}(\alpha)}{(1 + \beta^2 \tau_n^2)^{\frac{1}{2}}}\right\}. \end{aligned}$$

Substituting in (7) we obtain $\tau_{n+1}^2 = \tau_n^2 - b_n(1 - b_n)a_n^2$, which shows that there is an improvement by moving from x_n to x_{n+1} . Let $c_n = \beta a_n b_n(1 - b_n)$ and $\nu_n = \beta^2 \tau_n^2$. Then the *optimal Robbins-Monro procedure* can be written as

$$x_{n+1} = x_n - \frac{c_n}{\beta b_n(1 - b_n)}(y_n - b_n), \quad (10)$$

where

$$c_n = \frac{\nu_n}{(1 + \nu_n)^{\frac{1}{2}}} \phi\left\{\frac{\Phi^{-1}(\alpha)}{(1 + \nu_n)^{\frac{1}{2}}}\right\}, \quad b_n = \Phi\left\{\frac{\Phi^{-1}(\alpha)}{(1 + \nu_n)^{\frac{1}{2}}}\right\} \quad \text{and} \quad \nu_{n+1} = \nu_n - \frac{c_n^2}{b_n(1 - b_n)}, \quad (11)$$

with $\nu_1 = \beta^2 \tau_1^2$. Note that c_n and b_n are sequences that can be specified before the experiment. They can be easily computed once $\nu_1 \in (0, \infty)$ is specified. The procedure can work only with a finite value of ν_1 , which implies that a noninformative prior for θ cannot be used. Therefore the Bayesian formulation of the problem with a proper prior was very crucial in the development of the above procedure. After n experiments the best estimate of θ is x_{n+1} . We can also get a $(1 - \gamma)$ credible interval for θ as $x_{n+1} \pm \Phi^{-1}(\gamma/2)\tau_{n+1}$, where $\tau_{n+1} = \nu_{n+1}^{\frac{1}{2}}/\beta$.

PROPOSITION 1 *As $n \rightarrow \infty$, $\nu_n \rightarrow 0$, $b_n \rightarrow \alpha$, and $c_n \rightarrow 0$.*

Proof: From (11), we have

$$\nu_{n+1} = \nu_n - \frac{\nu_n^2}{1 + \nu_n} I\left\{\frac{\Phi^{-1}(\alpha)}{(1 + \nu_n)^{\frac{1}{2}}}\right\}, \quad (12)$$

where $I(u) = \phi^2(u)/(\Phi(u)\{1 - \Phi(u)\})$ is the Fisher information of u in binary data with probability of success equal to $\Phi(u)$. It is well known that $0 < I(u) \leq 2/\pi$. Thus $\nu_{n+1} < \nu_n$ and $\nu_{n+1} \geq \nu_n(1 - 2/\pi) \geq \nu_1(1 - 2/\pi)^n > 0$ for all n . Hence the sequence $\{\nu_n\}$ converges. Let $h(\nu_n)$ be the right side of (12). Then $\lim \nu_n$ satisfies the equation $\nu = h(\nu)$, for which 0 is a unique solution. Thus from (11), we obtain $b_n \rightarrow \alpha$, and $c_n \rightarrow 0$. \diamond

By Proposition 1 we have $\tau_n^2 \rightarrow 0$ and therefore $Z_n \xrightarrow{p} 0$ if M is equal to the normal distribution function given in (9). We can expect the result to hold even when the true model is different from it. Fortunately this hypothesis is true as stated in the following proposition.

PROPOSITION 2 *For the procedure in (10), $Z_n \xrightarrow{p} 0$.*

Proof: Let $n^* = 1/I\{\Phi^{-1}(\alpha)\}$ and $\nu^* = 1/I^2\{\Phi^{-1}(\alpha)\}$. We have $\dot{h}(0) = 1$. Since $\dot{h}(\nu)$ is a continuous function in ν , there exists a $\bar{\nu} \in (0, \nu^*]$ such that $\dot{h}(\nu) > 0$ for all $\nu \leq \bar{\nu}$. Also since $\nu_n \rightarrow 0$, there exists an \bar{n} such that $\nu_n < \bar{\nu}$ for all $n \geq \bar{n}$. Let $\tilde{n} = \max\{\lceil n^* \rceil, \bar{n}, \lceil \nu^*/\bar{\nu} \rceil\}$ and $\nu' = \tilde{n}\bar{\nu}$, where $\lceil x \rceil$ is the smallest integer greater than or equal to x . Thus $\nu_n \leq \nu_{\tilde{n}} < \bar{\nu} = \nu'/\tilde{n}$. Suppose $\nu_n \leq \nu'/n$ for some $n \geq \tilde{n}$. Since $\nu' \geq \nu^* > 1$, we have for $n \geq \tilde{n}$,

$$I\left\{\frac{\Phi^{-1}(\alpha)}{(1 + \nu'/n)^{\frac{1}{2}}}\right\} \geq I\{\Phi^{-1}(\alpha)\} = \frac{n^* + \nu^*}{(n^* + 1)\nu^*} \geq \frac{n + \nu^*}{(n + 1)\nu^*} \geq \frac{n + \nu'}{(n + 1)\nu'}.$$

Since $\nu_n \leq \nu'/n \leq \nu'/\tilde{n} = \bar{\nu}$ and $h(\nu)$ is increasing in ν for all $\nu \leq \bar{\nu}$, from (12) we get

$$\nu_{n+1} \leq \frac{\nu'}{n} \left[1 - \frac{\nu'}{n + \nu'} I\left\{\frac{\Phi^{-1}(\alpha)}{(1 + \nu'/n)^{\frac{1}{2}}}\right\}\right] \leq \frac{\nu'}{n + 1}.$$

Then by mathematical induction $\nu_n \leq \nu'/n$ for all $n \geq \tilde{n}$.

Let $\nu'' = \min\{1, \tilde{n}\nu_{\tilde{n}}\}$. Then $\nu_{\tilde{n}} \geq \nu''/\tilde{n}$. Suppose $\nu_n \geq \nu''/n$ for some $n \geq \tilde{n}$. Since $\nu'' \leq 1$, we have $I\{\Phi^{-1}(\alpha)/(1 + \nu''/n)^{\frac{1}{2}}\} < 1 \leq (n + \nu'')/((n + 1)\nu'')$. Because $\nu''/n \leq \nu_n \leq \nu'/n \leq \bar{\nu}$, from (12) we get

$$\nu_{n+1} \geq \frac{\nu''}{n} \left[1 - \frac{\nu''}{n + \nu''} I\left\{ \frac{\Phi^{-1}(\alpha)}{(1 + \nu''/n)^{\frac{1}{2}}} \right\} \right] \geq \frac{\nu''}{n + 1}.$$

Then by mathematical induction $\nu_n \geq \nu''/n$ for all $n \geq \tilde{n}$.

We have for all $n \geq \tilde{n}$

$$\begin{aligned} a_n &= \frac{\nu_n}{\beta(1 + \nu_n)^{\frac{1}{2}}} \frac{\phi\{\Phi^{-1}(\alpha)/(1 + \nu_n)^{\frac{1}{2}}\}}{\Phi\{\Phi^{-1}(\alpha)/(1 + \nu_n)^{\frac{1}{2}}\}[1 - \Phi\{\Phi^{-1}(\alpha)/(1 + \nu_n)^{\frac{1}{2}}\}]} \\ &\leq \frac{\nu'}{\beta(n^2 + n\nu')^{\frac{1}{2}}} \frac{1/\sqrt{2\pi}}{\alpha(1 - \alpha)} \leq \frac{\nu'}{\sqrt{2\pi}\alpha(1 - \alpha)\beta n}, \\ \text{and } a_n &\geq \frac{\nu''}{\beta(n^2 + n\nu'')^{\frac{1}{2}}} \frac{\phi\{\Phi^{-1}(\alpha)\}}{1/2(1 - 1/2)} \geq \frac{4\nu''\phi(\Phi^{-1}(\alpha))}{\beta(n + 1)}. \end{aligned}$$

Thus $\{a_n\}$ satisfies (2). Also $\sum_{j=1}^{n-1} a_j = O(\log n)$ as $n \rightarrow \infty$.

Using Taylor series expansion we obtain

$$b_n = \Phi\left\{ \frac{\Phi^{-1}(\alpha)}{(1 + \nu_n)^{\frac{1}{2}}} \right\} = \alpha - \nu_n \frac{\Phi^{-1}(\alpha)}{2} \phi\{\Phi^{-1}(\alpha)\} + O(\nu_n^2).$$

Thus $|b_n - \alpha| = O(1/n)$, which implies $a_n|b_n - \alpha| \sum_{j=1}^{n-1} a_j = O(\log n/n^2)$ and therefore (3)

holds. Thus by Theorem 1, $Z_n \xrightarrow{p} 0$. \diamond

It is clear from (11) that b_n always lies between α and $1/2$. The consequence of this is important. Even if we are interested in extreme quantiles, for small n , the optimal Robbins-Monro procedure is operated as though we are interested in a quantile between α and LD50. The search is moved closer to α as n gets larger (since $b_n \rightarrow \alpha$). This is markedly different

from the ordinary Robbins-Monro procedure in (1). Wetherill (1963) noted that the Robbins-Monro procedure performs miserably for the estimation of small or large quantiles. This is mainly because of the unequal up and down movements of the procedure. By using b_n instead of α this imbalance is somewhat mitigated. Consider an example of Wetherill. Suppose $\alpha = 0.25$ and the true model is a logistic function given by $M(z) = \{1 + 3 \exp(-z)\}^{-1}$. Suppose that the experiment started exactly at θ . If we observe $y_1 = 1$, which can happen with probability 0.25, then $x_2 = \theta - 1/\dot{M}(0)(1 - 0.25) = \theta - 4$. Thus x_2 is lower than θ . The search will move up if we observe $y = 0$. Suppose we observed $y_2 = \dots = y_n = 0$, then

$$x_{n+1} = \theta - 4 + \frac{0.25}{0.25 \times 0.75} \left(\frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n} \right),$$

which crosses θ from below at $n = 31$. In other words, it will take at least 30 steps to get back to the true value θ . Now consider the optimal Robbins-Monro procedure in (10). Here $\beta = \dot{M}(0)/\phi\{\Phi(0.25)\} = 0.59$. We select τ_1 such that x_2 is exactly the same as that in the Robbins-Monro procedure. This value is given by 4.475. Now $x_7 = \theta - 4 + 1.579 + 0.973 + 0.655 + 0.478 + 0.370 = \theta + 0.055$. Thus the sequence crosses θ from below in just 5 steps as opposed to the 30 steps taken by the Robbins-Monro procedure. This example clearly demonstrates the superiority of the optimal Robbins-Monro procedure over the Robbins-Monro procedure.

Note that we require the value of $\dot{M}(0)$ to compute β . In practice, it is unlikely that the experimenter will know the exact value of $\dot{M}(0)$ and the procedure has to be implemented using a guess value of $\dot{M}(0)$. It is clear from the proof of Proposition 2 that the procedure will work irrespective of the value of β . But the convergence can slow down if the guess

value is far away from the true value. If a good guess cannot be made on $\dot{M}(0)$, one can try to adaptively estimate it from the data by fitting a parametric model such as the one in (9). For the Robbins-Monro procedure such an adaptive procedure under some truncation rule gives the same asymptotic performance as that of the original procedure. See Lai and Robbins (1979) and Wu (1985) for more details.

3. SIMULATIONS

We now compare the performance of the optimal Robbins-Monro procedure in (10) with the Robbins-Monro procedure in (1) through simulations. For the Robbins-Monro procedure a_n is taken as $\{n\dot{M}(0)\}^{-1}$, which is the standard choice. Six models are selected for the simulation study:

$$\begin{aligned} \text{Normal distribution: } M(z) &= \Phi\{\Phi^{-1}(\alpha) + z\}, \\ \text{Uniform distribution: } M(z) &= \begin{cases} 0 & , z < -3\alpha \\ \alpha + z/3 & , -3\alpha \leq z \leq 3(1 - \alpha) \\ 1 & , z > 3(1 - \alpha) \end{cases} , \end{aligned}$$

$$\text{Logistic distribution: } M(z) = \left(1 + \frac{1 - \alpha}{\alpha} e^{-z}\right)^{-1},$$

$$\text{Extreme value distribution: } M(z) = 1 - \exp\{\log(1 - \alpha) e^z\},$$

$$\text{Skewed logistic distribution: } M(z) = \left(1 + \frac{1 - \sqrt{\alpha}}{\sqrt{\alpha}} e^{-z}\right)^{-2},$$

$$\text{Cauchy distribution: } M(z) = \frac{1}{2} + \frac{1}{\pi} \tan^{-1}\{z + \tan(\pi\alpha - \pi/2)\}.$$

These functions are plotted in Figure 1. Let $\theta = 0$. We choose 20 samples to estimate θ . Thus the best estimate of θ is x_{21} . We let $\tau_1 = 1$ as the initial uncertainty, which means

we are about 95% sure that the θ is in ± 2 of the starting value x_1 . The starting value is randomly generated from $N(0, \tau_1^2)$ and is used for both procedures. Then at different values of α between 0.1 and 0.9, 1000 simulations were performed. The mean square error of x_{21} is calculated and is plotted in Figure 2. We see that the optimal Robbins-Monro procedure performs uniformly better than the Robbins-Monro procedure with great improvement for the extreme quantiles. The performance of the two procedures around LD50 is comparable, which agrees with the findings of Wetherill (1963). Figure 3 shows the estimated density curves for x_{21} for different α values with logistic distribution as the true model. It clearly shows how the performance of the Robbins-Monro procedure deteriorates as we move towards the extreme quantiles. The entire process is repeated with $\tau_1 = 2$, which gave essentially the same conclusions.

This simulation study clearly demonstrate the superior performance of the optimal Robbins-Monro procedure over the Robbins-Monro procedure. The optimal Robbins-Monro procedure did a good job even for the estimation of extreme quantiles. The procedure appears to be robust against the model assumptions as it performed well for a wide range of distributions.

4. CONCLUSIONS

It is well known that the Robbins-Monro procedure is not useful for the estimation of extreme quantiles. One of the main reason for its failure is the use of the sequence $a_n = \{n\dot{M}(0)\}^{-1}$. This choice is based on asymptotic results and is obtained using a linear approximation to M , which is not good for a distribution function. We instead used a probit




Figure 1: Distributions used in the simulation study ($\alpha = 0.5$)

approximation and derived the optimal sequence. We also introduced another sequence $\{b_n\}$ into the procedure to improve the convergence. The optimal procedure is given in (10). A simulation study conducted using a wide range of distributions showed its superior performance over the Robbins-Monro procedure. The improvement obtained for the estimation of extreme quantiles is quite remarkable.

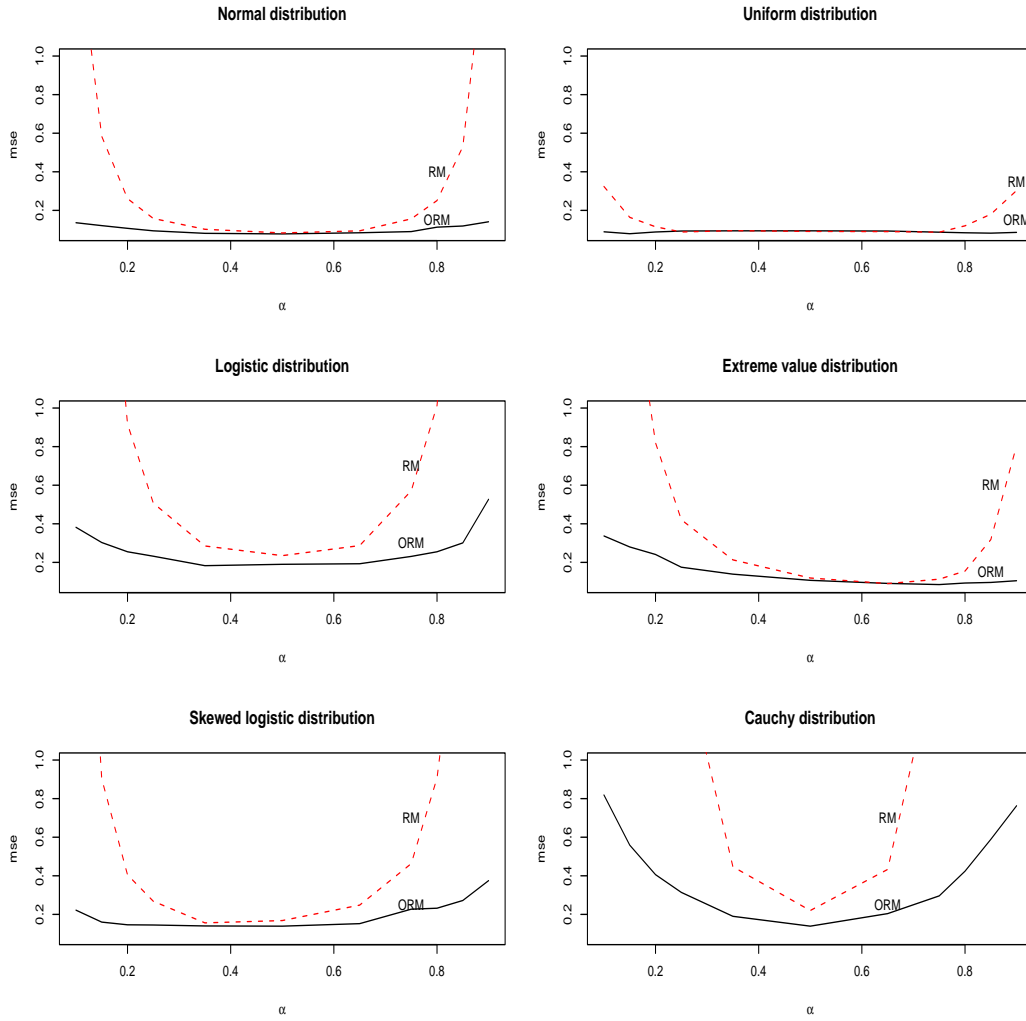


Figure 2: Mean square error against α . RM: Robbins-Monro, ORM: Optimal Robbins-Monro.

ACKNOWLEDGMENTS

I would like to thank two anonymous referees and Professor C. F. Jeff Wu for their valuable comments and suggestions.

REFERENCES

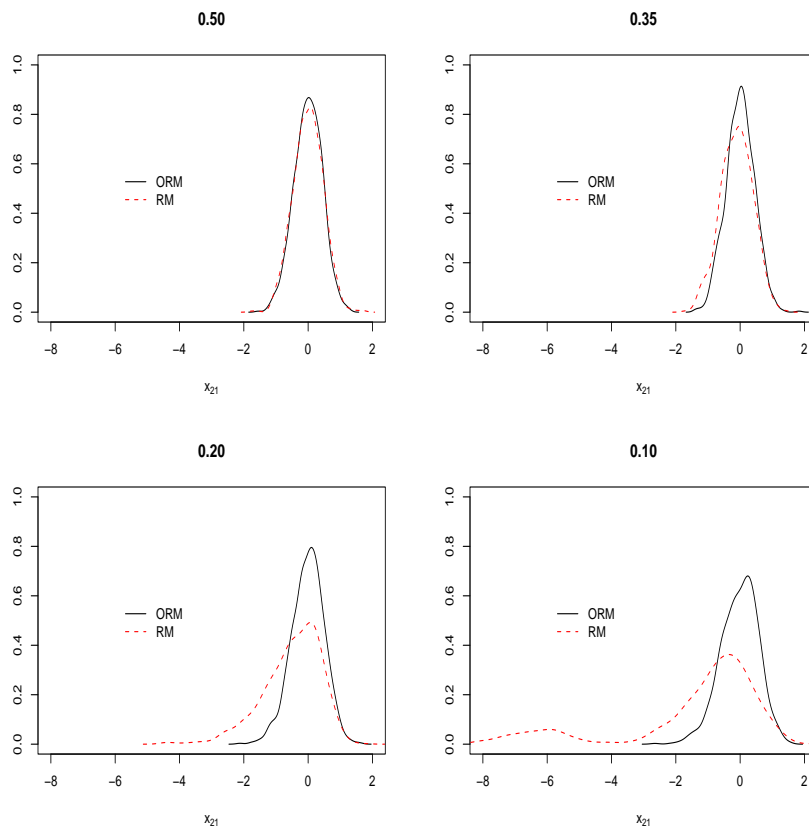


Figure 3: Density curves for different α with $M = \text{logistic distribution}$ and $\theta = 0$. RM: Robbins-Monro, ORM: Optimal Robbins-Monro.

Chung, K. L. (1954). On a Stochastic Approximation Method. *The Annals of Mathematical Statistics* **25**, 463-483.

Cochran, W. G. and Davis, M. (1965). The Robbins-Monro Method for Estimating the Median Lethal Dose. *Journal of Royal Statistical Society, Series-B* **27**, 28-44.

Hodges, J. L. and Lehmann, E. L. (1956). Two Approximations to the Robbins-Monro Process. *Proceedings of Third Berkeley Symposium* **1**, 39-55, University of California, Berkeley.

- Hu, I. (1997). Strong Consistency in Stochastic Regression Models via Posterior Covariance Matrices. *Biometrika* **84**, 744-749.
- Hu, I. (1998). On Sequential Designs in Non-linear Problems. *Biometrika* **85**, 496-503.
- Joseph, V. R. and Wu, C. F. J. (2002). Operating Window Experiments: A Novel Approach to Quality Improvement. *Journal of Quality Technology* **34**, 345-354.
- Kalish, L. A. (1990). Efficient Design for Estimation of Median Lethal Dose and Quantal Dose-Response Curves. *Biometrics* **46**, 737-748.
- Lai, T. L. and Robbins, H. (1979). "Adaptive Design and Stochastic Approximation," *The Annals of Statistics* **7**, 1196-1221.
- McLiesh, D. L. and Tosh, D. (1990). Sequential Designs in Bioassay. *Biometrics* **46**, 103-116.
- Neyer, B. T. (1994). D-Optimality-Based Sensitivity Test. *Technometrics* **36**, 61-70.
- Robbins, H. and Monro, S. (1951). A Stochastic Approximation Method. *The Annals of Mathematical Statistics* **29**, 373-405.
- Sacks, J. (1958). Asymptotic Distribution of Stochastic Approximation Procedures. *The Annals of Mathematical Statistics* **29**, 373-405.
- Sitter, R. R. and Wu, C. F. J. (1999). Two-Stage Design of Quantal Response Studies. *Biometrics* **55**, 396-402.

- Voelkel, J. G. (2003). Sequential Experimental Designs for Sensitivity Experiments. Revised for *Technometrics*.
- Wetherill, G. B. (1963). Sequential Estimation of Quantal Response Curves. *Journal of Royal Statistical Society- Series B* **25**, 1-48.
- Wetherill, G. B. and Glazebrook, K. D. (1986). Sequential Methods in Statistics. London: Chapman and Hall.
- Wu, C. F. J. (1985). Efficient Sequential Designs With Binary Data. *Journal of American Statistical Association* **80**, 974-984.
- Young, L. J. and Easterling, R. G. (1994). Estimation of Extreme Quantiles Based on Sensitivity Tests: A Comparative Study. *Technometrics* **36**, 48-60.