MONTE CARLO OPTIMIZATION BY STOCHASTIC APPROXIMATION (WITH APPLICATION TO HARVESTING OF ATLANTIC MENHADEN)

David Ruppert
Department of Statistics
University of North Carolina
Chapel Hill, North Carolina 27514

Rodney L. Reish
North Carolina Baptist Hospital
Winston-Salem, North Carolina 27103

Richard B. Deriso
International Pacific Halibut Commission
Seattle, Washington 98105

Raymond J. Carroll
Department of Statistics
University of North Carolina
Chapel Hill, North Carolina 27514

Key words: Robbins-Monro process, stochastic approximation, Monte Carlo, variance reduction, stopping rule, optimal harvesting.

This research was supported by contract from the National Marine Fisheries Service, Southeast Fisheries Center, Beaufort, North Carolina 28516. Also, the work of David Ruppert was partially supported by National Science Foundation Grant MCS 8100748, and the work of Raymond Carroll was supported in part by Air Force Office of Scientific Research Grant AFOSR 80-0080.

ABSTRACT

Stochastic approximation can be viewed as a methodology for designing a sequence of response surface experiments. Although stochastic approximation has not been employed widely by statisticians, several authors agree that it is potentially very useful for a variety of statistical problems. In a recent study of the Atlantic menhaden, a commercially important fish in the herring family (Clupeidae), we made extended use of stochastic approximation and were quite pleased with the results. This paper is intended to introduce stochastic approximation to those statisticians unfamiliar with the area. A stochastic simulation model of the menhaden population is used as an example, but the paper is not addressed to only those working in fisheries. In this model, two variables are used to define the harvesting policy. For any values of these variables, the model will produce a random catch, and for a specified utility function the objective is to find the values of the variables which maximize the expected utility of the catch. Therefore, this is a classical response surface problem. However, nonsequential response surface methods would be extremely expensive to apply here. We used stochastic approximation to estimate the policy maximizing the expected utility of the catch.

Much of the paper duscusses the application of known results, but there are also some new results. In particular, we show how the use of common random numbers, a standard variance reduction technique, can be applied to stochastic approximation.

1. INTRODUCTION

Atlantic menhaden are migratory fish caught from Florida to New England mainly by purse-seining. Menhaden are not palatable to humans, but they are processed into two important products, fish oil and fish meal. The oil is used in the manufacture of paints, lubricants, and cosmetics, and the meal is a high protein ingredient in animal feed. Together Atlantic and Gulf menhaden support the largest U.S. commercial fishery measured in pounds landed, and excluding shellfish, menhaden catches rank next to those of tuna and salmon in monetary value. At present the Atlantic menhaden fishery is subjected to some intrastate regulation, but it is not managed on a fishery-wide basis. Catches of menhaden experienced a rapid decline after 1962 and have only partially recovered, but some biologists feel that proper management could increase future catches substantially. To explore the potential effects of management, we developed a simulation model which incorporated several key features of menhaden population dynamics:

a) an age structured population

b) a stochastic spawner-recruit relationship

c) density dependent juvenile growth

d) size dependent egg production.

There are many analytical treatments of optimal harvesting for simpler population models, and Mendelssohn (1978), Reed (1980), and Ludwig and Walters (1981) serve as recent examples. Optimal harvesting strategies for models of the present complexity, however, cannot be fully studied by available analytic techniques. Therefore, we resorted by Monte Carlo simulation.

In this paper, we discuss our optimization methodology, which uses stochastic approximation and which is applicable not only to other fisheries models but also to a wide range of other simulation models. Our model will be

discussed here, but only in so far as it illustrates the stochastic approximation methodology. Further details of the model and the simulation results will be reported elsewhere.

In a pioneering paper on response surface methodology, Box and Wilson (1951) discussed the use of a series of experiments to locate the maximum of a regression function. Each experiment in the series is used to estimate the gradient of the function at the center of the design, and the next design center is chosen by moving along the estimated gradient direction (steepest ascent). The literature inspired by this paper has been mostly concerned with the design of the experiment at a single stage, not the sequential aspects of response surface methodology.

Not long after Box and Wilson's paper appeared, Kiefer and Wolfowitz (1952), motivated by the work by Robbins and Monro (1951) on stochastic approximation, introduced a somewhat different approach to maximization of a regression function. They considered only one independent variable and only simple, two point designs, and concentrated their study on how the sequence of design centers should be chosen. Clearly Box and Wilson and Kiefer and Wolfowitz were studying different sides of the same problem. Unfortunately, these two lines of research have been diverging ever since.

The work of Box and his collaborators emphasizes bias and variance properties of fixed designs, and much of this research has been applied in a nonsequential manner. There has been little or no asymptotic theory developed by these workers, but a great deal of attention has been focused upon practical applications.

The paper of Kiefer and Wolfowitz attracted the attention of many mathematicians who developed an extensive large sample theory, but the highly technical nature of these papers and the almost complete lack of finite sample results apparently has discouraged practitioners from using the

Kiefer-Wolfowitz method. The stochastic approximation literature has recently become even more oriented towards probability theory rather than practical application, since current work concerns questions such as invariance principles and laws of the iterated logarithm; cf. Ruppert (1982).

Mead and Pike (1975), in their extensive review of response surface methodology, saw stochastic approximation as a part of this methodology and they lamented the lack of attention it has received by applied workers. They also felt the need for papers reporting the use of stochastic approximation in concrete situations. We hope that this paper can, to some extent, fill this need.

In our work we did not use the Kiefer-Wolfowitz procedure, as would be expected for an optimization problem, but rather we applied the similar Robbins-Monro procedure directly to the gradient of the function to be maximized. The use of the Robbins-Monro method was facilitated by the variance-reduction technique of common random numbers and resulted in a considerable increase in efficiency.

2. THE MODEL

In our study, we decided to focus upon the total value or utility of menhaden harvested over a 25 year period. Since, for various economic or social reasons, the utility need not be proportional to biomass, we defined "total utility" as follows. Let C_i be the catch in biomass of menhaden in the ith year. Then C_i depends upon the harvesting policy being implemented and the recruitment into the fishery during years 1 to i. As we discuss below, our data show recruitment to be extremely variable, so to insure realism, a stochastic model of recruitment is necessary. Thus, C_i is a random variable. Let U be a function, called the utility function, which measures the value of a given catch for a year. Then the value for the ith year is $U(C_i)$ and the total utility during a 25 year period is

$$\sum_{i=1}^{25} U(C_i) \; . \tag{2.1}$$

The harvesting policies which we studied depend upon one or more control variables, e.g. for constant fishing mortality rate policies the constant value of the fishing mortality rate is the control variable. Thus, the expected value of (2.1) as a function of the control variables is a response surface. Our major concern was finding the values of the control variables which maximize the response surface.

We used two utility functions, $U(x) = x$, where utility is proportional to biomass, and $U(x) = \log(x + 1)$, which we hope is representative of the class of utility functions with decreasing marginal utility. Although utility functions play a central role in decision analysis, e.g. Raiffa (1968), we are unaware of previous introductions of utility functions into fisheries analysis.

The choice of a 25 year horizon was somewhat arbitrary. We wanted a horizon length which did not overemphasize the initial, transient behavior of the

population as it approached the equilibrium distribution for the harvesting strategy in effect.

Menhaden reach sexual maturity by age 2 or 3, and menhaden older than five years are extremely rare, so we felt that the equilibrium distribution would be reached well before 25 years. Later Monte Carlo studies, which will be reported elsewhere, showed that equilibrium is essentially reached within 8 to 10 years.

The population vector was initialized using our estimate of the 1976 population, the last year for which we had a complete estimate. We did not consider discounting of future catches, though discounting as well as a different horizon and many other changes could have easily been incorporated into our methodology.

One of the most significant features of menhaden population dynamics is a very large year to year variability in recruitment, even when egg production is nearly constant (Nelson et al., 1977). For this reason, deterministic models seemed inadequate for modeling the menhaden population. The spawner-recruit relationship in our model is based upon the Beverton-Holt (1957) model:

$$R_i = 1/(\alpha + \beta/E_i) \tag{2.2}$$

where E_i is egg production by the spawning stock during the winter of the ith year, and R_i is recruitment into the stock during the following fall. We estimated α and β by nonlinear least squares applied to equation (2.2) and linear least squares applied to the transformed equation:

$$1/R_i = \alpha + \beta/E_i \ . \tag{2.3}$$

The residuals from (2.3) appeared to be reasonably homoscedastic and close to normally distributed, whereas those from (2.2) exhibited marked heteroscedasticity and skewness.

Notice that (2.2) and (2.3) are the same deterministic model, but adding

homoscedastic, normal errors to (2.3) results in a very different model than adding such errors to equation (2.2). Equation (2.3), of course, is obtained by applying the inverse transformation simultaneously to the reponse, R_i, and the regression model, $1/(\alpha + \beta/E_i)$. Motivated in part by this particular example, Carroll and Ruppert (1982) have made a theoretical study of transforming nonlinear regression models by applying the same power function simultaneously to both the response and the model. As we see here, such a transformation allows the statistician to keep a chosen model for the median response, e.g. in this case the Beverton-Holt model, while changing the form of the residuals. Using this analysis, we chose as a model

$$R_i = 1/(\alpha + \beta/E_i + \sigma\,\varepsilon_i) \tag{2.4}$$

where the ε are i.i.d. N(0,1) and σ was estimated from the linear fit to (2.3), except that a slight modification of (2.4) was needed because the unbounded support of the normal distribution. This makes extremely large or negative recruitment possible, though unlikely, if equation (2.4) is used as is. Let R_{max} denote the largest menhaden recruitment that has been observed in the past. When the denominator in (2.4) was above $1/R_{max}$ we did not change it, but otherwise the denominator was modified so that it did not fall below $1/2R_{max}$. The choice of $1/2R_{max}$ is, of course, somewhat arbitrary. The modification was done in a smooth manner, not by an abrupt, nondifferentiable truncation. As discussed more fully later, the differentiability properties of such modifications have important implications for our stochastic approximation methodology, which utilizes first and second derivatives.

No component of our model other than the spawner-recruit relationship was stochastic, because the available data did not suggest that other aspects of the menhaden population dynamics were highly variable.

We studied three classes of harvesting strategies: 1) constant catch policies, where the same size catch (in biomass) is taken every year,

2) constant F policies, where the instantaneous fishing mortality rate F does not vary between years, and 3) adaptive policies. The latter were motivated by escapement policies, where all biomass above a certain level (the escapement) is harvested each year. Reed (1978) showed by analytic means that constant escapement policies are optimal for a class of non age-structured, stochastic models that he considered. However, for his models constant escapement is equivalent to constant egg production, but this is not true for age structured, stochastic models where egg production per unit of biomass is low in years where immature cohorts dominate. We wanted to investigate policies which control the amount of egg production that occurs. Therefore we defined a new class of policies which we call adaptive policies. Let E_i denote the potential egg production for year i, which is defined to be the egg production that would be obtained if there were no fishing mortality, only natural mortality, on the mature stock during the year. Our model has a subroutine which determines egg production per individual of a given age as a function of juvenile population density during the individual's first nine months of life (the juvenile period). The model also stores the number of fish in each age class. Thus E_i can easily be determined, as can the actual egg production, which depends upon the amount of depletion of the spawning stock by fishing mortality. In practice, of course, E_i will be known even approximately only if the sizes of the age classes have been estimated, say by mark-recapture studies.

Let C_i be the catch for the ith year. Then an adaptive policy is one which determines C_i by the equation

$$C_i = A(E_i - B)^{\theta}. \tag{2.5}$$

Here A, B, and θ are parameters which characterize the policy.

Setting $\theta = 0$ and $B = 0$, one obtains a constant catch policy. If $\theta = 1$

and A is properly chosen, then one has a policy close to a constant egg escapement policy, i.e. a policy where egg production is held constant. The amount of escapement is determined by B. Because $\theta = 1$ policies result in large yearly fluctuations in catch, we considered also policies with θ between 0 and 1. We expected that these policies would stabilize catches. We decided to study only $\theta = \frac{1}{2}$, which we believed would be representative of the class of policies with θ between 0 and 1.

To study constant catch policies, a smoothing spline produced by SAS/Graph (Council and Helwig, 1981), was used to estimate the probability of extinction during the 25 year period as a function of the constant catch level. We found that the probability of extinction increases approximately exponentially as the constant yearly catch level increases beyond 330,000 metric tons and reaches 0.5 when the level is 390,000 metric tons. We also studied "phase in" policies where the catch is increased in each of the first three years, and then held constant for years 4 through 25. Such policies perform somewhat better than policies where catches are constant for all 25 years, yet they are not as good as constant F and adaptive policies. Both splines are shown in Figure 1.

We also used a smoothing spline to study constant F policies by estimating the expected yearly catch over the 25 year period as a function of the constant F. We calculated 95% confidence limits for expected catch and these were also plotted as smoothed splines. The splines are shown in Figure 2, where we see that F approximately equal to 0.8 maximizes the average yearly catch at about 25,000 metric tons. This shows that constant F policies offer a considerable improvement over constant catch policies.

Incidentally, neither of the splines mentioned above looked similar to a low degree polynomial. In Figure 2, we see that the second spline increases rapidly as F increases from 0 to 0.8 and then decreases slowly as F increases further. This should serve as a warning against the use here of traditional,

polynomial based, nonsequential response surface methodology.

By "optimal" policy within a class, we mean a policy in the class which maximizes the expected total utility, i.e., the expectation of (2.1). We decided to contrast the adaptive policy optimal within the class with $\theta = \frac{1}{2}$ and the adaptive policy optimal within the class with $\theta = 1$. We could have allowed A, B, and θ to vary freely and found the policy optimal within this wider class, but we decided that perhaps more information could be gained by contrasting two different values of θ.

When studying the constant catch and constant F policies, we could afford to estimate the entire one-dimensional response surface. This required 200 simulations (each for 25 years) on each point of grids with 20 and 70 points, respectively. As we will see, the adaptive policies with their greater flexibility outperform the constant catch and constant F policies, but this greater flexibility requires two parameters, and with two parameters the study of the whole response surface is prohibitively expensive. What is needed here is sequential experimentation, where as knowledge accumulates about the location where the response surface is maximized, experimentation is concentrated about this point. Stochastic approximation is an extensively studied method of sequential experimentation and seems well suited for optimization of the adaptive policies.

3. STOCHASTIC APPROXIMATION

The basic ideas behind stochastic approximation are simple and the methods have potentially wide application, yet stochastic approximation has been little used. Mead and Pike (1975) believe that the lack of application has been caused by the focusing of the stochastic approximation literature to an audience primarily of mathematicians. Here we attempt a brief introduction for possibly less mathematically interested readers.

There are two basic types of stochastic approximation methods, the Robbins-Monro and the Kiefer-Wolfowitz procedures. The Kiefer-Wolfowitz procedure is intended for function maximization and the Robbins-Monro method for the location of the root of a function. However, for the type of simulation experiment we used, the Robbins-Monro procedure can also be used for maximization, and in fact will be more efficient than the Kiefer-Wolfowitz procedure. In order to show why the Robbins-Monro method is preferable, we will discuss both procedures.

Let us start with the Robbins-Monro (1951) procedure. Suppose we can perform experiments where we fix an independent variable x and observe a response y. Let f(x) be the expected value of y given x, assume that f is an increasing function, and suppose that we have chosen a number ρ and we must find the unknown value x* such that

$$f(x^*) = \rho \ .$$

For example, x might be the dosage of a drug which regulates blood pressure, while y is a patient's blood pressure after receiving the drug, and ρ is the level at which the patent's blood pressure should be maintained.

The Robbins-Monro procedure starts experimentation at an initial guess, x_1, of x*. By suitable choice of scale for y, we can assume that $\rho = 0$. If the response is positive (negative), then the next experiment is at x_2 which

is less (greater) than x_1. This procedure is iterated:

$$x_{n+1} = x_n - a_n y_n$$

where y_n is the response at x_n and a_n is a positive constant. If a_n does not converge to 0, then the process can oscillate indefinitely, and if

$$\sum_{n=1}^{\infty} a_n < \infty$$

then the process may slow up so quickly that x_n never reaches x^*, but otherwise x_n will converge to x^* under rather general circumstances. Starting with Chung (1954), considerable research, which is mostly of a highly technical nature, considers the choice of a_n and the asymptotic distribution of $(x_n - x^*)$. Later we will mention the results which are relevant to our application.

Suppose now that instead of being increasing, f is a unimodal function and x^* is its maximizer. Kiefer and Wolfowitz (1952) developed a method of locating x^*. Again let x_1 be an initial guess of x^*. Then choose a positive constant c_1 and perform experiments at $x_1 + c_1$ and $x_1 - c_1$ with responses $y_{1,1}$ and $y_{1,2}$, respectively. Then estimate $f'(x_1) = (d/dx)f(x_1)$ by

$$(y_{1,1} - y_{1,2})/(2c_1) .$$

If y_1 is positive (negative) then let x_2 be greater (less) than x_1. Continuing in this manner, let the process x_n be obtained by iteration of

$$x_{n+1} = x_n - (a_n/2c_n)(y_{n,1} - y_{n,2})$$

where $y_{n,1}$ and $y_{n,2}$ are responses at $x_n + c_n$ and $x_n - c_n$, respectively. The choice of c_n presents a dilemma. As c_n increases the bias

$$[f(x_n + c_n) - f(x_n - c_n)]/(2c_n) - f'(x_n)$$

increases, but the variance is proportional to c_n^{-2}. The best choice of c_n has been studied, but because of the need to compromise between bias and variance, the best rate of convergence of the Kiefer-Wolfowitz process is less than that

of the Robbins-Monro process. Fortunately, as will be discussed further, for simulation experiments, one can, by using the variance-reduction technique of common random number streams, maximize a regression function by applying the Robbins-Monro process directly to its derivative.

The Robbins-Monro and Kiefer-Wolfowitz processes have been extended (Blum, 1954) to the situations, respectively where x and y are k-variate and where x is k-variate and y is real-valued.

4. THE OPTIMIZATION METHODOLOGY

We have developed the means for tackling the following general problem. Let x be a k-variate parameter, let e be an m-variate random vector, and let f(x,e) be a real valued function of x and e. Define

$$f(x) = E\, f(x,e)$$

and suppose we wish to find a maximizer of f.

Throughout, the discussion will be confined to our specific application to adaptive policies. For fixed θ we want to find the values of A and B which maximize the expected utility. In this example, $e = (\varepsilon_1, \ldots, \varepsilon_{25})$ is the vector of 25 independent N(0,1) random variables which appear in equation (2.4) and $x = (A,B)$. To explicitly show dependence upon the control vector x and the random vector e, we let $C_i(x,e)$ represent the catch in year i. Then the function f is

$$f(x,e) = \sum_{i=1}^{25} U(C_i(x,e))$$

which, of course, depends on U and θ but these are fixed.

For any value of x we can, by simulation of the fishery, obtain a response y with the distribution of f(x,e). Therefore, we could apply Blum's (1954) multivariate extension of the Kiefer-Wolfowitz process. Let D(x,e) and D(x) be the gradients of f(x,e) and f(x), respectively. By experimenting at (A + c, B), (A - c, B), (A, B + c), and (A, B - c), using four independent replicates of e, to obtain respectively y_1, y_2, y_3, and y_4, we could estimate D(x) by

$$1/2c(y_1 - y_2,\ y_3 - y_4)' \quad . \tag{4.1}$$

Needless to say, the choice of c would involve the compromise between bias and variance that we encountered before.

However, we did not use this method of estimating D(x). Instead, we performed the experiments at (A + c, B), (A - c, B), (A, B + c), and (A, B - c)

using the *same* replicate of e and a very small value of c. Then (4.1) becomes D(x,e) obtained by numerical differentiation, and we can apply the Robbins-Monro procedure directly to D(x) provided

(4.2) $$E(D(x,e)) = D(x) ,$$

i.e., the interchange of differentiation and expectation is permissible.

Variance-reduction by the use of common random numbers is a standard technique of Monte Carlo simulations. Schruben and Margolin (1978) have explored the use of common, as well as antithetic, random number streams for factorial simulation experiments. However, there appears to be no general theory of such variance-reduction techniques and new, specific instances should be of interest.

If we had moved from x_n, our nth estimate of x*, to x_{n+1} along the estimated gradient direction, then our technique would have been similar to the steepest ascent method normally advocated for response surfaces. However, for deterministic functions it is well known that in general steepest ascent converges much more slowly than Newton-type methods based on the Hessian (matrix of second derivatives). See, for example, Bard (1974), page 88. Let H(x,e) and H(x) be the Hessians of f(x,e) and f(x), respectively. For fixed x and e, one can simulate at (A + c, B + c), (A - c, B + c), (A - c, B - c), (A + c, B - c), and (A,B) to obtain $y_5,\ldots, y_9$. These along with $y_1,\ldots, y_4$ are then used to calculate H(x,e) numerically:

$$(\delta^2/\delta A^2)\ f(x,e) \doteq (y_1 + y_2 - 2y_9)/c^2 ,$$

$$(\delta^2/\delta A \delta B)\ f(x,e) \doteq (y_5 - y_6 + y_7 - y_8)/(2c)^2 ,$$

and

$$(\delta^2/\delta B^2)\ f(x,e) \doteq (y_3 + y_4 - 2y_9)/c^2 .$$

A and B were scaled so that simulations took place in the ranges $1 \le A \le 3$ and $-80 \le B \le 130$. After some trial and error, we decided that c = 0.01 was satisfactory. The choice of c involves a compromise between roundoff error and bias

only, and both can be made much smaller than the statistical variability. One should make c as small as possible without appreciable roundoff error.

Suppose that $e_1, e_2, \ldots$ is a sequence of i.i.d. replicates of e. Then, for positive constants k_1 and k_2 to be discussed shortly our Robbins-Monro type estimation sequence satisfied

$$x_{n+1} = x_n - (n + k_1)^{-1} H_n^{-1} D(x_n, e_n) \tag{4.3}$$

$$H_{n+1} = H_n - (n + k_2)^{-1} (H(x_n, e_n) - H_n) \tag{4.4}$$

except for one modification. H_n is an estimate of $H(x^*)$, which will be negative definite. However, H_n may be either not negative definite or negative definite but poorly conditioned, i.e., having a very large ratio between the smallest and largest eigenvalues. Let $H_{n+1}^{(ij)}$ be the i,jth element of H_{n+1} and define

$$M_{n+1} = \begin{pmatrix} -1 & H_{n+1}^{(12)} / [H_{n+1}^{(11)} H_{n+1}^{(22)}]^{1/2} \\ H_{n+1}^{(12)} / [H_{n+1}^{(11)} H_{n+1}^{(22)}]^{1/2} & -1 \end{pmatrix}$$

M_{n+1} is just a rescaling of H_{n+1}. Let t_{n+1} be the largest eigenvalue of M_{n+1}. If $t_{n+1} > -0.01$, then we redefined H_{n+1} to equal H_n. Therefore, if we started with an H_1 such that $t_1 \leq -0.01$, then $t_n \leq -0.01$ for all n. This modification was suggested by the directional discrimination method given by Bard (1974) for deterministic optimization, but our method is considerably different than Bard's, because unlike users of deterministic methods, we needed to average the Hessian, which in our case is random, over many iterations to obtain a good estimate of $H(x^*)$.

The initial values of x_1 and H_1 were found by simulating the fishery on a rectangular grid of values of A and B. The location of the grid's center was based upon a calculation from a deterministic model, which was similar to the stochastic model, but analytically tractable. The size of the grid was

determined by experience gained while testing and debugging the program. The same fifty replicates of e were used at each point on the grid. Then a quadratic polynomial in A and B was fit to the data, H_1 came from the quadratic terms of the fitted polynomial, and x_1 was the maximizer of the polynomial.

The constants k_1 and k_2 should depend upon the accuracy of x_1 and H_1 as estimates of x^* and $H(x^*)$. We used $k_1 = k_2 = 50$, since 50 is the number of replicates of e used on the grid.

Let S be the variance-covariance matrix of $D(x^*,e)$. Then, since $D(x) = E(D(x^*,e)) = 0$ we have that

$$S = E\ D(x^*,e)\ D(x^*,e)' \ .$$

Using techniques from theoretical studies of stochastic approximation, e.g. Fabian (1971), in particular his Theorem 2.7 which applies to Kiefer-Wolfowitz processes similar to our Robbins-Monro process, one can formulate conditions sufficient that $x_n \to x^*$ and $H_n \to H(x^*)$ almost surely and

$$n^{1/2}(x_n - x^*) \xrightarrow{L} N(0, H(x^*)^{-1}\ S\ H(x^*)^{-1}) \ .$$

We will not pursue asymptotic theory here except to show how a stopping rule can be obtained. This rule may be new but is closely related to one of Sielken (1973). By a Taylor expansion

$$f(x_n) - f(x^*) \doteq 1/2\ (x_n - x^*)'\ H(x^*)(x_n - x^*) \tag{4.6}$$

since $D(x^*) = 0$. By (4.5) and (4.6), $[f(x_n) - f(x^*)]$ converges in distribution to a random variable with expectation

$$E[-\ (1/2)z'H(x^*)z] = -(1/2n)\ E\ (trHzz') = -(1/2n)\ tr[SH(x^*)^{-1}]$$

where $z \sim N(0,\ H(x^*)^{-1}\ S\ H(x^*)^{-1})$. Now $[f(x^*) - f(x_n)]$, of course, should be made small and we may wish to stop when the expected relative error, $[f(x^*) - E\ f(x_n)]/f(x_n)$ is smaller than some predetermined positive constant Δ. Then one can estimate $f(x^*)$ and S by f_n and S_n defined recursively by

$$f_{n+1} = f_n - (n + k_3)^{-1} [f(x_n,e_n) - f_n] \tag{4.7}$$

$$S_{n+1} = S_n - (n + k_4)^{-1} [D(x_n,e_n)D(x_n,e_n)' - S_n] \tag{4.8}$$

where k_3 and k_4 are chosen with the same considerations as k_1 and k_2 and f_1 and S_1 are estimated from the data generating x_1 and H_1. Then, one stops at the first integer N such that

$$-\mathrm{tr}(S_N H_N^{-1})/(Nf_N) < \Delta . \tag{4.9}$$

We didn't use f_n as given by equation (4.7). From prior experience with the model, we had a reasonably good estimate of the maximize of f(x), and we set f_n constantly equal to this estimate. Since f_n is only used to estimate relative error, we saw no need for a more accurate estimate. The sequence (4.7) is suggested for those users without a good prior estimate. We chose to use $\Delta = 0.001$, which may seem stringent but experimentation was inexpensive for our model. We found that H_n and S_n were unstable when n was small, say $n < 100$. Therefore, we required both $N \geq 100$ and (4.9) to hold. Our printout included x_n for $n = 1,\ldots, N$, S_1, H_1, S_N, H_N, and

$$N^{-1}\sum_1^N D(x_n,e_n) . \tag{4.10}$$

We did not rely entirely upon this stopping rule since it can only be justified asymptotically. If we felt that x_n was still changing systematically when n was close to N, S_1 and S_N or H_1 and H_N were rather dissimilar, or either coordinate of (4.10) was large, then we performed a second set of iterations. The iteration of (4.3), (4.4), (4.7) and (4.8) and the checking of the convergence criterion was done automatically by a loop in our computer program. Since this meant that inspection was not possible at each iteration, we only allowed a maximum of 200 iterations, so that money would not be wasted in the event of a programming error or other mishap. If the limit of 200 was reached before the stopping criterion was met, then also a second set of iterations was

performed. For the second set of iterations, x_1, H_1, S_1, and f_1 were x_N, H_N, S_N, and f_N from the first simulation. We generally obtained new values of k_1, k_2, k_3, and k_4 by adding N to the previous values. However, we sometimes used less than N when the sequences H_n or S_n seemed particularly unstable during the first set of iterations, or we were dissatisfied for some other reason. We did not modify the stopping criterion (4.7) but did require the rerun to perform at least 100 steps before stopping.

Perhaps we should summarize how the stochastic approximation algorithm used here differs from the Newton-Raphson or similar algorithms, which would be appropriate for a deterministic model. First, the factor $(n + k_1)^{-1}$ in equation (4.3) is replaced by 1 in deterministic algorithms. In deterministic situations, the gradient $D(x_n)$ converges to 0 as x_n converges to the optimum and therefore the steps sizes decrease to 0 even without the factor $(n + k_1)^{-1}$. In stochastic situations, a damping factor which converges to 0 is needed since the estimated gradient $D(x_n,e_n)$ is random and does not converge to 0 even if x_n approaches the optimum. Theory shows that $(n + k_1)^{-1}$ is the optimal damping factor. Second, equations (4.4), (4.7), and (4.8) estimate $H(x^*)$, $f(x^*)$, and S respectively, by average across all prior iterations. This averaging is unnecessary in deterministic situations where $H(x^*)$, $f(x^*)$, and S are best estimated using only the most recent iteration which is presumably close to the optimum and (unlike in stochastic settings) has no random variations. Third, the stopping rule here is based upon the random variability in the estimate of $D(x^*)$, not (as for deterministic algorithms) upon the size of the gradient at x_n or the difference between x_n and x_{n-1}.

5. RESULTS

In general, we were very satisfied with our experiences using stochastic approximation, but we learned enough that future experiences should be even more successful. There were four adaptive policies, since we used two utility functions and two values of θ. All the adaptive policies had average yearly catches around 525,000 metric tons, which shows that the adaptive policies offer considerable improvement over the best constant F policy. Other implications of the results for harvesting of menhaden will be discussed in the fisheries literature. Here we will try to convey the information relevant to future users of stochastic approximation.

Generally, during the first 30 or so iterations, the sequence x_n moved rapidly to a neighborhood of the point of maximization. Then the sequence x_n slowly converged in an oscillatory manner as the statistically variability became large relative to $D(x_n)$. Often several hundred iterations were necessary before the convergence criterion was met. The cost of experimentation was about $5.00 for 200 iterations, each performing a complete 25 year simulation at each point on the 9 point grid used for numerical differentiation. This sum was quite small compared to the cost of the data analysis needed to develop the model. Also our stringent stopping rule seemed justified, since strong consideration is being given to regulation of the menhaden fishery and our research may have important policy implications.

Part of our reason for using so many iterations was a desire to see how stable the algorithm actually was. It was quite comforting, when it appeared that the maximizer had been found, that the sequence x_n remained quite constant during further iteration. Although asymptotic theory provided us with some confidence in our methodology, there was always the question of whether the asymptotic results were relevant for the sample sizes we were considering.

To illustrate the behavior of our stochastic approximation method, we will give an example. Our model, includes natural mortality which is assumed to occur at a constant instantaneous rate, m. Our estimate, m = 0.502 on a yearly basis, came from the analysis of mark-recapture data. To study the effects of possible estimation error, we duplicated our simulation analysis with m = 0.319, θ = 1, and U(x) = x. The initial values where A_1 = 2.36 and B_1 = 77.6, which were the previously determined optimal values of A and B when m = 0.512, θ = 1, and U(x) = x. The convergence criterion (4.9) was first met at N = 155.

Figures 3 and 4 are plots of A_n and B_n, respectively, against n for n = 1 to 155. Notice the large overshoots on the first two iterations. After these A_n moves quickly to the optimum, A = 1.11, but B_n moves away from the optimum, B = 63.5. We believe that a somewhat poor initial value of H accounts for this behavior of B_n. When the estimated partial derivative with respect to A is negative with a large absolute value, then too large an absolute value of H_n^{12} (which is negative) can force B_n to increase even when the estimated partial derivative with respect to B_n is also negative. By n = 10, both A_n and B_n are decreasing toward the optimal values. In Figures 5 and 6, we show A_n and B_n plotted against n for n = 105 to 155 only. This allows an expanded vertical scale which illustrates the oscillatory behavior of A_n and B_n. Figure 7 plots the average 25-year catch at (A_n,B_n) versus n. The catches for each n depend, of course, upon 25 random recruitments, and Figure 7 shows great variability among the catches. The fact that A_n and B_n converge rapidly to their optimal values despite this variability demonstrates the power of the variance-reduction technique we employed.

Because B_n appeared to still be decreasing at n = 155, we performed another set of iterations. The final estimates from the first set of iterations were A = 1.11 and B = 73.0. The final estimates from the second set were A = 1.11 and B = 63.5. In terms of expected catch, the performance of A = 1.11 and

B = 73.0 is very similar to A = 1.11 and B = 63.5, and we now believe that the second set of iterations was not necessary. Looking at Figure 7, one can see that the catches increased during the first 10 to 20 iterations, but after that point no pattern among the catches is visually discernible amidst the variability. The reader may have noticed that compared to previously mentioned average catches, even those for adaptive policies, almost all catches in Figure 7 are large. This is simply because m is smaller here.

There was also some doubt whether all assumptions of the asymptotic theory were actually met. In particular, we have mentioned that (4.2) must hold. A similar interchange of differentiation and integration must be assumed valid in order to prove that $H_n \to H(x^*)$. At first we believed that these were purely mathematical assumptions with no real ramifications in practice. But, as we will now describe, this was not the case.

Initially, the functions (one for each class) expressing the effects upon size at a given age of density dependent juvenile growth were step functions, each with two jumps. This was because we had classified juvenile density into three categories and had modeled the effects as constant within each category. With the model in this form the algorithm was very unstable. If for fixed x_n and e_n, any of the 25 cohorts generated during the simulations of spawning had a density near one of the two boundaries between density categories, then the numerically obtained derivatives could be enormous. On some points of the grid used for differentiation the cohort would be on one side of the boundary while for other points it would be on the other side. We had hoped that such occurrences would be sufficiently rare to be insignificant. As it turned out, a single occurrence could push x_n into a region of infeasible values of x, where the population is driven to extinction or the harvesting policy makes no sense (e.g., if $A < 0$).

After removing the jump discontinuities, we still found the sequence H_n poorly behaved. A frequent occurrence was for all entries of H_n to increase a hundred or even a thousand fold on one iteration. Then, as would be expected from equation (4.4), since H_n^{-1} was extremely small, x_n would hardly change even if it was quite far from the maximizer. This problem was corrected by changing all functions in the model, for example those determining density dependent effects upon size and the one modifying equation (2.4) as described in section 2, so that they had continuous first derivatives and were piecewise twice differentiable.

After these modifications, our experience with stochastic approximation was very satisfactory. One real advantage of stochastic approximation was that once the maximizer for a given model and utility function was determined, it was a good starting value when the model was modified or a new utility was used.

Our project involved a major investigation of the risks associated with estimation errors for fundamental biological parameters, such as the natural mortality rate and the parameters α, β, and σ in equation (2.4). This required us to find not only the harvesting policy optimal if the biological parameters were equal to our estimates, but also the policies optimal for various perturbations of the parameters about our estimates. With stochastic approximation, this risk analysis was very quick and inexpensive.

We have two suggestions for further work. As we have mentioned, we used 9 simulations for numerical differentiation at each iteration. Only 6 are necessary. Initially we tried a regular pentagonal design with a center point, a design we took from Box and Wilson (1951). We first thought that the design was causing considerable roundoff errors, but we later discovered that the errors were due to a programming mistake in calculating the derivatives. Also, when experimenting with the pentagonal design, we were still using step

functions to model the effects of density dependent juvenile growth. At the point when we realized that the pentagonal design might produce satisfactory numerical derivatives, it seemed better to proceed with the design we were using. It seems worthwhile in the future to try designs which use less simulations, especially if x is more than two dimensional.

Our stopping rule is based upon a number of large sample approximations. Monte Carlo studies of this stopping rule should be performed. At first, we did not have a great deal of confidence in the stopping rule, and we tended to continue iteration even after the stopping criterion had been met. After this experience, we now feel that the stopping rule may have been working better than we thought.

As in most other areas of statistics, there is no substitute for experience when applying stochastic approximation. We spent several months (on a part-time basis) working with the algorithm before we were able to use it well, but we feel that we gained valuable skills in the process. We hope that the present paper will help future users of stochastic approximation learn these skills with less time and effort than we required.

REFERENCES

Bard, Y. (1974). *Nonlinear Parameter Estimation.* New York: Academic Press.

Beverton, R.J.H. and Holt, S.J. (1957). *On the Dynamics of Exploited Fish Populations.* London: Her Majesty's Stationery Office.

Blum, J. (1954). Multidimensional stochastic approximation methods. *Annals of Mathematical Statistics. 25,* 737-744.

Box, G.E.P. and Wilson, K.B. (1951). On the experimental attainment of optimum conditions. *Journal of the Royal Statistical Society, Series B,* 1-45.

Carroll, R.J. and Ruppert, D. (1982). Power transformations when fitting theoretical models to data. *Institute of Statistics Mimeo Series, #1383.* Chapel Hill, NC: Department of Statistics, University of North Carolina.

Council, Kathryn A. and Helwig, Jane T., ed. (1981). *SAS/Graph User's Guide.* Cary, NC: SAS Institute.

Chung, K.L. (1954). On a stochastic approximation method. *Annals of Mathematical Statistics. 25,* 463-483.

Fabian, V. (1971). Stochastic approximation. In *Optimizing Methods in Statistics,* J.S. Rustagi, ed. New York: Academic Press.

Kiefer, J. and Wolfowitz, J. (1952). Stochastic estimation of the maximum of a regression function. *Annals of Mathematical Statistics, 23,* 462-466.

Ludwig, D. and Walters, C.J. (1981). Measurement errors and uncertainty in parameter estimates for stock and recruitment. *Canadian Journal of Fisheries and Aquatic Sciences. 38:* 711-720.

Mead, R. and Pike, D.J. (1975). A review of response surface methodology from a biometric viewpoint. *Biometrics, 31,* 803-851.

Mendelssohn, Roy. (1978). Optimal harvesting strategies for stochastic single-species, multiage class models. *Mathematical Biosciences, 41.* 159-174.

Nelson, W.R., Ingham, M.C., and Schaaf, W.E. (1977). Larval transport and year class strengths of Atlantic menhaden, *Brevoortia tyrannus. Fisheries Bulletin, 75.* 23-41.

Raiffa, H. (1968). *Decision Analysis. Introductory Lectures on Choices Under Uncertainty.* Reading, MA: Addison-Wesley.

Reed, W.J. (1980). Optimal age-specific harvesting in a nonlinear population model. *Biometrics, 36.* 579-593.

Robbins, H. and Monro, S. (1951). A stochastic approximation method. *Annals of Mathematical Statistics, 23,* 400-407.

Ruppert, David. (1982). Almost sure approximations to the Robbins-Monro and Kiefer-Wolfowitz processes with dependent noise. *Annals of Probability, 10.* 178-187.

Schruben, L.W. and Margolin, B.H. (1978). Pseudorandom number assignment in statistically designed simulation and distribution sampling experiments. *Journal of the American Statistical Association, 73,* 504-519.

Sielken, R.L., Jr. (1973). Stopping times for stochastic approximation procedures. *Zeitschrift fur Wahrscheinlichkeitstheorie und verwandte Gebiete, 26,* 67-75.

LIST OF FIGURES

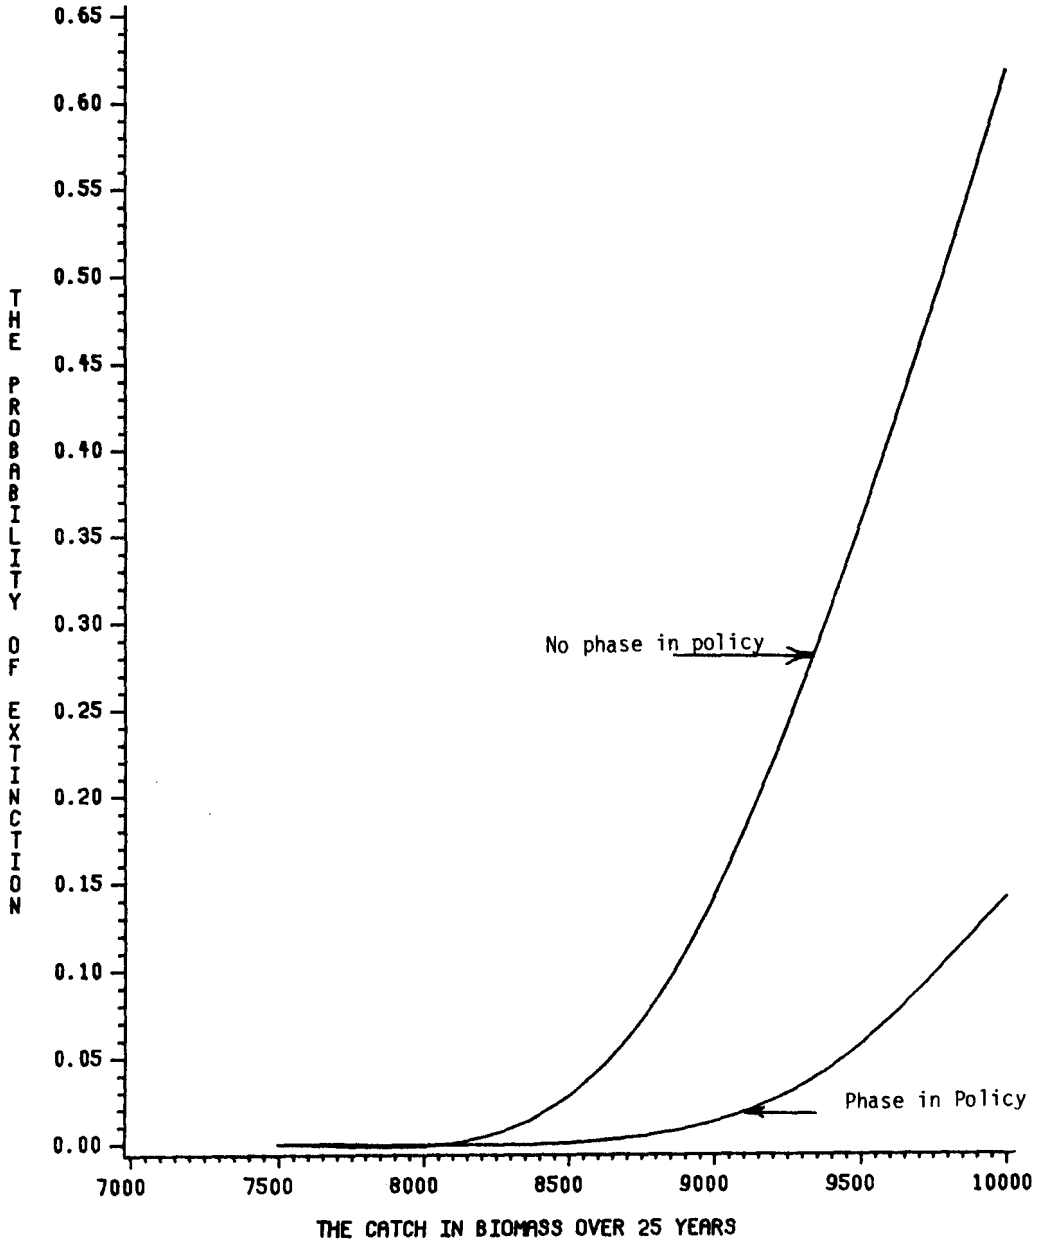
0.65
0.60
0.55
0.50
0.45
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
THE PROBABILITY OF EXTINCTION
No phase in policy
Phase in Policy
7000
7500
8000
8500
9000
9500
10000
THE CATCH IN BIOMASS OVER 25 YEARS

Figure 1

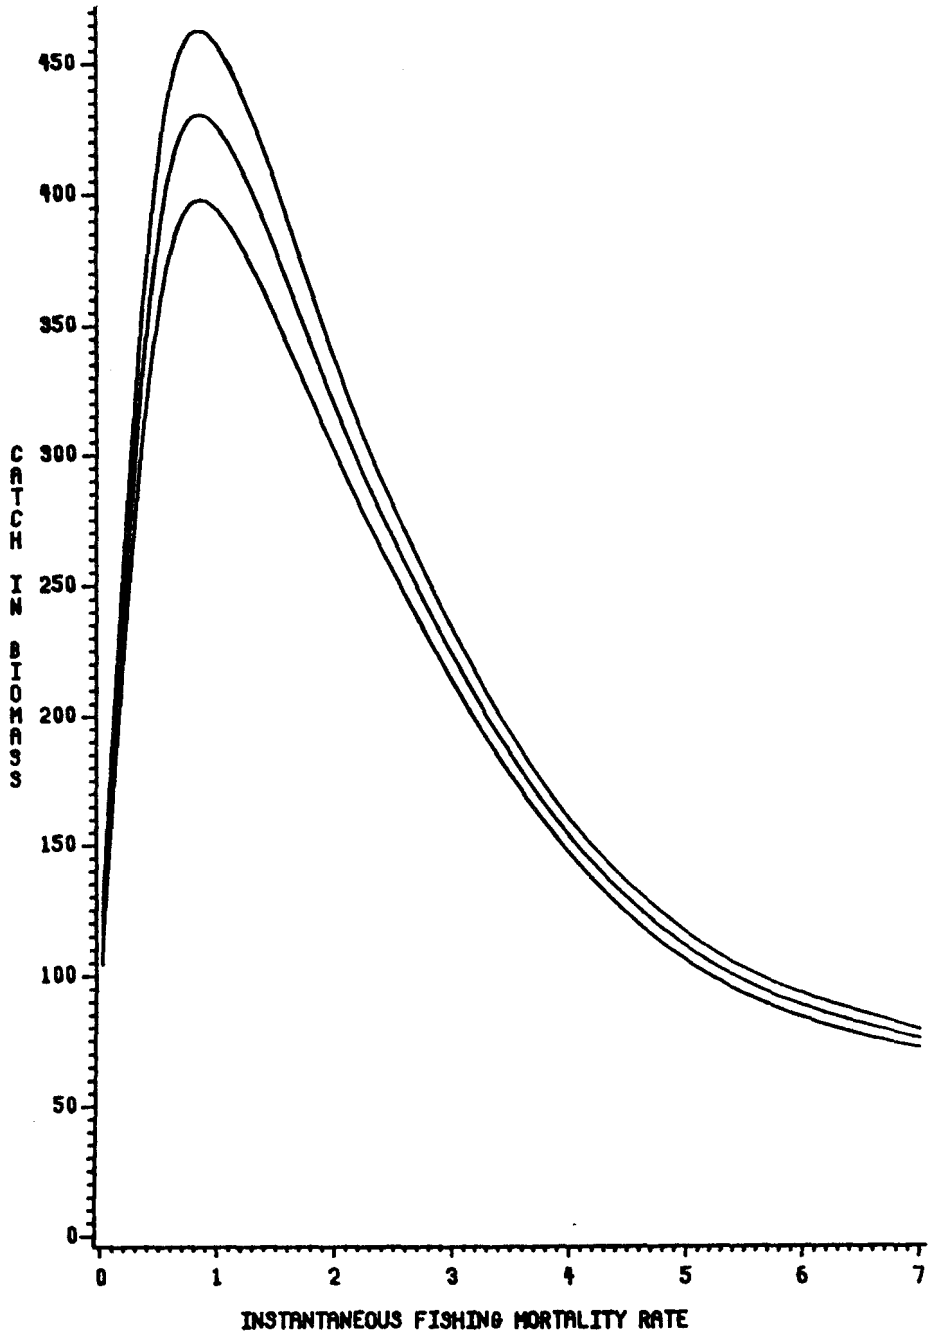
450
400
350
300
250
200
150
100
50
0
CATCH IN BIOMASS
0
1
2
3
4
5
6
7
INSTANTANEOUS FISHING MORTALITY RATE

Figure 2

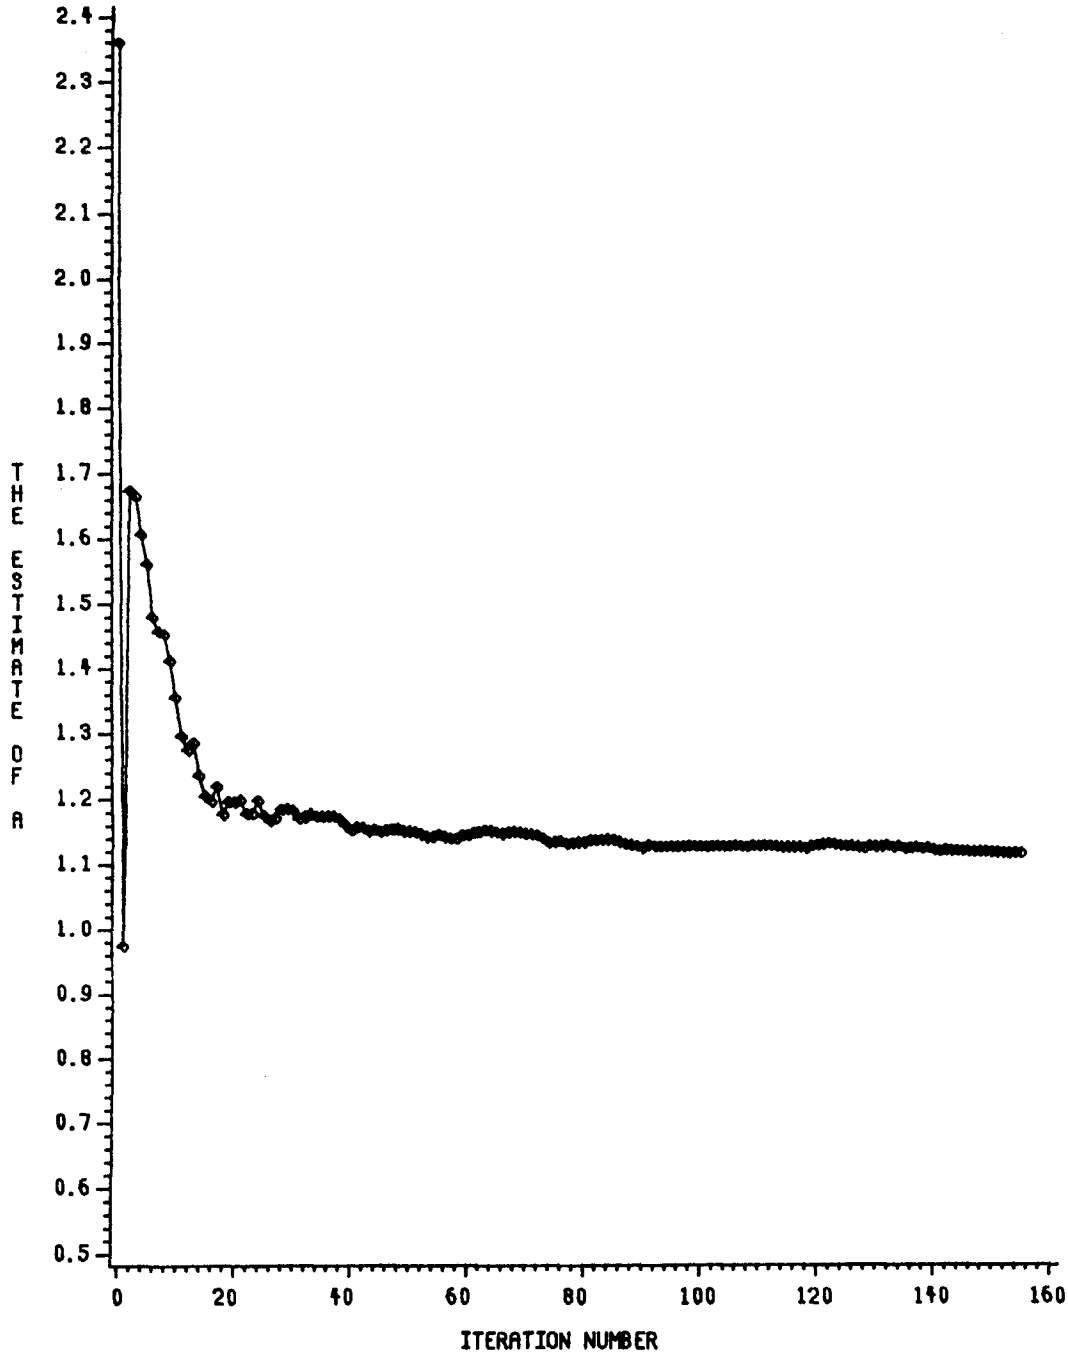
THE ESTIMATE OF A
2.4
2.3
2.2
2.1
2.0
1.9
1.8
1.7
1.6
1.5
1.4
1.3
1.2
1.1
1.0
0.9
0.8
0.7
0.6
0.5
0
20
40
60
80
100
120
140
160
ITERATION NUMBER

Figure 3

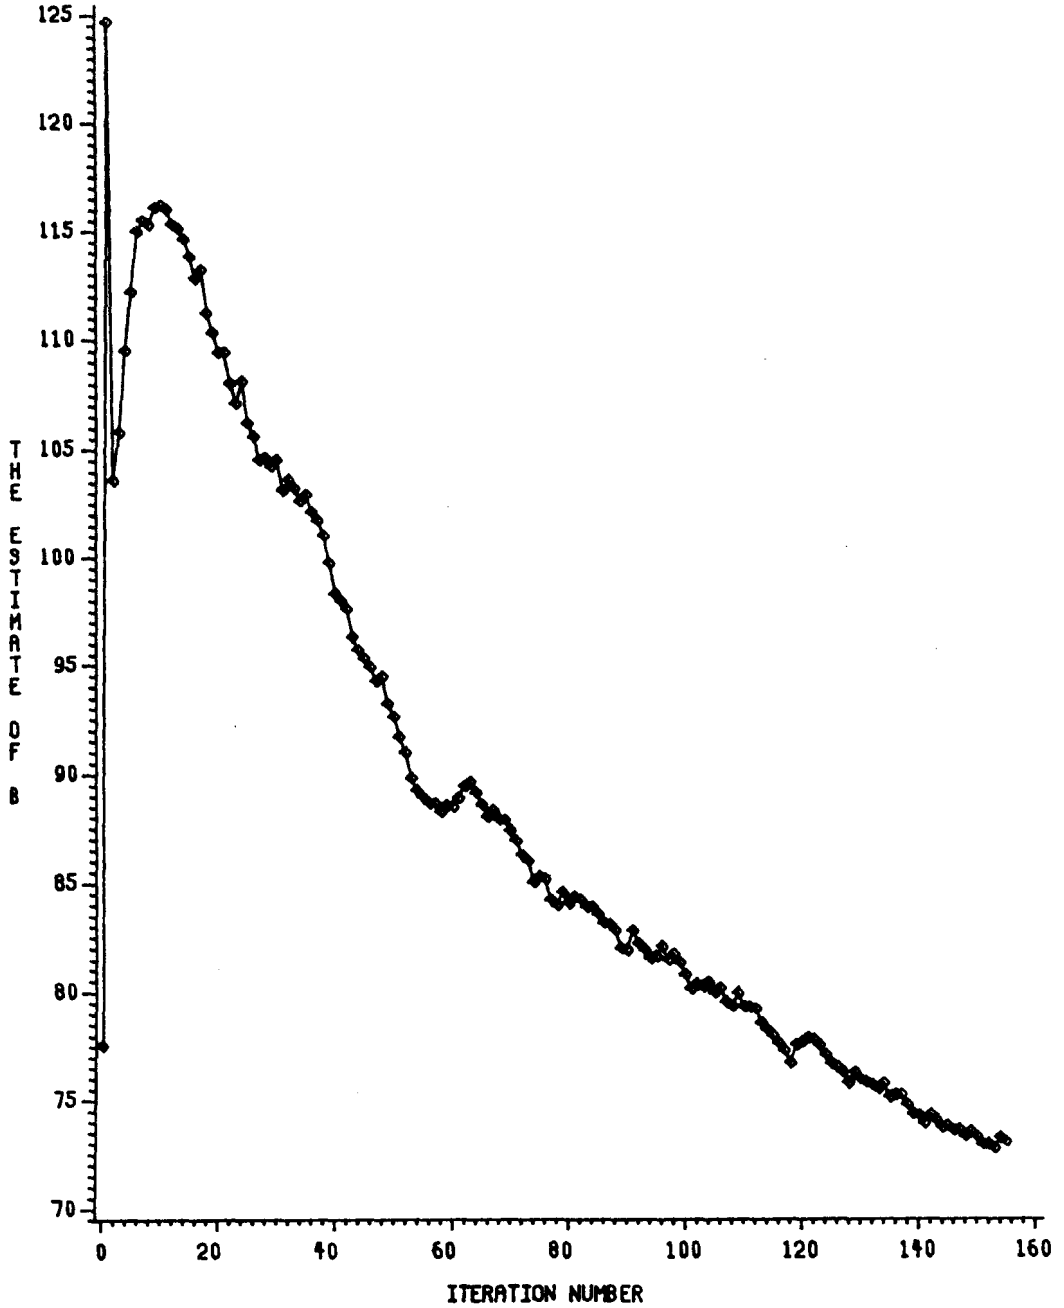
125
120
115
110
105
100
95
90
85
80
75
70
THE ESTIMATE OF B
0
20
40
60
80
100
120
140
160
ITERATION NUMBER

Figure 4

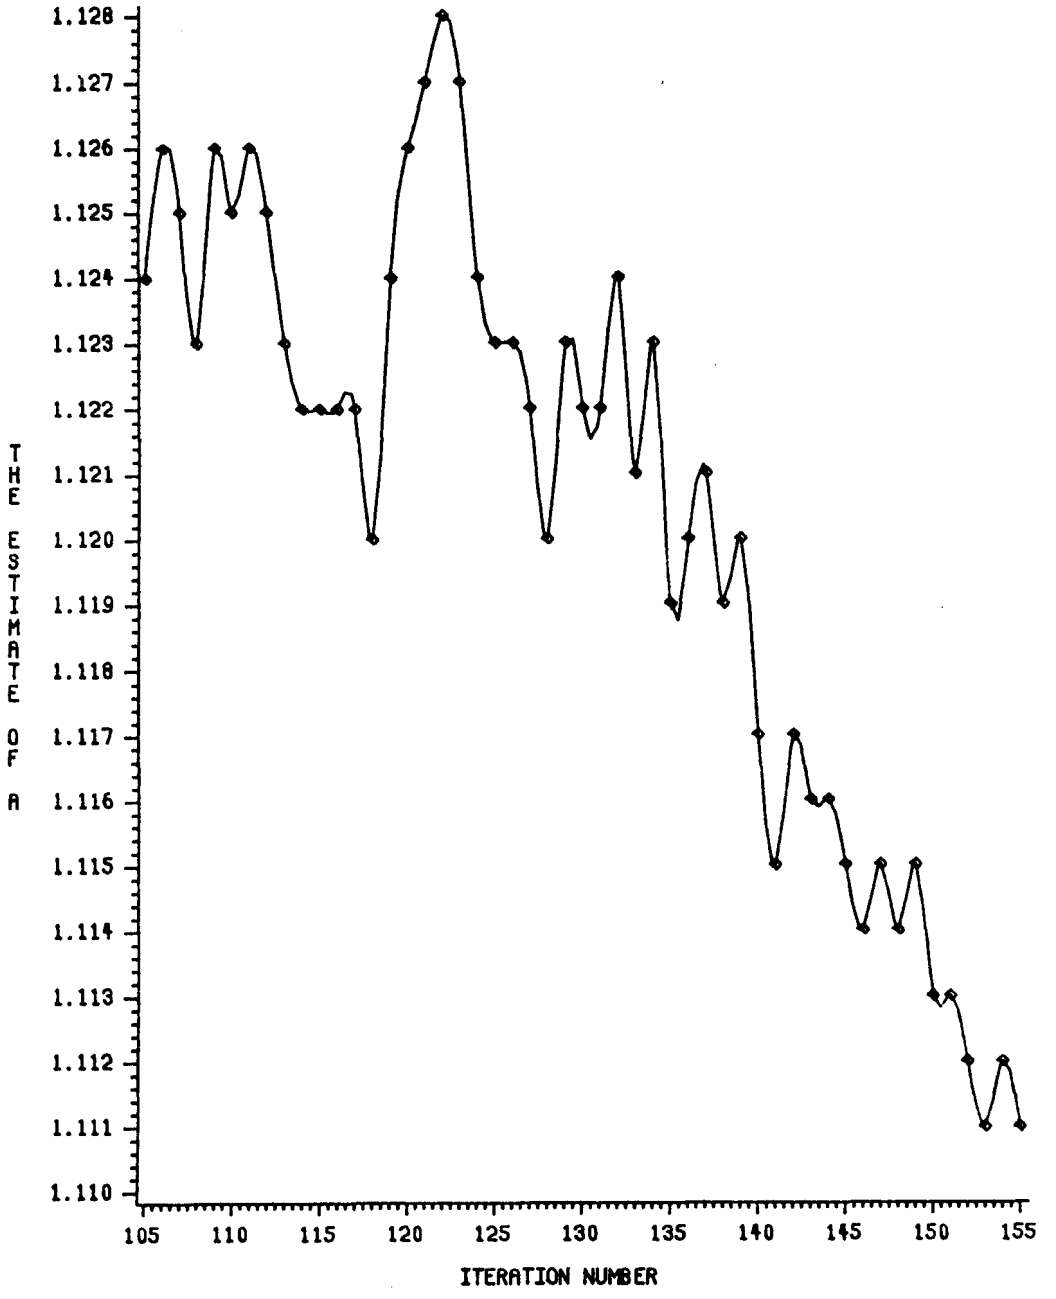
1.128
1.127
1.126
1.125
1.124
1.123
1.122
1.121
1.120
1.119
1.118
1.117
1.116
1.115
1.114
1.113
1.112
1.111
1.110
THE ESTIMATE OF A
105
110
115
120
125
130
135
140
145
150
155
ITERATION NUMBER

Figure 5

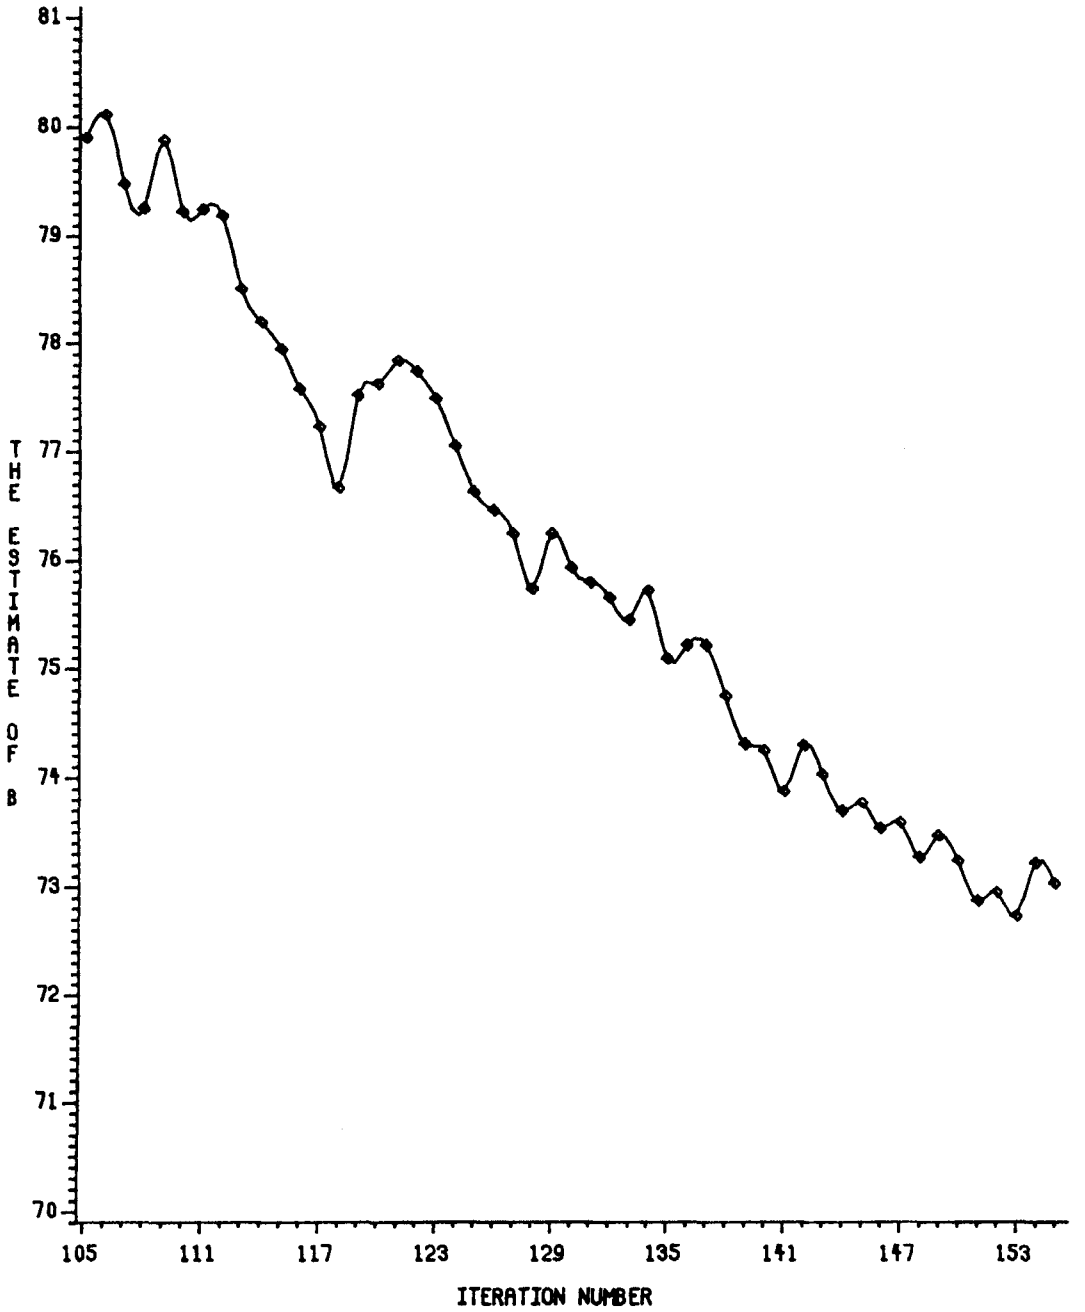
81
80
79
78
77
76
75
74
73
72
71
70
THE ESTIMATE OF B
105
111
117
123
129
135
141
147
153
ITERATION NUMBER

Figure 6

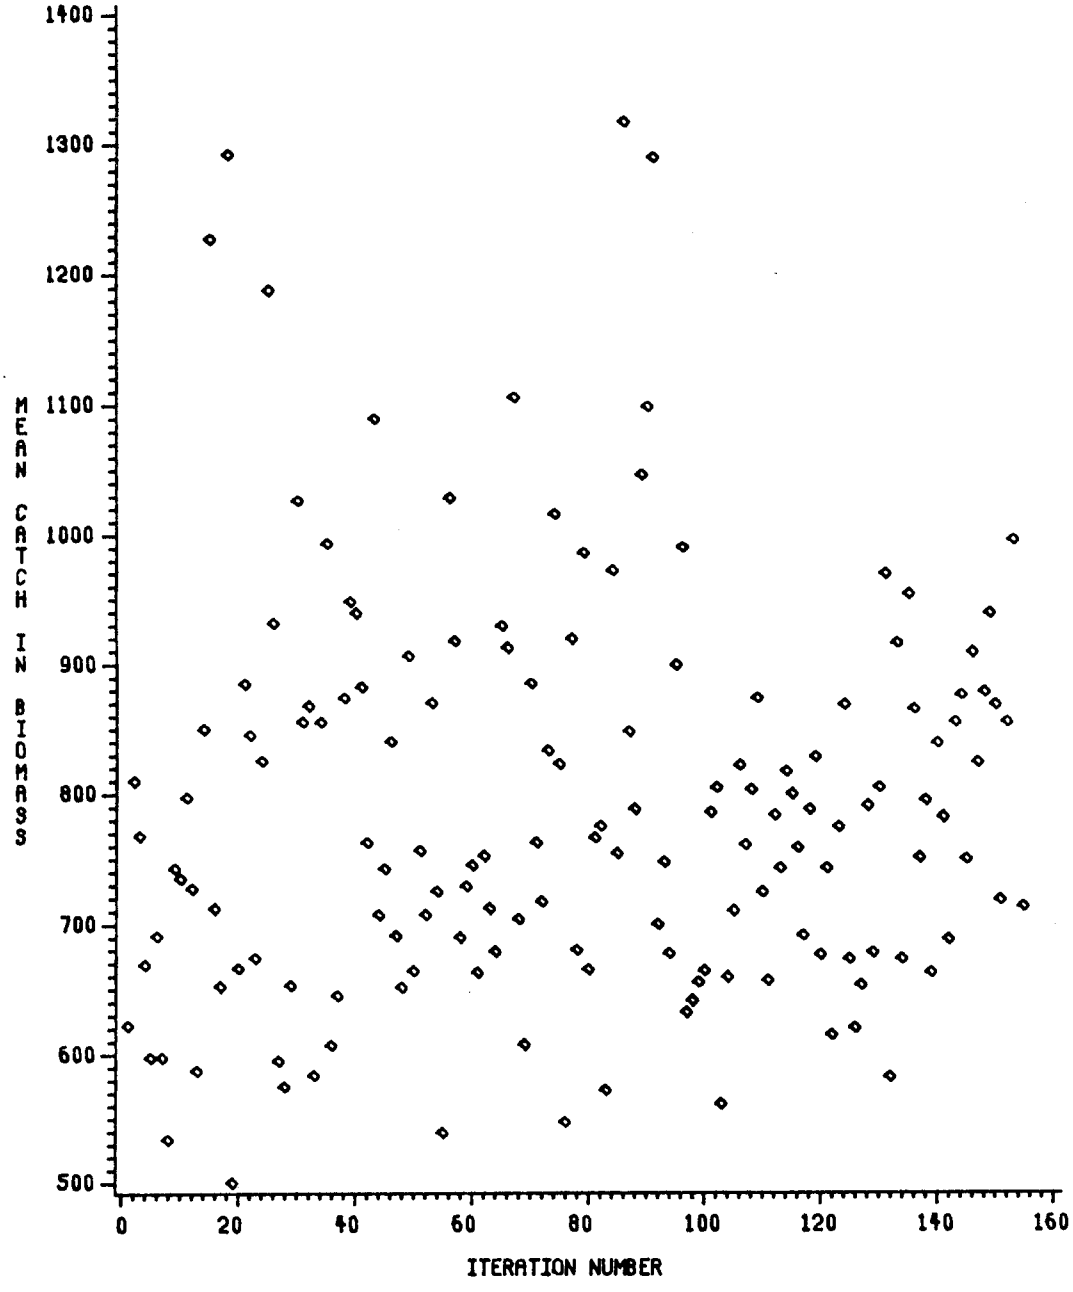
1400
1300
1200
1100
1000
900
800
700
600
500
MEAN CATCH IN BIOMASS
0
20
40
60
80
100
120
140
160
ITERATION NUMBER

Figure 7