[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

and

[illegible]

ECDO R & D Report No. 1977-2

University of Minnesota
College of Education
Education Career Development Office
Minneapolis, Minnesota

April, 1977

(Presented at the 1977 AERA Annual Meeting: New York)

Abstract

This research investigates stochastic approximation procedures of the Robbins-Monro type. Following a brief introduction to sequential experimentation, attention is focused on formal methods for selecting successive values of a single independent variable. Empirical results obtained through computer simulation are used to compare several formal stochastic approximation techniques and "stopping rules". Marked differences were found between the five approximation procedures studied. One procedure using a "finite memory" had utility. Two procedures suggested in the literature were impractical under various conditions.

Introduction

Characteristically, an experiment involves a collection of treatment conditions (i.e. treatment levels or treatment combinations), a collection of experimental units, and an explicit plan for assigning treatment conditions to units. For purposes of exposition, we can divide experiments into those in which _time_ plays an important role and those in which it does not. _Time_ may enter into the experimental plan in several ways. For example, 1) at some point during a sequence of repeated measurements of the experimental unit a treatment condition may be introduced, as in trend analysis, 2) the experimental material may be successively exposed to several pre-specified treatment conditions and measured after each, as when assessment of order or residual effects is of interest, 3) treatment conditions may be administered to experimental units over time in such a way that previous treatment conditions and responses to them are used in determining the treatment conditions which follow. Note that in examples two and three treatment conditions are administered over time. But in the second example the exact treatment conditions are determined _a priori_, while in example three, they are determined during the experiment as a function of accumulating data. For convenience, we label the three examples as instances of _repeated measurement_, _serial_, and _sequential_ designs, respectively.

The present research is concerned with sequential experimentation. Experimental designs which are sequential in nature require that the experimenter consider both how the ensuing treatment conditions will be changed or adjusted and how the process will be discontinued, i.e. a "stopping rule". Sequential experiments can be differentiated from one another by considering whether a formal or informal procedure is used when adjusting treatment conditions, whether the stopping rule is formal or informal, whether or not more than one factor is used (i.e. multifactor experiment employing several different treatments), whether the independent variable or dependent variable is continuous or discrete, and by considering the purpose of the procedures (e.g. locating maxima). (For a general review and bibliography of recent work on experimental design, including the topics dealt with here, see Herzberg and Cox (1969). For a current review of the design of sequential experiments, see Chernoff (1975). Wetherill (1975) provides a useful introduction to the subject of our paper).)

Examples of applications of sequential designs are not plentiful in the educational research literature. Meyer (1963) presents an application of response surface methodology. This methodology is seen as sequential in nature by Chernoff and by Wetherill. Response surface designs are factorial in nature, employing several quantitative independent variables. The dependent variable is often assumed continuous and a polynomial function of the independent variables. Purposes of these designs include locating maxima or estimating parameters of the polynomial. Decision rules which specify the "design points" to use in the next stage and when to stop the process tend to be informal.

In contrast to response surface methodology are stochastic approximation techniques in which a single continuous independent variable is investigated and where values of chat independent variable are determined formally as a function both of the preceding values and the responses that were obtained when they were administered. A technique due to Robbins and Monro (1951) is an example. Its purpose is to find that value of the independent variable, say θ, such that the expected value of the dependent variable given θ is equal to some predetermined constant.

Our research investigates two examples of the Robbins-Monro process and three variant procedures which were motivated by it. Much of the previous research in this area has been focused on asymptotic properties. Chernoff (1975) gives a brief and readable review of this work. Of particular interest here is a paper by Hodges and Lehmann (1956) because it suggests assuming a linear relationship between the independent and dependent variables and also assumes that the slope parameter is known. While these two conditions would seldom be met in practice, their theoretical and numerical results provide a basis of comparison for empirical findings.

The Robbins-Monro procedure has been modified by some researchers so that two values of the independent variable are employed at each step (e.g. see Venter (1967)). This procedure has certain advantages, but only the case in which a single value of the independent variable is used at each step is studied here.

Presently, we know of no application of Robbins-Monro procedures in an educational experiment. However, the technique has been applied to a measurement problem by Lord (1971a, 1971b). Those two papers dealt with quantal responses, a subject not dealt with here. (For this reason and because we did not want to define the values of the independent variable *a priori* we have not considered the "Up and Down" method of stochastic approximation.) The present area of investigation has aimilarities with sequential estimation, but also aome important differences. For the estimation problem only the "stopping rule" need be considered, for no independent variable is manipulated.

The Problem

Assume that the experimenter's goal is that the value of a particular population mean is to be changed from its present value, δ, to a different value, α. For example, a population of adults may on the average score $\delta = 100$ on a particular standardized reading test and the goal is to increase that average to $\alpha = 116$. The experimenter has in mind a treatment variable (say, number of hours of individual tutoring) which he knows can affect the average reading score, but the exact nature of the relationship between reading score and tutoring is unknown. In other words, the "end" is known, but not the specific "means", and therefore, the appropriate value of the independent variable, or treatment condition, must be found. More formally, the expected value of the reading score is a function of the independent variable, $E(y) = f(x)$, and the experimenter wishes to determine that specific value of the independent variable, $x = \theta$, for which $E(y) = \alpha$, or $E(y(x = \theta)) = \alpha$. For present purposes it is assumed that if $x > \theta$ then $E(y(x)) > \alpha$, and if $x < \theta$ then $E(y(x)) < \alpha$. Given this situation the experimenter can select an initial value x_1 and thereafter choose the value of the independent variable as $x_{n+1} = x_n - a_n(y_n(x_n) - \alpha)$. The a_n are selected to have several characteristics, the most intuitively important of which is that $a_n \rightarrow 0$ as $n \rightarrow \infty$ "at a suitable rate". One possible definition is $a_n = \frac{1}{n}$. If appropriate a_n are chosen, such as $\frac{1}{n}$, Robbins and Monro (1951) proved $n \rightarrow \infty$, $x_n \rightarrow \theta$. The experimenter, of course, must have some feel for the speed of convergence, and how this convergence is affected by the choice of x_1, the relationship between $E(y(x))$ and x, and the density of $y(x)$. He also must have some idea of when to stop the experimentation. Most of the results in the literature to date, however, are asymptotic in nature, with relatively little work being done on stopping rules (Chernoff (1975) offers no citations, but see Farrell (1962)). The literature, as it appears to us, provides little if any practical guidance for the experimenter.

Methods

Initial results were obtained with an interactive empirical approach using computer simulation techniques on a time-shared CDC Cyber 74. Many computer runs were made as the researchers sought to understand the importance of the numerous parameters which can be considered. Following this first phase of computer runs, during which all the values produced from a single sequential experiment were often observed, more traditional Monte Carlo experiments were performed, replicating the experiments a number of times to obtain estimates of how the procedures operate "in the long run". In summary, the approach used combined both an interactive search during which the researchers observed the behavior of various functional relationships during a single replication and more traditional "fixed" type of experiments in which a number of replications of an experimental situation were made to obtain stable estimators.

All pseudo random numbers were obtained from either NORMAL or RAN3F which are a normal (N) random number generator and a uniform (U) random number generator, respectively. One thousand random numbers were generated per each call of these routines and following generation they were immediately permuted by an independent randomization procedure using the program PERMUTE. All routines are maintained by the University of Minnesota Computer Center.

Design

The model for the random variable y was $y = \mu_{y.x} + \varepsilon$, where $\mu_{y.x} = E(y(x)) = f(x) = \beta_0 x^0 + \beta_1 x^1 + \beta_2 x^2 + \beta_3 x^3$ and where ε is independently identically distributed either as $N(0, \sigma^2_{y.x})$ or $U(0, \sigma^2_{y.x})$. Following the interactive search in which many parametric specifications and stopping rules were studied, certain choices of parameters and rules were made for the more standard type of Monte Carlo investigation. These included:

1. Four definitions of a_n. They were $a_n = \frac{1}{n}, \frac{1}{n\beta'_\theta}, \frac{1}{n\hat{\beta}}, \frac{s^k}{n}$, where β'_θ is the first derivative of $f(x)$ evaluated at θ, $\hat{\beta}$ is the usual slope estimator, and $s^k = |\sum_{i=j}^{n} z_i|^k$, $k = \text{INT}[s/\frac{1}{2}c]$, where $c > 0$ and even, $j = \max(1, n - c + 1)$ and $z_n = 1$ if $y(x_n) < \alpha$, or $z_n = -1$ if $y(x_n) > \alpha$. (INT means "integer part of.") For the procedure employing $\hat{\beta}$, $a_n = \frac{1}{n}$ for $n < 20$ and $\frac{1}{n\hat{\beta}}$ otherwise. In the definition $a_n = \frac{s^k}{n}$ a "finite memory" is introduced into the approximation process, and successively positive or negative values of $y_n(x_n) - \alpha$ cause larger adjustments x_{n+1} than is the case with the other definitions. Both s^k and $\hat{\beta}$ are random variables and this results in a variant of the Robbins-Monro procedure in that it assumes the a_n to be "a fixed sequence of positive constants."

2. Two stopping rules. They were

R1: Stop if $n \geq 20$ and if α contained in $\hat{\mu}_{y.x} \pm \hat{\sigma}_{y.x} \left(\frac{1}{n} + \frac{(x_n - \bar{x}_n)^2}{(\Sigma(x_i - \bar{x}_n)^2} \right)^{\frac{1}{2}} t_{p/2}$

(where $\hat{\mu}_{y.x} = \hat{\delta} + \hat{\beta}x_n$) or $n = 200$.

R2: Stop if $n \geq 20$ and, considering the last 20 values of z, if $\Sigma z = 9$, 10, or 11, and the number of "runs" is 9, 10, 11, 12, or 13, or if $n = 200$.

3. Three sets of β_0, β_1, β_2, β_3. They were [100, .14142, 0, 0] [100, .34641, 0, 0], and [100, .12686, .0058512, -.000023767].

4. Two conditional variances. They were $\sigma^2_{y.x} = 100$ and 25.

Most "final" experiments were based on 500 replications.[1] Based on these replications, the mean and variance were computed for $(x_n - \theta)$ at $n = 30$, 50, 100 steps and for both rules, R1 and R2. Additionally, for both rules the mean and variance of the number of steps needed to stop were also computed.

$a_n = \frac{1}{n}$ was included in the experiment because it was suggested in Robbins and Monro's original paper. Hodges and Lehmann provide results on $a_n = \frac{1}{n\beta_1}$ when the regression is in fact linear, and it has certain optimal characteristics and therefore was included as a basis for comparisons. In discussing the preceding work, Chernoff (1975) remarked that "In the stochastic approximation case using sequences $a_n = \frac{c}{n}$, there is no prior knowledge of θ to insure that $c = \beta^{-1}$. However, as data accumulate one would hopefully obtain a satisfactory estimate of β providing the successive x_n are not too close to each other (p. 70)."[2] We interpreted these comments to mean that when one has "sufficient" information one would estimate β_1 using the least squares estimator, $\hat{\beta}$. Initial results demonstrated that the instability of $\hat{\beta}$ for small n caused erratic adjustments and poor convergence. This lead to the procedure $a_n = \frac{1}{n}$ for n = 1, . . ., 19 and $\frac{1}{n\hat{\beta}}$ thereafter. The $a_n = \frac{s^k}{n}$ procedure was developed during the interactive part of the present research. It seemed reasonable to specify an adjustment procedure which would make larger adjustments if $E(y_n(x_n)) - \alpha$ were judged to be large. Considering only the sign of $y_n(x_n) - \alpha$ and taking c = 4, for the patterns (+ + + +) or (- - - -), $s^k = 16$, for patterns like (+ - + +) or (- - - +) $s^k = 2$, and for patterns with two pluses and two minuses, $s^k = 1$. This type of adjustment assumes that the error distributions are symmetric so that the probability of a plus at $x_n = \theta$ is $\frac{1}{2}$. During the interactive phases of this research, c = 4 and c = 10 were found to work well.

Stopping rule R1 employs the standard confidence interval for estimating $\mu_{y.x}$. This seemed a reasonable approach to consider, especially when $a_n = \frac{1}{n\hat{\beta}}$ is employed. The confidence coefficient, p, used was .60. This value was chosen during the interactive phase on the basis of performance.

Stopping rule R2 comes from reasoning similar to that used in developing the s^k procedure. At $x_n = \theta$, for symmetric error distributions the sign of $y_n(x_n) - n$ would be independently distributed as a Bernoulli variable with parameter $\frac{1}{2}$. R2 essentially tests two hypotheses, one concerning "randomness" and the other that the proportion of "pluses" is $\frac{1}{2}$.

Results

Results of the runs are presented in Tables 1-8. The mean and variance of the bias $(x_n - \theta)$, are reported for the conditions studied as well as the mean and variance of the "number of steps to stop" for R1 and R2. The average squares bias, $(x_n - \theta)^2$ is not reported, but it can be easily obtained by squaring the average bias and adding this to the variance of the bias (i.e. $E(x_n - \theta)^2 = V(x_n - \theta) + (E(x_n - \theta))^2$).

The method employing $\frac{1}{n\beta_a'}$ was superior to the other methods, but since β_a' would seldom be known, results associated with β_a' will be of greatest utility as "benchmarks". It is clear from the results that generally, the procedure $a_n = \frac{1}{n}$ had the poorest performance. In situations where little or no prior information is available about the relationship between the independent and dependent variables, the results would lead us to use $a_n = \frac{s^k}{n}$ with c = 4. This typically does as well as or better than the other procedures not employing β_a'. It is markedly better when there is a weak relationship between the independent and dependent variables and a poor start is made (see columns 1 and 2 of Tables 1-7). In an attempt to determine the behavior of $a_n = \frac{s^k}{n}$ when a "good start" is made, the experiments reported in Table 8 with $x_1 = \theta$ were carried out. We believe the procedure did reasonably well under these circumstances.

If more information is available, one might profitably choose one of the procedures studied. Neither R1 nor R2 is uniformally better with respect to bias and number of steps to stop. There also appears to be at least some interaction with the definition of a_n, and this complicates matters in a few instances. Here we can only recommend that one make a best guess about conditions and use that stopping rule which would be best.

Educational Significance

One potential area of application for stochastic approximation is that of formative evaluation. Stochastic approximation can suggest values of the independent variable which would attain programmatic goals, and this information could be fed to persons directly involved in program development. Within the framework developed by Sanders and Cunningham (1974), stochastic approximation could provide "external information" for "formative interim evaluation activities". When a summative evaluation is planned, perhaps using one of the more standard experimental designs, design points can be chosen in the region suggested through sequential experimentation, thereby increasing the likelihood that the program will demonstrate its effectiveness.

In general, stochastic approximation would appear to be a useful technique in any area where individuals have a goal firmly in mind but lack sufficient knowledge of the independent variable to design an efficient, more traditional experiment. Education is goal oriented, and information about how to achieve a goal is often more important than, say, information about the exact nature of the relationship between an independent and dependent variable. Stochastic approximation can provide useful information about an independent variable, even when its defined over a broad range of values, while requiring relatively few subjects for its implementation.

Table 1

Mean and Variance of the Bias $(x_n - \theta)$ at 30 Steps Where $x_1 = 4$ †

a_n	$\sigma^2_{y.x}$	$[\beta_0,\beta_1,\beta_2,\beta_3]$: [100,.14142,0,0]		[100,.34641,0,0]		[100,.12686,.0058512, −.000023767]	
		$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$
$\frac{1}{n\beta'_\theta}$	100	-.49	.38	-.13	.20	-.96	-.79
		160.93	177.75	27.18	29.84	12.85	14.54
	25	-.19	.24	-.08	.10	-.53	-.48
		40.85	44.77	6.81	7.46	3.39	3.81
$\frac{1}{n}$	100	-61.16	-61.11	-9.73	-9.25	-6.57	-6.68
		72.00	81.11	27.65	32.02	20.82	22.91
	25	-61.27	-60.96	-9.46	-9.62	-6.48	-6.36
		18.89	20.48	7.69	7.62	4.74	5.59
$\frac{1}{n\hat{\beta}}$	100	-44.86	-41.86	-5.65	-4.55	-3.14	-2.36
		1436.76	1385.58	140.41	180.92	108.84	188.85
	25	-36.23	-35.41	-5.72	-5.57	-3.07	-3.10
		490.86	577.45	65.93	51.08	43.88	50.68
$\frac{s^k}{n}(c=4)$	100	-1.92	-3.58	.42	.02	-.18	-.19
		298.62	452.01	72.73	97.26	44.55	61.85
	25	-.88	-.98	.44	.19	.18	-.32
		71.44	96.15	18.46	22.93	10.77	14.73
$\frac{s^k}{n}(c=10)$	100	-10.85	-17.03	-.81	-1.93	-.06	-.89
		835.47	764.57	110.47	87.05	47.69	49.23
	25	.37	-1.98	-.67	-.75	-.10	.26
		378.17	419.80	24.05	25.86	24.08	11.04

† For Tables 1–8 the upper number in each cell is the mean and the lower number is the variance.

Table 2

Mean and Variance of the Bias $(x_n - \theta)$ at 50 Stepc Where $x_1 = 4$

		$[\beta_0, \beta_1, \beta_2, \beta_3]$					
		[100,.14142,0,0]		[100,.34641,0,0]		[100,.12686,.0058512, −.000023767]	
a_n	$\sigma^2_{y.x}$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$
$\frac{1}{n\beta'_\theta}$	100	.30	.34	.17	.16	−.43	−.58
		93.84	109.69	15.75	18.35	7.21	8.67
	25	.18	.20	.08	.08	−.28	−.25
		23.66	27.54	3.94	4.59	1.96	2.25
$\frac{1}{n}$	100	−56.87	−56.83	−8.02	−7.69	−5.04	−5.17
		62.61	69.82	20.59	23.88	13.69	15.49
	25	−56.92	−56.61	−7.92	−8.04	−4.98	−4.88
		16.49	17.82	5.44	5.88	3.12	3.67
$\frac{1}{n\hat{\beta}}$	100	−20.43	−22.29	−1.38	−1.51	−.63	−.67
		1036.14	994.74	158.51	104.35	33.50	67.36
	25	−20.56	−20.91	−2.97	−3.20	−1.47	−1.43
		231.29	236.72	25.52	18.06	12.84	13.58
$\frac{s^k}{n}$(c=4)	100	−1.48	−1.99	.07	−.06	.26	−.34
		157.02	262.19	39.42	50.39	25.33	33.99
	25	−.46	−.47	.27	.11	.04	.07
		40.04	58.47	9.83	13.56	6.11	8.22
$\frac{s^k}{n}$(c=10)	100	−7.93	−12.07	−.38	−1.36	−.12	−.31
		476.52	513.41	54.41	46.98	32.12	22.28
	25	.33	−1.81	−.10	−.02	.20	.16
		169.92	229.20	12.92	12.72	6.35	6.72

Table 3

Mean and Variance of the Bias $(x_n - \theta)$ at 100 Steps Where $x_1 = 4$

		$[\beta_0, \beta_1, \beta_2, \beta_3]$					
		[100,.14142,0,0]		[100,.34641,0,0]		[100,.12686,.0058512, -.000023767]	
a_n	$\sigma^2_{y.x}$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$
$\frac{1}{n\beta'_\theta}$	100	.35	.00	.16	.01	-.16	-.27
		48.75	58.39	8.14	9.75	3.56	4.65
	25	.19	.01	.08	.00	-.11	-.16
		12.20	14.63	2.03	2.44	.95	1.08
$\frac{1}{n}$	100	-51.56	-51.46	-6.31	-6.05	-3.53	-3.63
		53.31	57.95	13.19	15.68	7.21	8.42
	25	-51.56	-51.36	-6.20	-6.28	-3.48	-3.42
		13.64	14.56	3.50	3.83	1.71	2.07
$\frac{1}{n\hat{\beta}}$	100	-8.62	-9.93	-.43	-.40	-.16	-.31
		333.15	308.16	44.44	34.35	8.90	18.12
	25	-10.44	-10.67	-1.39	-1.53	-.50	-.58
		74.54	71.06	7.19	5.60	3.08	3.30
$\frac{s^k}{n}(c=4)$	100	-.28	-1.03	.25	.15	.09	.02
		78.34	115.32	21.10	27.90	11.89	15.02
	25	-.11	-.27	.05	-.04	-.05	.02
		19.31	28.79	5.08	6.84	3.10	4.06
$\frac{s^k}{n}(c=10)$	100	-5.75	-7.97	-.35	-.61	.04	-.29
		227.50	299.75	19.40	35.03	11.50	14.37
	25	-.26	-1.61	.00	-.06	.12	-.01
		71.40	113.45	5.18	8.52	2.98	3.15

Table 4

Mean and Variance of the Bias ($\bar{x}_n - \theta$) When

Stopped with Parametric Rule (R1) Where $x_1 = 4$

		$[\beta_0, \beta_1, \beta_2, \beta_3]$					
		[100,.14142,0,0]		[100,.34641,0,0]		[100,.12686,.0058512, -.000023767]	
s_n	$\sigma^2_{y.x}$	$\varepsilon\sim N(0,\sigma^2_{y.x})$	$\varepsilon\sim U(0,\sigma^2_{y.x})$	$\varepsilon\sim N(0,\sigma^2_{y.x})$	$\varepsilon\sim U(0,\sigma^2_{y.x})$	$\varepsilon\sim N(0,\sigma^2_{y.x})$	$\varepsilon\sim U(0,\sigma^2_{y.x})$
$\frac{1}{n\beta'_\theta}$	100	-.70	.66	-.17	.33	-1.07	-1.09
		237.97	265.39	40.34	44.80	18.41	21.33
	25	-.25	.40	-.10	.16	-.49	-.55
		60.46	67.14	10.08	11.18	4.75	5.50
$\frac{1}{n}$	100	-47.45 *	-46.62 *	-7.82	-7.49	-4.42	-4.54
		61.09	49.33	23.91	24.05	13.57	14.73
	25	-46.58 *	-46.76 *	-4.94 *	-5.04 *	-2.60	-2.54 *
		11.45	12.53	2.87	2.70	1.38	1.26
$\frac{1}{n\hat{\beta}}$	100	-14.39	-12.61	-5.28	-5.11	-2.70	-2.56
		924.61	1000.61	78.38	96.76	44.00	47.93
	25	-5.16	-5.46	-2.04	-1.92	-.68	-.74
		112.84	71.87	13.92	12.36	5.93	3.91
$\frac{s^k}{n}(c=4)$	100	-2.36	-4.32	.06	.41	-.80	-1.15
		352.92	460.81	59.17	67.44	28.71	32.94
	25	-.89	-1.20	.10	.22	-.32	-.65
		91.96	110.92	14.93	15.90	7.28	7.61
$\frac{s^k}{n}(c=10)$	100	-13.25	-14.75	-1.91	-2.41	-.42	-.71
		498.27	502.03	77.19	69.56	32.00	33.87
	25	-1.88	-1.79	-.48	-.64	.42	.39
		129.16	114.31	14.53	17.57	8.27	8.20

*The estimates in these cells are based on 100 replications instead of 500.

Table 5

Mean and Variance of the Number of Steps

When Stopped with Parametric Rule (R1) Where $x_1 = 4$

		$[\beta_0,\beta_1,\beta_2,\beta_3]$					
		[100,.14142,0,0]		[100,.34641,0,0]		[100,.12686,.0058512, -.000023767]	
a_n	$\sigma^2_{y.x}$	$\epsilon\sim N(0,\sigma^2_{y.x})$	$\epsilon\sim U(0,\sigma^2_{y.x})$	$\epsilon\sim N(0,\sigma^2_{y.x})$	$\epsilon\sim U(0,\sigma^2_{y.x})$	$\epsilon\sim N(0,\sigma^2_{y.x})$	$\epsilon\sim U(0,\sigma^2_{y.x})$
$\frac{1}{n\beta'_\theta}$	100	20.18	20.21	20.13	20.18	23.13	22.23
		2.37	1.79	1.30	1.54	338.95	194.61
	25	20.17	20.13	20.18	20.14	26.28	24.96
		2.07	1.08	2.33	1.09	622.57	464.60
$\frac{1}{n}$	100	196.44 *	198.24 *	87.64	85.30	92.26	95.61
		627.30	309.76	7087.82	7003.51	7320.44	7372.14
	25	200.00 *	200.00 *	191.49 *	196.05 *	191.22	194.62 *
		0.00	0.00	1393.06	572.69	1416.44	622.76
$\frac{1}{n\hat{\beta}}$	100	89.66	93.18	28.76	30.11	28.78	29.03
		5019.12	5523.06	788.86	1132.37	695.83	716.53
	25	177.87	178.80	76.85	82.71	78.21	80.39
		2849.49	2840.74	2214.93	2414.49	1589.28	1468.92
$\frac{a^k}{n}$(c=4)	100	24.86	27.53	25.68	26.46	26.02	27.16
		338.35	610.14	210.89	262.41	165.28	228.26
	25	25.03	27.37	25.50	26.75	25.87	28.97
		477.47	711.50	155.27	213.57	137.34	405.82
$\frac{s^k}{n}$(c=10)	100	53.20	65.79	24.32	25.61	23.96	26.59
		3807.48	5113.46	399.95	492.67	250.35	654.13
	25	83.42	111.00	27.24	29.36	27.01	26.12
		6015.15	6715.59	775.31	1136.49	777.34	646.42

*The estimates in these cells are based on 100 replications instead of 500.

Table 6

Mean and Variance of the Bias $(x_n - \theta)$ When

Stopped with Nonparametric Rule (R2) Where $x_1 = 4$

		$[\beta_0,\beta_1,\beta_2,\beta_3]$					
		[100,.14142,0,0]		[100,.34641,0,0]		[100,.12686,.0058512, -.000023767]	
a_n	$\sigma^2_{y.x}$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$	$\varepsilon \sim N(0,\sigma^2_{y.x})$	$\varepsilon \sim U(0,\sigma^2_{y.x})$
$\frac{1}{n\beta'_\theta}$	100	-.21	.05	-.01	.07	-.94	-.80
		173.52	179.81	29.40	29.87	12.59	14.97
	25	.00	.08	.00	.03	-.44	-.54
		44.00	44.88	7.34	7.48	3.42	4.09
$\frac{1}{n}$	100	-48.05	-50.04	-7.75	-8.02	-4.68	-5.52
		42.17	51.09	16.73	25.63	11.06	17.34
	25	-46.74	-46.54	-6.17	-7.01	-3.66	-4.10
		11.18	12.08	3.10	4.97	1.75	2.90
$\frac{1}{n\hat{\beta}}$	100	-9.77	-14.57	-2.86	-3.25	-1.55	-2.16
		526.20	631.61	54.60	84.38	18.48	30.35
	25	-9.70	-11.74	-2.53	-3.42	-1.48	-1.83
		56.80	82.45	8.40	11.26	4.10	5.39
$\frac{s^k}{n}(c=4)$	100	-.46	-2.30	.53	-.20	-.47	-.60
		249.88	453.42	77.44	106.10	46.65	65.27
	25	-.04	-.50	.49	.07	.12	-.11
		61.75	92.97	16.49	25.41	10.77	14.87
$\frac{s^k}{n}(c=10)$	100	-8.11	-12.36	-.34	-1.73	.29	-.25
		391.91	502.21	42.68	64.64	30.43	33.71
	25	.49	-1.48	.00	-.12	.61	.52
		138.40	215.96	10.91	16.19	6.55	7.37

Table 7

Mean and Variance of the Number of Steps

When Stopped with Nonparametric Rule (R2) Where $x_1 = 4$

		$[\beta_0,\beta_1,\beta_2,\beta_3]$					
		[100,.14142,0,0]		[100,.34641,0,0]		[100,.12686,.0058512, -.000023767]	
a_n	$\sigma^2_{y.x}$	$\epsilon\sim N(0,\sigma^2_{y.x})$	$\epsilon\sim U(0,\sigma^2_{y.x})$	$\epsilon\sim N(0,\sigma^2_{y.x})$	$\epsilon\sim U(0,\sigma^2_{y.x})$	$\epsilon\sim N(0,\sigma^2_{y.x})$	$\epsilon\sim U(0,\sigma^2_{y.x})$
$\frac{1}{n\beta'_\theta}$	100	32.42	31.79	32.40	31.86	35.06	33.45
		226.32	187.77	223.18	187.39	294.07	247.42
	25	32.60	31.98	32.60	31.98	34.81	33.31
		224.71	187.41	224.71	187.41	240.51	237.54
$\frac{1}{n}$	100	173.71	136.65	61.13	50.38	60.90	48.90
		2234.70	3301.21	1369.02	917.82	1359.13	756.51
	25	200.00	199.66	112.97	84.67	101.78	79.81
		0.00	34.43	2832.98	1962.52	2279.99	1652.54
$\frac{1}{n\hat{\beta}}$	100	84.30	70.84	51.21	43.98	47.11	44.06
		1188.04	999.99	1039.26	681.22	596.34	642.90
	25	110.54	94.07	66.81	57.02	59.28	53.61
		1751.84	1258.61	718.70	614.92	536.85	471.00
$\frac{s^k}{n}(c=4)$	100	39.48	36.76	31.71	30.31	30.70	30.02
		365.17	341.61	203.34	156.66	191.41	172.76
	25	36.85	36.51	35.16	31.93	31.08	30.98
		239.20	297.97	310.12	171.47	140.91	162.65
$\frac{s^k}{n}(c=10)$	100	58.92	54.05	39.62	36.75	35.52	35.71
		994.45	745.88	431.32	358.61	307.83	340.59
	25	64.06	59.23	39.79	37.49	37.07	34.87
		1192.78	953.26	407.40	324.22	364.19	315.87

Table 8

Means and Variances Where $x_1=\theta$ and $\epsilon \sim N(0,\sigma^2_{y.x}=100)$

$[\beta_0,\beta_1,\beta_2,\beta_3]$	a_n	Bias $(x_n-\theta)$ at 30 steps	Bias $(x_n-\theta)$ at 50 steps	Bias $(x_n-\theta)$ at 100 steps	Bias $(x_n-\theta)$ when stopped with R1	No. of steps when stopped with R1	Bias $(x_n-\theta)$ when stopped with R2	No. of steps when stopped with R2
[100,.14142,0,0]	$\frac{1}{n\beta'_\theta}$	-.49	.30	.35	-.88	20.33	-.22	32.59
		160.93	93.84	48.75	235.51	1.11	172.67	222.70
	$\frac{s^k}{n}(c=4)$	.30	.10	.39	-.03	23.82	-.25	33.11
		234.71	148.69	77.78	243.27	80.88	232.17	289.16
[100,.34641,0,0]	$\frac{1}{n\beta'_\theta}$	-.13	.17	.16	-.26	20.30	-.01	32.57
		27.18	15.75	8.14	39.86	1.07	29.24	219.56
	$\frac{s^k}{n}(c=4)$	.25	.17	.08	-.06	25.45	.66	30.56
		72.55	36.14	19.10	63.29	176.93	77.57	203.33

References

Chernoff, H. Approaches in sequential design of experiments. in J. N. Srivastava (Ed.), A survey of statistical design and linear models. Amsterdam: North-Holland Publishing, 1975.

Farrell, R. H. Bounded length confidence intervals for the zero of a regression function. Annals of Mathematical Statistics, 1962, 33, 237-247.

Herzberg, A. M. & Cox, D. R. Recent work on the design of experiments: A bibliography and a review. Journal of the Royal Statistical Society, 1969, Series A, 132, 29-67.

Hodges, J. L. & Lehmann, E. L. Two approximations to the Robbins-Monro process. In J. Neyman (Ed.), Proceedings of the Third Berkley Symposium on Mathematical Statistics and Probability. Berkeley, California: University of California Press, 1956.

Lord, F. M. Robbins-Monro procedures for tailored testing. Educational and Psychological Measurement, 1971, 31, 3-31.

Lord, F. M. Tailored testing, An application of stochastic approximation. Journal of American Statistical Association, 1971, 66, 707-711.

Meyer, D. L. Response surface methodology in education and psychology. The Journal of Experimental Education, 1963, 31, 329-336.

Robbins, H. & Monro, S. A stochastic approximation method. Annals of Mathematical Statistics, 1951, 22, 400-407.

Sanders, J. R. & Cunningham, D. J. Formative evaluation: Selecting techniques and procedures. In G. D. Borich (Ed.), Evaluating educational programs and products. Englewood Cliff, New Jersey, Educational Technology Publications, 1974.

Venter, J. H. An extension of the Robbins-Monro procedure. Annals of Mathematical Statistics, 1967, 38, 181-190.

Wetherill, G. B. Sequential methods in statistics. London: Chapman and Hall, 1975.

Footnotes

[1]The original plan was to have 500 replications for each set of conditions, however, given the value of x_1 used here, $a_n = \frac{1}{n}$ converged slowly and for some conditions the rule "stop if $n = 200$" was used for virtually every replication. We decided to use only 100 replications in these instances, and those runs are noted in the tables.

[2]The "c" in this quote is not defined in the same way as the "c" in the definition of $a_n = \frac{c^k}{n}$.